ATCA-Based TP Strawman Design

Draft version 0.2 (January 27, 2010)

Table of Contents

Introduction	1
Design constraints	2
ATCA crate form factor and infrastructure	2
Virtex-6 FPGA speed limitations	2
Supporting hardware	2
CMM++	3
TP-CTPI	3
TP Modules and organization	4
External signals	4
TP-QM optical inputs	5
TP-GM electrical outputs to CTP (or TP-CTPI)	5
TP-ROD optical outputs	5
TP-TCM front-panel interfaces	6
Backplane Pinouts	6
TP-GM backplane pinout	6
TP-QM backplane pinout	6
TP-ROD backplane pinout	7
TP-TCM backplane pinout	7
Other design considerations	7
Geographical addressing	7
Network-based control and configuration	7
TTC distribution and decoding	7
DCS interface	8
Readout to TP-ROD	8
FPGA configuration and management	8
Figures (conceptual)	9

Introduction

This document outlines a strawman specification for a topological processor (TP) for the Atlas L1Calo phase-1 upgrade. The current design assumptions include the following:

- ROI data are transmitted from the L1Calo and Muon trigger processors to the TP via 12-fiber optical cables at 6.4 Gbit/s/fiber, and received by SNAP-12 modules.
- The TP design is a multi-board system housed in a crate with an ATCA form factor and power supply and equipped with full-custom, high-density backplanes
- Control and configuration is performed over gigabit ethernet, and each module in the crate has a local controller running a TCP/IP stack.
- The design is based on available FPGAs from the Xilinx Virtex-6 product family.

We begin with an overview of the technologies with an emphasis on their impact on design constraints. Afterwards we present the strawman design.

Design constraints

ATCA crate form factor and infrastructure

The ATCA form factor for a 19" chassis has an 8U front panel height and 600mm frame depth. The backplane slot pitch is nominally 1.2", thus accommodating up to 14 boards. The backplane is divided into three horizontal zones (from bottom to top):

- Zone 1 contains redundant -48V power, hardware address pins, and two I2Cbased shelf management signal buses.
- Zone 2 has a connector height of up to 125mm, and is nominally populated by up to 5 ZD connectors, each of which provides 40 differential signals that can be driven at rates up to 5 Gbit/s.
- Zone 3 is a user-definable backplane area with up to 80mm connector height. It is usually used for connection to rear transmission modules without the need for an interposed backplane. But FCI (for example) produces a high-density connector family (AirMax), which can accommodate up to 200 differential pairs or 300 single-ended signals in 80 mm of connector height, and some parts of the AirMax family are suitable for backplane use.

Between the Zone 2 and Zone 3 backplane connectors are guide pins and horizontal extrusions, so it is not possible to increase the available vertical connector height without severely compromising the ATCA chassis architecture. It also means that all interconnections between modules must stay within a single zone.

Virtex-6 FPGA speed limitations

The Virtex-6 has single-ended I/O pins that function up to 800 MHz, while most of the internal FPGA features (memory, DSP) function at up to 600 MHz.

This probably sets the practical limit for large, complex trigger algorithms at 320 MHz ($8 \times$ LHC clock rate), with I/Os driven at double data rate (640 Mbit/s). We use this assumption for purposes of this design.

Supporting hardware

The TP system receives data from redesigned common merger modules (CMM++) located in the CP and JEP crates and possibly redesigned muon interfacs octant modules (MIOCT) in the muon CTP interface.

Output to the CTP will be carried over a number of LVDS cables whose number and arrangement will almost certainly change multiple times over the course of commissioning. A flexible CTP interface (TP-CTPI) to optimally route and monitor trigger signals to the CTP would provide flexibility on where and how trigger bits are produced, simplifying firmware development for the CMM++, TP and CMM.

The full specifications for the CMM++ and TP-CTPI are beyond the current scope of this document, but as they provide the data source and sink for the topological processor we provide a short overview of them here.

CMM++

The CMM++ is a 9U×400mm module designed to fit in the CMM positions in the CP and JEP processor crates. It is capable of being configured as a drop-in replacement for the current CMM, but contains additional features for topological trigger running.

- Backplane signals
 - 400 single-ended inputs from up to 14 CPMs or 16 JEMs at data rates up to 160 Mbit/s
 - 84 differential pairs (LVDS) to the CMM rear-transition module for electrical fan-in/out of crate level results to system CMM for legacy operations
 - o VME-- interface, geographic address pins, TTC and CAN
- Front panel connections
 - Three 100-pin MegArray footprints for up to three SNAP12 transmitter modules for data fan-out to TP crate
 - Three 100-pin MegArray footprints for up to three SNAP12 receiver modules for possible use in Stage 2a topological processing
 - One optical readout fiber to ROI ROD
 - Two optical readout fibers to DAQ ROD
 - Two 68-way D-subminiature outputs to CTP
 - o Coaxial 40 MHz clock output to CAM module

ТР-СТРІ

The TP-CTPI is a small rack-mounted subsystem located near the CTP. A single large FPGA (Xilinx Virtex 6) receives differential LVDS trigger bits at nominally 40 or 80 Mbit/s from the energy-sum CMM, the CMM++ modules, and the outputs from the TP (TP-GM, see below). The trigger bits are optimally rearranged and routed to the CTP over 104 differential pairs on six LVDS cables. The nominal system connectivity would be:

- Input connections
 - o Two (2) 68-way D-subminiature LVDS cables from sum-E CMM
 - o Two (2) 68-way LVDS cables from jet CMM
 - o Two (2) 68-way LVDS cables from EM and Hadron CMMs
 - Nine (9) 68-way LVDS cables from three TP-GM modules
- Output connections
 - Six (6) 68-way LVDS cables to CTP
- Other interfaces
 - RJ-45 connector for gigabit Ethernet interface (for control and configuration)
 - Optical G-link output for readout to DAQ ROD
 - Optical TTC input
 - o 9-pin D-sub CAN connector to DCS

The TP-CTPI may also be used to monitor and histogram all CMM++ and TP outputs without pre-scaling, even for bits produced but not sent to the CTP.

TP Modules and organization

For logistical and board complexity concerns, the strawman TP design is based on several different modules connected by a full-custom backplane. For the purposes of this exercise we list them here along with convenient acronyms:

- The TP global merger (TP-GM) receives and processes quadrant-level results via the backplane and reports them to the CTP.
- Two TP quadrant input modules (TP-QM) per TP-GM each receive a full set of ROI data from two opposing quadrants on seven SNAP12 optical receiver modules. The TP-GMs preprocess the input data, sharing information between them if needed, and send selected results to the TP-GM.
- Readout of the TP-QMs and TP-GMs is performed by one or more TP readout-drivers (TP-ROD) in each crate. The current working assumption is one TP-ROD per "triplet" (comprising 2 TP-QMs and 1 TP-GM).
- A timing and control module (TP-TCM) receives and fans out a differential TTC signal to each module in the crate, provides a CAN bus for interfacing to DCS, and serves as a gigabit Ethernet for control and configuration of all modules.

For the purposes of this strawman design, one TP chassis is assumed to accommodate three identical TP "triplets", each with its own TP ROD for readout. These modules are serviced by a single TP-TCM (see Figure 1).

We also assume for the moment custom Zone 2 and Zone 3 backplanes, each with up to 200 differential pairs per slot (for a total of 400 pairs per module).



Figure 1: Possible TP crate organization

External signals

The TP system receives real-time data from L1Calo and L1Muon subsystems and produces real-time results to be sent to the CTP. In this section we list the external signals received and produced by the different modules

TP-QM optical inputs

The TP-QM receives real-time ROI data on seven (7) optical cables with 12 fibers each. They are converted to 6.4 Gbit/s differential signals by SNAP12 receiver modules.

Cable #	Fibers	Total	Comments
	used	data bits	
1	12	1344	EM crate 0 (1) 14 CPMs \times 96 bits
2	12	1344	EM crate 2 (3) 14 CPMs \times 96 bits
3	12	1344	HAD crate 0 (1) 14 CPMs \times 96 bits
4	12	1344	HAD crate 2 (3) 14 CPMs \times 96 bits
5	12	1536	JEM crate 0 (1) 16 JEMs \times 96 bits
6	8	576	4 MIOCTs (Quadrant 0) × 144 bits / MIOCT
7	8	576	4 MIOCTs (Quadrant 2) × 144 bits / MIOCT
Total	76	8064	

TP-GM electrical outputs to CTP (or TP-CTPI)

The TP-GM module produces L1 real-time results for transmission to the CTP. The CTP currently accepts 104 bits of L1Calo results at 40 Mbit/s on six differential LVDS cables. The data volume to the CTP may be doubled by running the input cables at twice the speed (80 MHz) at the price of reduced flexibility in routing trigger bits to the CTP inputs.

The 68-way D-subminiature connectors used for the current L1Calo to CTP interface occupy a front-panel height of about 64mm. The 8U front panel of the TP-GM will comfortably accommodate three such cables, each of which can carry up to 33 signal pairs to the CTP representing 31 data signals, a clock signal and 1 parity bit.

These 93 output bits (or 186 at double data rate) allow each TP-GM to drive nearly all of the non-energy-sum CTP inputs allocated to L1Calo. Under normal circumstances a single TP-GM would use just a portion of these available outputs. The TP-CTPI may be used to receive and rearrange the signals to send to the appropriate CTP inputs.

TP-ROD optical outputs

The TP-ROD will need one or more readout links to send results to the DAQ system. The current 9U L1Calo ROD uses up to four commercial HOLA S-link boards mounted on a rear-transition module. To simplify hardware and firmware development, the TP-ROD may well be built to accommodate one or more HOLA boards (which are still available from CERNTECH in Hungary (www.cerntech.hu). Alternatively the S-link protocol may be emulated in an FPGA and transmitted by board-edge laser modules.

TP-TCM front-panel interfaces

The TP-TCM serves as the interface to the TTC and DCS systems, as well as a remote network controller. The front panel of the TP-TCM therefore include at least the following three ports:

- A RJ-45 1000BASE-T Ethernet jack
- A fiber receiver for the TTC fiber
- A DB9 connector for CANbus connection to DCS

Backplane Pinouts

The largest bottleneck in a backplane-based TP is the input to the TP-GM. We will therefore begin with the backplane pinout for this module.

TP-GM backplane pinout

The total pinout to the TP-GM is 400 differential pairs. A few of these are required for control, timing and readout, while the remainder is usable for global merging.

Port name	# pairs	Comments
Ethernet	4	Standard Gbit Ethernet port for control/configuration
TTC	1	From the TP-TCM
CAN	1	CANbus interface
Readout	4	To TP-ROD
Merging	390	From the TP-QMs. Up to 6240 bits at 640 Mbit/s
Unused	0	

Without parity, this suffices for a 32×32 eta-phi map with 6 bits per cell, plus 96 bits. Note that this assumes that all QM-to-GM lines are used entirely for data transmission, which is unrealistic. The six links representing the "extra" 96 bits might well be used to carry up to 3 data clock signals each from the two TP-QMs for input synchronization The data volume of 6240 bits into the TP-GM corresponds to about one third of the total data volume received by the TP-QMs.

TP-QM backplane pinout

The total pinout to the TP-QM is up to 400 differential pairs. Assuming two TP-QM modules per TP-GM, 195 pairs are required for sending output to the TP-GM. This leaves 195 signal lines that can be routed between the two TP-QM modules in the same triplet for sharing information such as features at the quadrant boundaries, etc.

Port name	# pairs	Comments
Ethernet	4	Standard Gbit Ethernet port for control/configuration
TTC	1	From the TP-TCM
CAN	1	DCS interface
Readout	4	To TP-ROD in ~50 ticks/BC at 320 Mbit/s
To GM	195	Up to 3152 bits at 640 Mbit/s
To other QM	195	Can choose the direction of each connection
Unused	0	

TP-ROD backplane pinout

The total pinout to the TP-ROD is up to 400 differential pairs. Each TP-ROD accommodates one TP "triplet" (two TP-QMs and one TP-GM). Each processing module sends four differential signals to its respective TP-ROD. The bit rate of these signals is currently unspecified, but the connectors and backplane could accommodate multiple-Gbit links.

Port name	# pairs	Comments
Ethernet	4	Standard Gbit Ethernet port for control/configuration
TTC	1	From the TP-TCM
CAN	1	DCS interface
Readout inputs	12	From 2 TP-QMs and 1 TP-GM
Unused	382	

TP-TCM backplane pinout

The total pinout to the TP-TCM is up to 400 differential pairs. The Ethernet and TTC signal distribution is performed through fan out of point-to-point links to the other module positions, while the CAN is a differential bus connected to all modules.

Port name	# pairs	Comments
Ethernet	52	Point-to-point distribution to other 13 positions
TTC	13	From the TP-TCM
CAN	1	DCS interface
Unused	324	

Other design considerations

Geographical addressing

The ATCA Zone 1 connector includes a geographical address, available at power-up, unique for each module in the crate. This address can be used to:

- Select the correct firmware to be loaded
- Assign a fixed IP address for the Ethernet interface
- Set fixed CAN and TTC addresses

Network-based control and configuration

The TP subsystem uses Ethernet for the control and configuration interface. This means that each module must have a local Ethernet controller. One option is a singlechip embedded system with a software TCP/IP stack. A second, possibly favored alternative is a hardware-based GigE TCP/IP Offload Engine (TOE), which handles incoming and outgoing data at the full link speed and interfaces easily with an FPGA. To simplify software development, each module has a fixed IP address set by its unique geographic address (see above).

TTC distribution and decoding

The TP-TCM receives and electrically fans out the incoming TTC signal to all modules in the crate. A TTCrx chip is therefore required on each module. An attractive solution is the existing L1Calo TTCdec board, which produces two 40 MHz

deskew clocks, with or without PLL cleaning. The TTCrx can have a unique address (and other internal parameters) set through an I2C interface managed by an FPGA or microcontroller. The TTC address to be used is determined by the module's geographic address.

Additional PLL-based devices (e.g. VCO) on each board condition and clean selected clock signals to provide low-jitter timing for multi-Gbit transceivers.

DCS interface

The TP-TCM is interfaced with the DCS by a front-panel CAN connector, and with a CANbus on the backplane shared by all modules in the crate. The CAN interface in all the current processor modules is implemented in a common microcontroller (Fujitsu MB90FG94). To simplify development, a common microcontroller (identical or compatible with the MB90FG94) should be used on all TP modules. As for TCP/IP and TTC, the CAN address is set at power-up according to the modules' geographic address.

Readout to TP-ROD

The TP-ROD receives and processes readout data from a TP triplet over four multi-Gbit electrical links from each module. As the main processor FPGAs on the TP-QM and TP-GM may not have sufficient high-speed transceivers available, one simple solution is using a 10-Gbit Ethernet transceiver (XGMII-XAUI) to read out each module over four 3.2 Gbit XAUI links.

The XGMII interface includes 4 lanes of 8 data bits clocked at 320 Mbit/s, plus a data clock and four "control" bits that designate whether a byte contains data or an idle control character for lane alignment. By sending idle bytes between valid readout frames, the arrival of a new readout frame at the TP-ROD can be signaled by the falling edge of the control characters, in the same way that G-link readout uses the rising edge of the DAV bit.

FPGA configuration and management

The FPGAs on each module are configured through an advanced firmware management solution (Xilinx SystemAce or equivalent). As for existing L1Calo modules, the configuration memories on board all TP modules of a single type include all possible firmware sets. The correct set of firmware is selected and loaded according to the modules' geographic address.

Figures (conceptual)







Figure 3: TP quadrant module layout (TP-QM). Two FPGAs each receive and process data from one quadrant. Single-ended links between the two FPGAs, as well as differential backplane links between the two QMs in each triplet carry environmental and other data for preprocessing, while 195 differential backplane links carry selected results to the GM.



Figure 4: TP global merger layout (TP_GM). A single FPGA receives 195 differential links from each of two QM modules in the triplet.



Figure 5: TP readout driver (TP_ROD). Three multi-gigabit transceiver chips (GTX) each receive four serial readout streams from two TP-QM and one TP-GM modules in a triplet and forward them to a processor FPGA. The optical output is in S-link format.