NUCLEAR PHYSICS DIVISION
FAST ELECTRONICS GROUP

Description and Requirements
for the

Hall B Trigger Processor

DRAFT

By cuevas at 9:47:01 AM, 2/10/2014

Scott Kaneta (skaneta@jlab.org)
Chris Cuevas (cuevas@jlab.org)
Ben Raydo (braydo@jlab.org)

10 February 2014

Table of Contents

1 - Introduction

The new Level 1 Trigger system being designed for the Jefferson Lab 12GeV experimental halls features a hierarchical path for trigger data. Over the course of development, the distinction between the Global Trigger Processor (GTP) and the Crate Trigger Processor (CTP) has been reduced to a point where the benefits of merging their designs have become clear. The GTP prototype includes minor differences that allow it to serve in both roles using different firmware configurations. Particularly for Hall B, this consolidation of hardware will reduce costs and increase functionality, capitalizing on the newest technologies available.

The rapid advancement of FPGA technology has already yielded a significant improvement in the performance and power of the level 1 trigger, as shown in the HPS test run. This flexibility will allow one module to serve in two distinct roles, reducing the number of unique hardware designs and increasing the production quantity of one, ultimately reducing the cost by taking advantage of the economy of scale. In addition, a firmware platform encompassing the common elements of both designs will make modification and support easier as new applications are developed.

2 – Trigger Requirements

While the specific applications will be continually developed for years to come, many elements of the hardware must be fixed to allow the system to be produced and deployed. The known experimental specifications help drive the hardware requirements. The most rigorous application thus far has been for the Heavy Photon Search (HPS) experiment which pushed the current CTP prototype design to its limits.

2.1 - Heavy Photon Search (HPS)

The limitations of the CTP prototype had a measureable impact on the results of the HPS test run. The bandwidth of its high speed links to the Flash ADCs (fADC) and the Sub-System Processor limited the granularity of the trigger data. In addition, the split FPGAs reduced the effective processing power, prohibiting the addition of desired features and increasing firmware development time. The production run imposes even greater requirements to which the hardware and firmware must adhere.
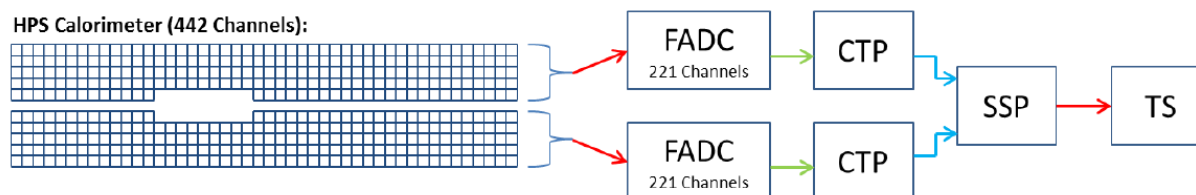
Figure 2.1a: HPS Trigger Configuration

As shown in figure 2.1a, the CTP collects data from fourteen fADCs which provide energy pulse information for one half of the calorimeter. It must then calculate the energy sums for hits on multiple channels within a 3x3 window in order to detect the total energy from a single event. The position and energy are reported to the SSP for coincidence with hits on the other half of the calorimeter.

HPS ECAL

The Electromagnetic Calorimeter (ECAL) is split into two halves, each containing 221 crystals arranged in a 46x5 grid with an offset gap near the beam pipe which removes nine. Figure 2.1b shows the arrangement of the calorimeter. The crystal position on the X and Y axes are shown on the bottom and left sides. The numbers within the black line represent the number of channels that contribute to the 3x3 cluster computation centered at that location.

y \ x	-22	-21	-20	-19	-18	-17	-16	-15	-14	-13	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
4	4	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	4
3	6	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	6
2	6	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	6
1	6	9	9	9	9	9	9	9	9	9	9	8	7	6	6	6	6	6	6	6	7	8	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	6
0	4	6	6	6	6	6	6	6	6	6	6	6	5										5	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	4

Figure 2.1b: HPS ECAL (1 of 2)

Cluster Finding

In order to find 3x3 clusters, the data coming from multiple fADCs must be mapped to cluster processors, firmware elements designed to operate on a single 3x3 window. Taking advantage of the overlap of windows on the perimeter, only 125 cluster processors are required to cover the entire space. Given an arbitrary channel, its data must be shared with up to nine cluster processors, creating a web of interconnectivity between cluster processors.   Due to various characteristics of the calorimeter and trigger system, a programmable time window must be created allowing the cluster computations to include data that arrives within a certain period to ensure hit energy is correctly calculated.

Specifications
- Cluster Processors: 125
- Cluster Window: programmable (0-28ns, 4ns steps)
- Inputs
  - 221 channels
  - 13 bit energy + 3 bit time every 32 ns per channel
  - **8 Gbps data rate**

- Outputs
    - 16 bit cluster energy
    - 9 bit time
    - 9 bit hit pattern (3x3)
    - 15 bit position
    - **16 Gbps data rate**

2.2 Global Trigger Processing

The Trigger Algorithm Processing is the final stage of the global Level 1 trigger formation, and the GTP is connected to   all subsystem data streams. The GTP can create up to 32 triggers that can be routed to any of the 32 front panel outputs of the GTP board. These front panel signals are cabled to a rear transition module that is directly coupled to the Trigger Supervisor (TS).

Currently the following subsystems and variables are supplied to the Trigger Algorithm Processors:

Level 1 Subsystem	Level 1 Subsystem variable
Forward Calorimeter	ADC sum
Barrel Calorimeter	ADC sum
Time of Flight	number of tracks
Tagger	min, max, track count
Start Counter	number of tracks

Ultimately these subsystem variables will be put into a trigger equation to decide whether or not to create a Level 1 trigger, but first they must be conditioned to take into account physics events that span several 250MHz clock cycles.

2.2.1 - ADC Sum & Track Count Variables

The subsystem ADC variables report the sum of all ADC channels for that subsystem in the current clock cycle. These values must be integrated over a time window to form the energy from events (same idea as gating of an ADC). Figure 2.2.1a shows how this can be accomplished by feeding the current subsystem ADC sum into a programmable length shift register (implemented as a FIFO in the hardware). The shift registers and energy will start out as zero in value and as the ADC sums enter the shift register they will be added to the subsystem energy variable. As the ADC sums exit the shift register they will be subtracted from the subsystem energy variable. The result is that the subsystem energy variable contains the integrated ADC values for the subsystem for the length of time determined by the programmable length of the shift register. The forward and barrel calorimeter FADC will be processed in this manner.

The track count variable is treated exactly the same as the ADC sum variables. They are fed into a programmable length shift register that sums the shift register contents to integrate the variable

over time (see Figure 2.2.1a). The result is the total track count over the sliding window period. The time of flight and start counter subsystems are processed in the same way.
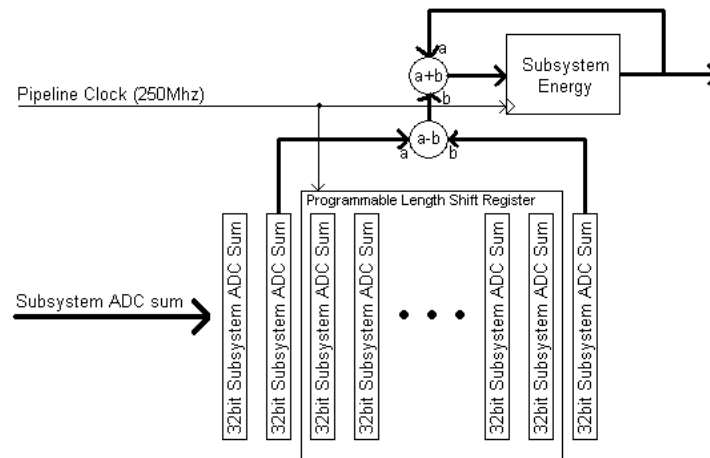
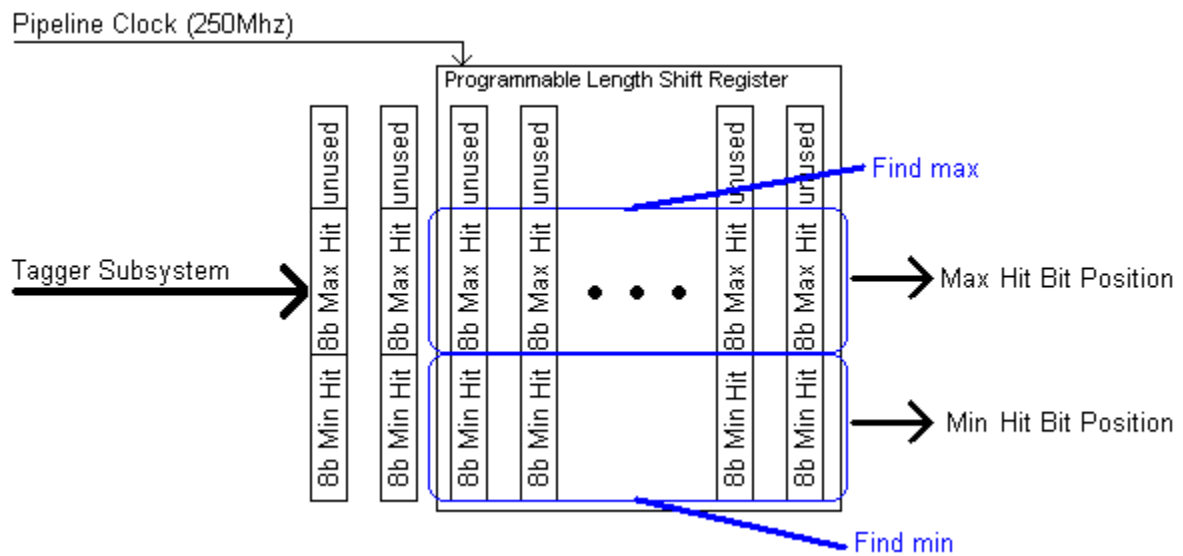Figure 2.2.1a: Energy & Track Count Processing

Figure 2.2.1b: Tagger Processing

2.2.4 - Trigger Equations

The final stage of the Trigger Algorithm Processor takes the conditioned subsystem variables and uses them to compute an arbitrary equation to generate a trigger or not. As a realistic example an equation, taken from Dave Doughty's Level 1 Trigger talks (good spot for a reference footnote), was coded and synthesized into one of the Trigger Algorithm Processors (Virtex 5 LX220 FPGA). This equation has the following form:

$$f(TTOF, EFCal, EBCal) = TFM*TTOF + EFM*EFCal + RM*(EFCal+1)/(EBCal+1)$$

TTOF: Tracks Forward TOF

EFCal: Energy Forward Calorimeter

EBCal: Energy Barrel Calorimeter

TFM,EFM,RM: Programmable Coefficients

Every 4ns a trigger condition can be computed as follows:

if (f(TTOF, EFCal, EBCal) >= Z) and (MinTaggedHit > Y) and (MaxTaggedHit < X)

Trigger = 1

else

Trigger = 0

Inside the FPGA, the 32bit subsystem variables are received and converted from 32bit integers to 32bit floating point numbers. The example function, f, was synthesized into the FPGA using pipelined, floating point arithmetic (32bit "real" numbers) and ran at over 300MHz with a computational latency of 46 clock cycles (184ns @ 250MHz). This project used up roughly 3% of a Stratix IV GX 180 FPGA's resources. This project clearly demonstrates that each of the four Trigger Algorithm Processor FPGAs has room for multiple and more complex trigger equations.

Trigger processing produces up to 30 trigger bits that can be routed to various positions on the output. The 30 trigger outputs and two clocks will be sent to the TS using four 8-channel Densishield connectors over 100 ohm differential pairs.

2.3 BCAL

TBD

2.4 ECAL

TBD

2.5 DCRB

TBD

3 – Configuration and Status

As a member of the larger trigger system, a common control interface is critical to successful integration. Since VXS Switch modules are not connected to the VME bus, the primary mechanism is via an I^2C link to the Trigger Interface (TI) module which transfers Switch module data to its VME connection. The Switch modules are allocated space within the TI's memory map for register access. However, this serial interface is slow, has limited space, and requires access to the VME bus which may not be ideal during readout.

3.1 Registers

The register interface, which is exposed through the TI's VME connection via I^2C, has some disadvantages during normal operation. However, it is well suited for configuration before an experimental run begins and ensuring Ethernet connectivity despite changes in its configuration. In general, some registers are exposed via I^2C while the rest are available through Ethernet. Refer to the register map document for more details.

3.2 Diagnostics and Monitoring

Scalars

In addition to normal registers, scalars will also be supported to provide rate and counting measurements. Scaler properties are outlined in table 3.2a.

	Sensitivity	Threshold	Latch Condition	Gate Condition
Trigger Processing	Edge/Level	Yes	Register	
Sync/Trig1/Trig2	Edge	No	Register	
Clock (250MHz)	Edge	No	Register	
Error Rates	Level	No	Register	

Table 3.2a: Scaler Properties

Histograms

TBD

3.3 FPGA Configuration

The FPGA will have support for up to seven user-defined configurations and one factory image. Configurations that are already written to non-volatile memory can be loaded via software control.

New images may be written to specified pages during maintenance periods to allow for new functionality and reverting back to known configurations if necessary. Exposing this functionality will be restricted.

3.4 Ethernet

A communication path independent of the VME bus allows access to the module regardless of other activity within the crate. Ethernet is well suited for this since it can use existing infrastructure and protocols to simplify integration. In addition, it allows much greater payloads of data for configuration, monitoring, and diagnostics.

Ethernet Overview

There are many approaches for the implementation of Ethernet. As a first attempt, a low cost VHDL TCP/IP stack will be used. A firmware solution has its share of advantages and disadvantages over a software approach as indicated in Table 3.4a. The bolded entries indicate the stronger benefit.

	Firmware	Software
Bandwidth/Throughput	**Limited by client**	Limited by processor speed of host and client
High Level Protocols	Implemented using VHDL state machine or embedded processor	**Software libraries exist for CMSG, may be portable**
Operating System Support	None	**Linux, RTOS, etc**
Cost	**$1500, JLab may already have license**	Variable depending on solution, can be >$10k

Table 3.4a: Comparison of Firmware and Software Ethernet Solutions

The Ethernet interface will be developed in several stages. The firmware stack is currently being evaluated but has been successful in managing TCP sockets. A simplified protocol will be used initially for register access and will be refined to conform to the CMSG protocol for publishing its status information.

3.5 User Interfaces

The current graphical user interface (GUI) environment used by the Fast Electronics group is ROOT. The I2C and Ethernet interfaces will adhere to existing interface specifications to simplify GUI development with other environments.

3.6 PCIe

4 – Trigger System Overview

The high level trigger system architecture is well defined. The new Trigger Processor is designed to fulfill two distinct roles using one hardware design. Figure 4a shows the locations of the Trigger Processors (TP) in the Front End and Global Trigger crates. In the Front End crate, it serves at the Crate Trigger Processor and must receive data over the VXS backplane from 16 fADC modules and

send to the Sub-System Processor via its four-channel fiber optic transceiver. In the Global Trigger crate, it functions as the Global Trigger Processor, receiving its inputs from up to 16 SSP modules and sending up to 30 trigger bits and two clocks to the Trigger Supervisor.
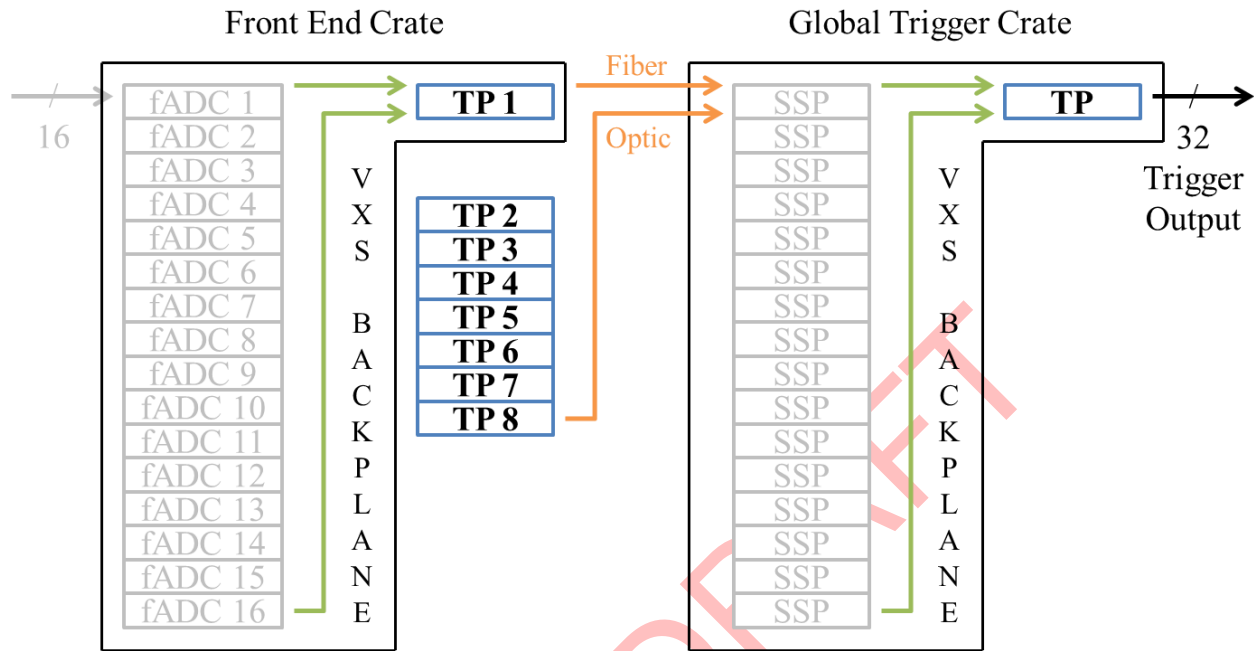

Figure 4a: Crate Configurations

5 – Functional and Hardware Design

The GTP unit can receive up to 16 simultaneous subsystem variable streams to generate Level 1 triggers. Each subsystem variable stream provides a 32 bit quantity every 4ns (250MHz clock) indicating the current status of that subsystem (i.e. the total energy, number of tracks, or min/max energy for that subsystem in the current clock cycle). The eight 32 bit system variables are evaluated in up to 30 trigger equations and each can be mapped to a trigger output bit high or low as defined by the result of the trigger equation. All logic and computations run at the global 250MHz clock and are pipelined to allow continuous trigger decisions.
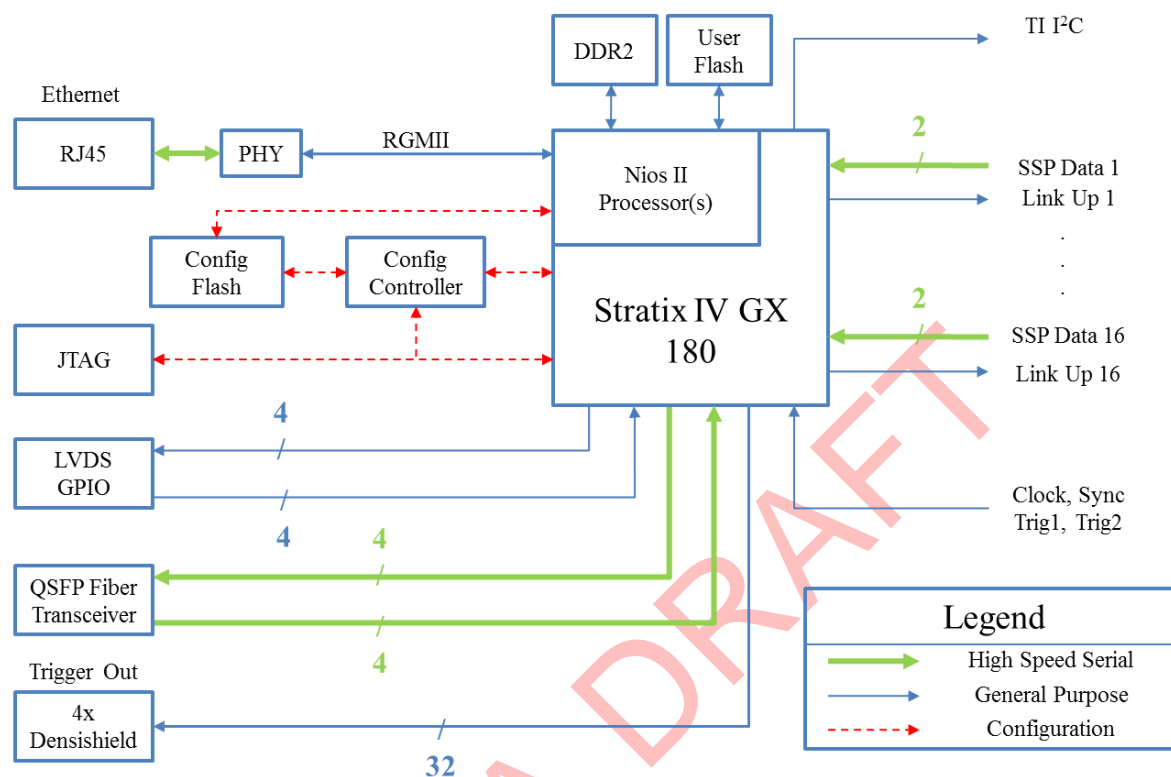
Figure 3a: Global Trigger Processor Hardware Block Diagram

5.1 Trigger Data Paths

In accordance with the Level 1 Trigger system architecture, the Trigger Processor must use the VXS backplane port definition, a Fiber Optic transceiver, and Densishield Trigger Outputs in order to interface with existing trigger modules.

Payload Port Transceivers

The SSP modules each communicate to the GTP via a 2 lane high speed serial VXS connection. Each lane operates at 5 Gbps providing a 10Gbps link. Using the Xilinx Aurora protocol this 10Gbps channel allows each SSP to provide a 32bit "instantaneous" subsystem quantity every 4ns. This channel operates in simplex mode where data is sent only in the SSP->GTP direction. Transceiver FPGAs using all four lanes to each of the 16 payloads is cost prohibitive and appears to be unnecessary given current requirements.

During initialization of the serial links, it is convenient for the SSP modules to know when the GTP has initialized its receiver on the 10Gbps channel. To facilitate the initialization of this simplex channel an asynchronous status line (shown as "LINKUP_SSPx" in the diagram above) is sent from the

GTP to each SSP that signals when the channel has been properly initialized (i.e. when the receiver PLL has locked onto the data stream, aligned the data, and "bonded" the lanes to form a channel). While the SSP unit sees the LINKUP_SSPx signal low it will continue to transmit the sequences that allow the GTP to properly initialize its receiver. The SSP or GTP can eventually timeout after the initialization has started if it is not able to establish a link and report an error to the CPU indicating where the fault may exist.

The Xilinx Aurora protocol uses the standard 8b10b line encoding scheme found in modern high speed networking protocols. Some of the main features of this encoding scheme include: DC balanced data stream, limited run length, and single bit error detection. Prototyping with the high speed serial links has revealed extremely low error rates when used on the VXS backplanes with the Xilinx Aurora protocol. While there is available bandwidth in the SSP->GTP links to provide capability for error correction, the error rates have shown to be low enough such that we can reliably use error detection and either ignore the erroneous data or restart the Level 1 trigger system to avoid use of corrupted data.

As shown in Figure 3a, one FPGA will process all sixteen sets of SSP serial data streams. It is responsible for retiming the subsystem data streams such that a fixed latency trigger can be formed by the final computational stage. The retiming is done by using the fixed latency "SYNC" signal from the Trigger Inteface (TI). When the Level 1 trigger system starts, by deassertion of the "SYNC" signal, it will buffer the subsystem streams in a FIFO and after a fixed number of cycles (the number of cycles is determined by ensuring each SSP has been able to send at least to the GTP 2 subsystem words) the FIFOs will be read out simultaneously, which turns the SSP streams into a fixed latency stream for the final stage in the Level 1 trigger.

Deskewing the subsystem data streams such that correlated physics events arrive at the Trigger Algorithm Processors at the same time will simplify the final processing stage.

In order to offload some of the processing in the Trigger Processor, the unused serial lanes may be used to allow payload ports to communicate directly with each other. This could provide data about the boundaries between adjacent payload modules with regard to the mapping of the detectors.

Note: The transmit-to-receiver latency can vary by a few clock cycles between the Xilinx Aurora links each time the receive locks to the incoming data stream, but once locked the latency is fixed. Ultimately the GTP retimes the final trigger decision such that its latency is fixed with respect to the physics event (accurate to within roughly 4ns).

Fiber Optic Transceiver

A new QSFP (quad single fiber pluggable) fiber optic transceiver has been selected for use on the SSP which allows data rates above 5 Gbps per lane. With all four lanes operating at 5 Gbps, this meets the data rate specification of 16 Gbps indicated in section 2.1 above.

Trigger Outputs

In order to interface with the Trigger Supervisor (TS), the Trigger Processor must have four Densishield Trigger Output connectors, each containing eight LVPECL differential pairs. This allows transmission of up to 30 trigger bits and two clocks at the system clock rate of 250 MHz.

5.2 Trigger Control Paths

Clock, Sync, Trig1, Trig2

All modules in the Level 1 trigger accept the common signals Clock, Sync, Trig 1 and 2. All trigger algorithms are run using the distributed system clock running at a rate of 250 MHz. Sync provides system synchronization provides global timing references. Trig1 and 2 are the trigger signals that are used for data acquisition and readout.

Ethernet

Ethernet hardware consists of a Gigabit Ethernet PHY (physical layer device) and an Ethernet jack with integrated magnetics. The PHY interfaces with a MAC (medium access controller) which is implemented in the FPGA as VHDL or IP core.

PCIe

PCIe support requires further discussion. The current FPGA has no remaining transceiver lanes so a significant change would be required to provide any PCIe support. This is also dependent on the use of a crate controller which supports PCIe over the VXS backplane. Currently, the model from Concurrent is the only candidate and is significantly more expensive than the GE competitor.

I2C

TI Link

The Switch A -> TI link is a differential pair that provides a faster communication path than the single-ended I^{2}C interface is capable of. The content of this link is currently undefined.

LVDS I/O

In addition to the LVPECL Trigger Output cables, there are four LVDS outputs and four LVDS inputs for use as general purpose I/O. Applications may include external triggers, monitoring or diagnostics.

5.3 FPGA

Given the dynamic nature of trigger processing requirements, the FPGA should allow as much flexibility in implementing new trigger algorithms while remaining within cost and time budgets. The GTP prototype uses a Stratix IV GX 180 which is the largest single FPGA of any of the Level 1 Trigger modules at Jefferson Lab. It is footprint compatible with several other densities, providing a low risk upgrade path.

6 - Specification Sheet
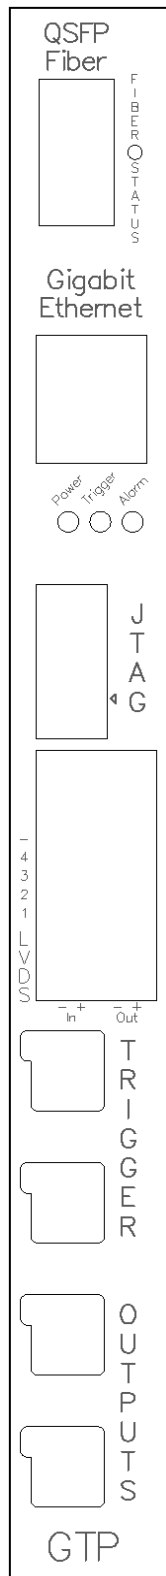
MECHANICAL
- Single width VITA 41 (VXS) Switch Module

HIGH SPEED SERIAL I/O:

From Switch B (Signal Distribution module)
- 250MHz Clock (LVPECL)
- Trig1, Trig2, Sync (LVPECL)

To Payload Ports 1-16 (SSP/spare modules)
- 2x 5 Gbps SSP Stream (Xilinx Aurora)
- LINKUP out (LVTTL)
- BUSY out (LVTTL)

Outputs:
- Trigger output decision time <512ns

Front Panel:
- 1x RJ45: Ethernet
- 4x Densishield: Trigger Signals (LVPECL)
- 1x QSFP Fiber Transceiver
- 1x JTAG

Indicators: (Front Panel)
- Power – Blue LED
- Trigger – Amber LED
- Alarm – Red LED

Programming:
- On board JTAG Port
- 100/1000Mbps Ethernet Interface
- Trigger equation and processing for up to 16 JLAB SubSystem Processors

Power Requirements:
- +5v @ 10 Amps (From Backplane)
- Local regulators for other required voltages

Environment:
- Forced air cooling: N CFM
- Commercial grade components (85°C max)

7 - Power Supply and Current Consumption
8 - VXS Pinout Table

VXS Port	15	13	11	9	7	5	3	1
RX0	PP13	PP11	PP9	PP7	PP5	PP3	PP1	PP2
TX0								
RX1		PP15	PP13	PP11	PP9	PP7	PP5	PP3
TX1								
RX2	SSP	SSP	SSP	SSP	SSP	SSP	SSP	SSP
TX2	FADC	FADC	FADC	FADC	FADC	FADC	FADC	FADC
RX3	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB
TX3								
SCL	Busy	Busy	Busy	Busy	Busy	Busy	Busy	Busy
SDA	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp

VXS Port	2	4	6	8	10	12	14	16
RX0	PP4	PP6	PP8	PP10	PP12	PP14	PP16	
TX0								
RX1	PP1	PP2	PP4	PP6	PP8	PP10	PP12	PP14
TX1								
RX2	SSP	SSP	SSP	SSP	SSP	SSP	SSP	SSP
TX2	FADC	FADC	FADC	FADC	FADC	FADC	FADC	FADC
RX3	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB	DCRB
TX3								
SCL	Busy	Busy	Busy	Busy	Busy	Busy	Busy	Busy
SDA	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp	LinkUp

VXS Port	17	18	B1	B2	B3	B4
RX0	Diff Pair*		Trig 1			Clock
TX0						
RX1			Trig 2			
TX1		TI GTP RX				
RX2			Sync			
TX2						
RX3		TI Busy				
TX3		TI GTP TX				
SCL	SCL*	TI_SCL				
SDA	SDA*	TI_SDA				