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 IPbus is going to 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 an ATLAS standard IPMC. The requirements with respect to general board control are:

- 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.
- 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.

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.
- 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 DAQ and two fibres for ROI transmission will be provided on L1Topo. The requirements with respect to DAQ and ROI interface are:

- Provide two identical optical ROI output channels to the front panel as required by the future ROI Builder.

* provide an embedded (firmware) ROD for data into DAQ and ROI

- ~~Provide an optical DAQ output channel to the front panel.~~
- Use emulated S-Link in the control FPGA.
- Connect them to MGTs on the control FPGA.
- Provide a separate ~~40.0789 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 (12 channels, bi-directional).
- Provide an electrical busy output lemo-00 connector on the front panel.

About ROD firmware development see Sect. 4.2.

2.8 Extension mezzanine

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:

- Break-out and signal conditioning for electrical trigger signals (32 lanes to CTP).
- Route TTC clock towards TTCdec module.
- Receive two 4-pair Ethernet signals from backplane zone 2 (base interface).
- 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.

* connect 8 lanes to each hub slot (fabric interface zone 2

or do we say: to the front panel?

2.9 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 (optional)
- Connect to Ethernet via the extension module (optional)
- Connect power control / alarm on the ATCA power brick
- Supply the IPMC with the management voltage generated on the power bridge

2.10 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

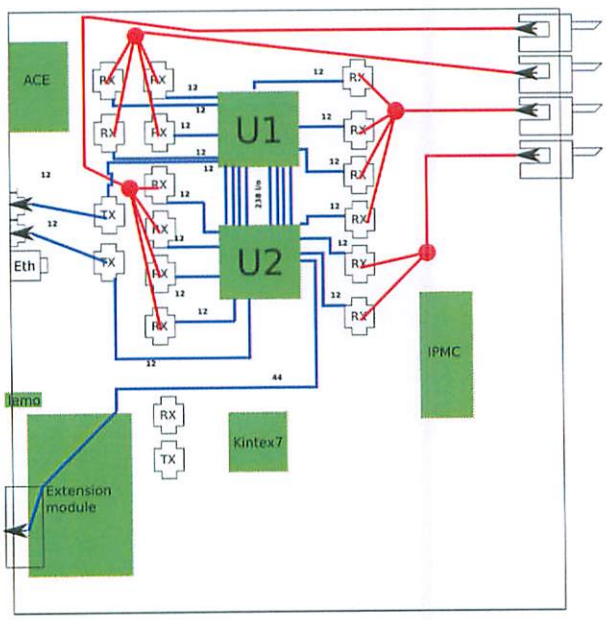

Figure 2: L1Topo real-time data path: (red) pigtails from the back panel, (blue) RTDP links.

3 Implementation

This section describes some implementation details of the L1Topo module.

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 Fig.2 .

3.1 Modular design

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.

On the back panel optical connectors, four shrouds are mounted. At the envisaged density of 48 fibres per connector a capacity of 192 fibres is provided. This is sufficient to to supply the 160 input channels (with 72-

so as to minimise ~~cable~~ fibre length and latency, a small number of input fibres (for L1Muon) is fed through the front panel.

way MTP/MPO bundles, 288 channels would be possible). The back panel connection scheme has been chosen to simplify maintenance on the module.

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. An alternative jitter cleaner device can be mounted on the extension mezzanine if needed. A detailed block diagram of the clock path is shown in Appendix C.

3.3 Control

A detailed block diagram of control paths is shown in Appendix D.L1Topo module control is done via standard Ethernet access from the front panel. 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 or the front panel, 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 [14]. The hardware provided on L1Topo is compatible to that scheme and IPbus is adopted. The specific firmware and software required has been developed.~~

The IPbus firmware/software suite is used and has been chosen for L1CALO Phase1 module control as well

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 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 40.0789 MHz. Parallel differential I/O operates at multiples up to 1Gb/s. Multi-Gigabit (MGT) links operate at 6.4. This is only possible on matched-impedance lines. Differential sink termination is used throughout. All FPGA inputs are internally terminated to 100Ω or to 50Ω||50Ω, according to the manufacturers guidelines. All lines carrying clock signals must be treated with particular care. In the appendix there is a checklist for the detailed module design.

x do we have to say :
the scheme used on the prototypes
has been modified for lower ripple
and ~~the~~ the effect of the improvement
is currently being assessed

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 and boundary scans, IBERT/ChipScope tests, and firmware/software controlled playback/spy tests with any data source/sink device available. Initially the GOLD [7] 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.

Optical 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. Module control will be based on the IPbus firmware (VHDL based UDP stack) and software suite developed at the University of Bristol.

The test environment available in Mainz 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/ROSs. Therefore all system level tests will have to be done in an appropriately equipped L1 test lab.

4.1 Algorithmic firmware

There are several classes of physics algorithms currently being explored with help of physics simulations [11]. Most algorithms consist of identification of energetic objects and subsequent cuts on $\Delta\eta$ and ΔR. L1Topo system will be a processor crate equipped with two L1Topo modules. At Phase-0 an identical copy of the input data can be driven into each L1Topo processor.

is this current data ?

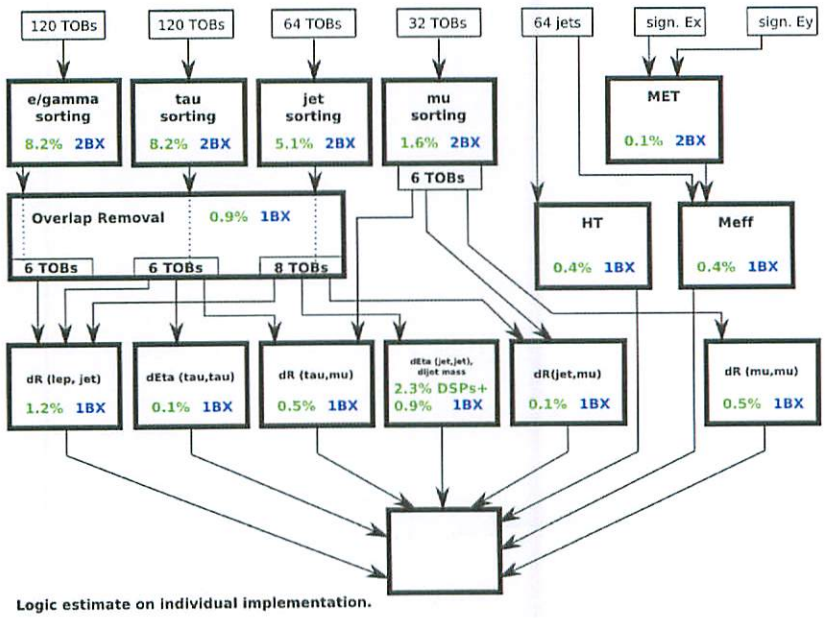

Figure 3: Algorithm implementation in one FPGA and corresponding logic usage (green) and latency (blue).

As an example, a block scheme of the algorithmic firmware implemented on a single processor is shown in Fig.3 including performance in terms of logic utilization and latency derived from Xilinx simulation tools. Incoming objects (TOB data, see Appendix H) are initially sorted followed by topological cuts 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. Further algorithms will be coded and simulated in several participating institutes. There will be a separate document on software and on algorithmic firmware [15] produced at a later date.

4.2 ROD firmware

ROD firmware is under development [12].

4.3 Tests and service firmware/software

The L1Topo will be initially tested to verify the functioning of the external hardware interface, board level connectivity and the basic operation of the hardware devices.

Initial test are aimed at:

1. Check on the power regulators and component fails.
2. Boundary scan performed using the GOEPEL SYSTEM CASCON tool, interfaced via USB.
3. All optical links integrity (miniPOD RX and TX ~~plus SFPs~~) using IBERT cores and loopback when possibile. As external source of any specific PRBS paths and check performance according to SFF-8431 standard [13]. ~~The GOLD [7] module or the~~ VC707 evaluation module equipped with a special mezzanine [9] can be used. VC709
4. Statistical eye and margin analisys and settings optimization on MGT and Avago miniPODs.
5. Check for signal attenuacy due to cabling and patching.
6. Test of Parallel connectivity among the two processor FPGAs.
7. Implementation of Playback and spy memories.
8. Test of the GTX phase shift and MGT clocks jitter analysis.
9. Latency measurements and algorithms performance comparison with simulation.

Module control will be achieved via standard Ethernet access.

please read through that again . . .

5 Interfaces

5.1 Internal interfaces

The L1Topo module carries the following daughter modules:

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

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

3. **Extension module**: 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).

4. **External JTAG programming module**: This module is for the system ACE and host the SD card for pre-configuration of the Kintex7 device. or an alternative configuration controller

5.2 External interfaces and communications protocol

L1Topo interfaces via optical fibers to the following external systems: CMXs, CTP and RODs for both DAQ and Level-2 ROI. Initially at Phase-0 an electrical connection to the CTP is provided. Control is provided by IPbus. A summary of the external interfaces is shown on Tab.1. For running at 12.8 Gb/s, 13Gb/s speed grade MGT and 14Gb/s speed grade miniPOD recievers are used. On the real-time data path fibre-bundles with multiples of 12 multimode fibres are used. Opto-electrical translation on the real-time path is made with miniPOD 12-channel transceivers. Far end transceivers are of the same type to minimize optical power budget. No optical splitting or other transceiver types are foresee on the real-time data path. However, optical power budget should be checked carefully. The external connections are made by MTP/MPO connectors. L1Topo is equipped with male MTP connectors.

receiver

(except micron fibres!)

footnote: unless via front panel

limited to 80Mb/s at CTP

Table 1: External interfaces

I/O	From/to	bandwidth		
RT input	various	160 × 6.4 (12.8) Gb/s	Opto / MTP 48	miniPOD, 8b/10b
RT output	CTP	12 × 6.4 Gb/s	Opto / MTP	miniPOD, 8b/10b
RT electrical	CTP	32 × up to 160Mb/s	via mezzanine	LVDS
Control		2×1Gb/s	Eth elect. Zone 2	
Spares		≈1Gb/s	miniPod	
DAQ	D-ROD	≈1Gb/s	miniPod	
ROI	R-ROD	2×≈1Gb/s	miniPod	
DAQ/ROI S-Link	ROD	12× up to 8Gb/s	miniPod	
IPMB			Zone 1	
LHC clock	TTC etc.		Optical	front panel

2. →

5.2.1 Communication Protocol on RTDP

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 has been considered. This might also simplify initial link start-up. The current version of encoding scheme is described in Appendix 5.2.1.

scheme

adopted

The baseline deserialization scheme is relying on a de-multiplexing the 6.4Gb/s data stream to 4 sub-ticks (word slices) of 32 bits wide, transmitted on a 160MHz clock. 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 40.0789 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.

→ I believe we are actually assuming we have 1 DAQ link?! and two copies of ROI link at standard slink rate (1.6 Gbps) everything else is spare!