DRAFT
November 15, 2006

L1Cal2B NIM Article

Michael Mulhearn

Abstract

My contribution to the L1Cal2B NIM article:

```
e. ADF-to-TAB Transfer..............................Mike
   - synchronization scheme
   - cables
f. TAB/GAB System...................................Mike
   - TABs
   - GAB
   - VME/SCL
   - TAB/GAB crate
   - TAB/GAB outputs
   - Diagnostic capabilities
g. Readout
   - L2/L3..........................................Mike
   - Cal-Track......................................Mike
h. Performance......................................Mike
   - BER limits
   - latency
   - readout deadtime
```

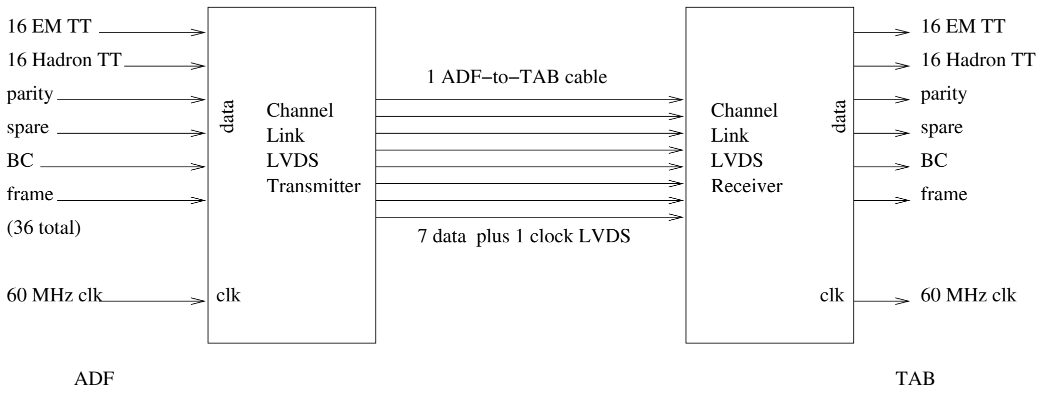

Figure 1: The ADF-to-TAB transmission.

1 ADF-to-TAB Transfer

The digitized trigger-tower data from each ADFs four-by-four η-ϕ region is sent to the TAB cards for further processing. To accommodate the high density of input on the TABs, the 8-bit serial trigger-tower data is transmitted using the channel-link LVDS chipset `DS90CR48X` [10, 9] which serializes 48 CMOS/TTL inputs and the transmission clock onto eight LVDS channels plus a clock channel. In our application, the input to the transmitter is 60 MHz TTL ($8\times$ crossing rate) which is stepped up to 420 MHz (60 MHz$\times$7) LVDS.

Each ADF outputs 36 8-bit words: the digitized E_{t} of 16 EM and 16 hadron trigger towers and four control words. The bunch-crossing (BC) number control word indicates which accelerator crossing produced the data payload and is used throughout the system for synchronization. The frame-bit control word signifies the least significant bit or the phase of the serial data, i.e. it is always "1". There is an unused spare control word. The parity control word is the logical XOR of every other word, except the spare word, and is used to check the integrity of the data transmission.

As there are only 36 total data words transmitted, the ADF-to-TAB transmission (See Figure 1 only uses eight LVDS channels–seven data and one clock–which enables a convenient 2×4 LVDS cable design using Gore Z-PACK 2MM HM Cable Connector (2MMS02XX)[8].

While the ADF logic is 8-bit serial (60 MHz) the TAB logic is 12-bit serial (90 MHz). To cross the clock domains, the data passes through a dual-

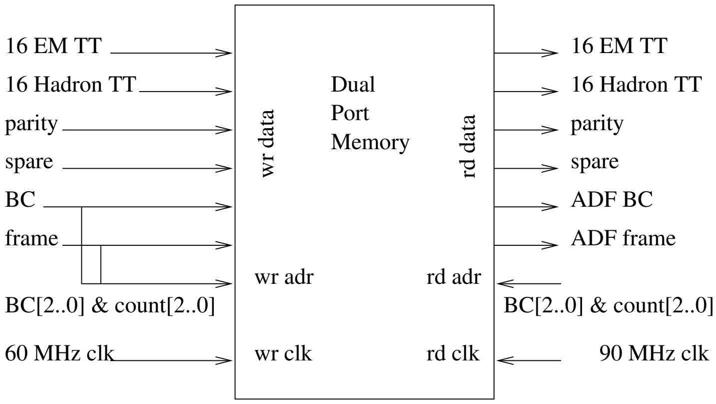

Figure 2: The ADF-to-TAB dual-port memory.

port memory with the additional four bits set to zero on the reading (90 MHz) side. These additional bits are not needed for the input data, but are required by the sliding windows algorithm, as summing many 8-bit words produces additional carry bits. The output words also takes advantage of these additional available bits.

The write address is calculated from the frame and bunch crossing words of the ADF data. The least significant bits are a bit count reset by the frame signal, while the most significant are the first three-bits of the bunch crossing number. This means that the memory is large enough to contain eight events of eight-bit serial data.

By calculating the read address in the same fashion, but from the TAB frame and and bunch crossing words, the dual-port memory crosses 60 MHz/90 MHz clock domains, maintains the correct phase of the data, and synchronizes the data to within 8 crossings all at the same time. This means the TAB timing can range between a minimal latency setting where the data is retrieved just after it is written and a maximal latency setting where the data is retrieved just before it is overwritten. If the TAB timing is outside this range, the data from eight crossing previous or later will be retrieved.

Although we used off-the-shelf components within their marketed operating conditions, we found operating 240 such links reliably was challenging. For longer cables the chipset includes a DC-balance feature and a pre-emphasis feature which we do use, and a deskew feature which we do not use. The DC-balance feature inverts the output polarity of a channel if it would

ADF crate	nominal cable length	installed cable length	pre-emphasis (Ω)
A	365 cm	365 cm	?
B	245 cm	245 cm	?
C	305 cm	500 cm	?
D	425 cm	425 cm	?

Table 1: ADF-TAB LVDS cable lengths and pre-emphasis levels.

cause too much DC current. The resistor-tunable amount of pre-emphasis injects additional current during transitions to create shorter rise-times and a larger eye-pattern at the receiving end.

As the channel-link LVDS is a serial-parallel system (high-frequency serial-data transmitted on multiple wires) channel-to-channel skew on the LVDS cables can be a problem. When the deskew feature is used, the transmitter sends a training pattern which the receiver uses to detect and correct for channel-to-channel skew. We found the deskew feature to be unreliable, and instead use cables with tightly controlled skew and leave the feature off.

Even with excellent quality cables, we found reliability to be a problem for lengths exceeding (365 cm). For this reason, we chose to purchase cables with lengths customized to the ADF crate location (Table 1). Unfortunately, we received cables that were longer than neccessary for crate C but opted to install them anyway due to the installation schedule. For the longer length cables in crate C and D, we have had to selectively replace a handful of problematic cables with shorter varieties to achieve reliable performance. Our experience has been that changing the transmitter, cable, or receiver can effect link performance, so we have extra cables installed in case of problems when modules are replaced in the future. We have also bought replacement short cables for crate C and $365cm$ cables–which are long enough in most cases–for crate D. We may opt to install these cables in the future if long-term reliability of the longer cables becomes a problem. This has not been the case in the first few months of running.

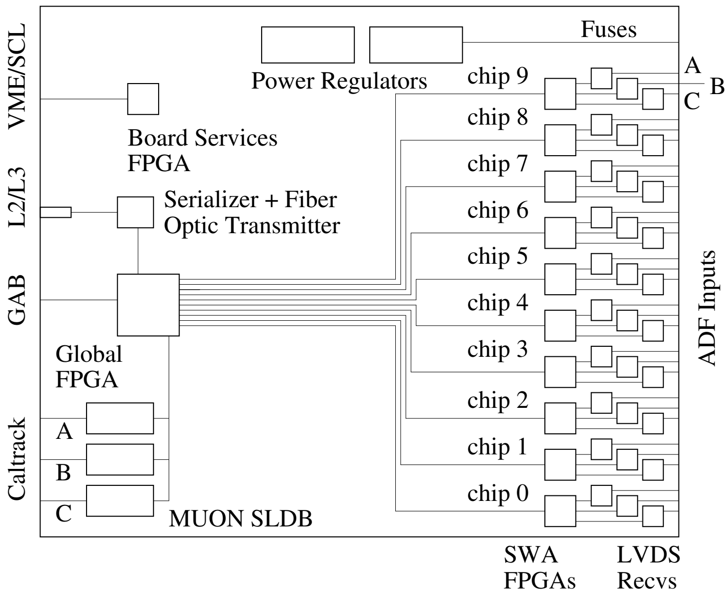

Figure 3: Schematic overview of the TAB.

2 Trigger Algorithm Board

The primary responsibility of the eight trigger algorithm boards, or TABs, is to calculate the sliding windows algorithm to identify EM, Jet and Tau physics objects. The TAB is a double-wide 9U card designed for a custom backplane (see Figure 3) ADF data is received from the back of the card on LVDS receivers and sent to 10 TAB sliding windows algorithm FPGAs for processing. The algorithm output is sent to a single TAB global FPGA. The global FPGA calculates regional sums and sends the results out the front of the board to the global algorithm board (GAB). See Table 2 for the part numbers and relevant data sheets of the major TAB components.

The global FPGA also sends 3 copies of the Jet and EM object information to the Caltrk system for processing using Muon Serial Link Daughter Cards [6]. Upon receiving a Level-1 Accept from the D0 trigger, the TAB global chip outputs data for the Level-2 trigger and Level-3 readout on a serial fiber-optic link.

Low-level board services are provided by the TAB cyclone chip. Upon powering-up the TAB, the cyclone FPGA is the only FPGA powered and

TAB Component Name	Vendor	Part Number
LVDS Channel Link Receiver[10]	National	DS90CR484VJD
TAB SWA[4]	Altera	EP1S20F780C7
TAB Global[4]	Altera	EP1S20F780C6
TAB Cyclone[3]	Altera	EP1C6Q240C7
Cyclone Configuration[2]	Altera	EPCS4SI8
Clock Distribution[5]	Cypress	CY29948AI
L2/L3 serializer[1]	Agilent	HDMP1022
L2/L3 transmitter[11]	Stratos Lightwave	M2T-25-4-1-L

Table 2: Major TAB components

configured by an on-board serial configuration device (Table 2). The specialized serial VME/SCL control link is received by the cyclone chip. When instructed over VME, the cyclone chip is capable of powering up and configuring the remaining stratix chips. The cyclone also receives SCL timing information fanning out the 132 ns detector clock to the stratix FPGAs using an on-board clock distribution device. Serial VME commands are forwarded to the stratix chips for setting operating parameters, reading diagnostic memories, and reading and setting status registers.

The cyclone chip, global chip, and SWA chip 4 have JTAG connections allowing for direct configuration and use of a vendor-software supplied signal-tap feature which builds a virtual logic-analyzer–complete with trigger–into the chip.

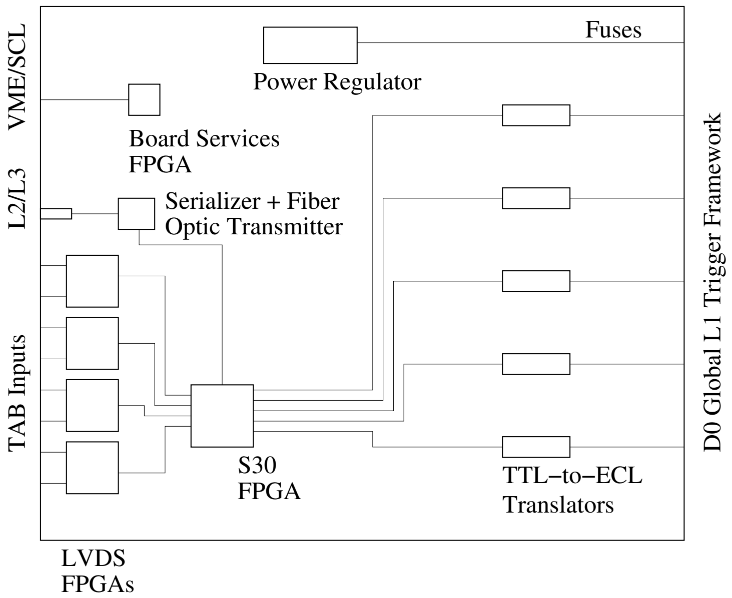

Figure 4: Schematic overview of the GAB.

3 GAB

The global algorithm board (see Figure 4) receives data containing regional sums of jet, EM, and tau physics objects calculated by the trigger algorithm boards (TABs) and produces a menu of trigger terms, called and/or terms in the DØ level 1 framework. Like the TAB, the GAB is a double-wide 9U card designed for a custom backplane. The part numbers and data sheets for the major components of the GAB are detailed in Table 3.

Four LVDS receivers each receive the output of two TABs, synchronizing the data to the GAB 90 MHz board clock using a dual-port memory. Both the TABs and the GAB calculate a run signal which goes high only after an SCL init signal has gone low and a beginning of turn signal is received. On the LVDS FPGAs, the memory is written to the circular memory as soon as the TAB run signal is received. It is read as soon as the GAB run signal is calculated. By setting the GAB SCL timing parameters to add slightly more latency, the TAB data will arrive just before the GAB LVDS FPGAs attempt to read it.

The synchronized TAB data from all four LVDS FPGAs is forwarded to

GAB Component Name	Vendor	Part Number
GAB LVDS[4]	Altera	EP1S20F780C6
GAB S30[4]	Altera	EP1S20F780C7
GAB Cyclone[3]	Altera	EP1C6Q240C7
Cyclone Configuration[2]	Altera	EPCS4SI8
Clock Distribution[5]	Cypress	CY29948AI
L2/L3 serializer[1]	Agilent	HDMP1022
L2/L3 transmitter[11]	Stratos Lightwave	M2T-25-4-1-L
TTL-to-ECL translator[7]	Fairchild	100324SC15

Table 3: Major GAB components

the single GAB S30 FPGA which calculates and/or terms, and sends them to the trigger framework through TTL-to-ECL converters. Their are five 16-bit outputs on the GAB, but only four are used by the framework.

Much like the TABs, upon receiving a L1-A, the GAB S30 outputs and/or term data for Level 3 readout on a serialized fiber-optic link. Also as on the TABs, the cyclone FPGA provides low-level board services. On the GAB, the cyclone chip, global chip, and LVDS chip 4 have JTAG connections.

cable:	C		B				A		
$\Delta\phi$: $\Delta\eta$	-2	-1	0	1	2	3	4	5	6
0	0	1	2	3	4	5	6	7	8
1	9	10	11	12	13	14	15	16	17
2	18	19	20	21	22	23	24	25	26
3	27	28	29	30	31	32	33	34	35

Table 4: Explanation of TAB SWA trigger-tower index number.

4 TAB/GAB trigger path

Eight TABs cover the entire 40×32 η-ϕ trigger-tower grid, with each TAB handling the entire η range for a wedge of four ϕ. This allows each TAB to concentrate the data by summing objects in three eta ranges (North, Central, and South) before sending to the next stage. Each of the TAB sliding windows algorithm chips handles a 4×4 η-ϕ grid. In what follows, when referring to the entire grid we use coordinates $\eta[0, 39]$ and $\phi[0, 31]$ but when referring to a single TAB SWA chip we use local coordinates $\Delta\eta[0, 3]$ and $\Delta\phi[0, 3]$ where

$$\phi = 4 \times (\text{TAB } n) + \Delta\phi, \quad (1)$$

$$\eta = 4 \times (\text{SWA chip } n) + \Delta\eta. \quad (2)$$

ADF data from 3 cables, labeled A,B, and C (See Figure 3) is sent to each of the TAB SWA chips. Each ADF cable contains a 4×4 η-ϕ grid of EM and hadron TT data. While the B cable provides the data for 4×4 η-ϕ grid for which the SWA chip is responsible, the A and C cables provide the upper and lower ϕ-overlap regions. Also, each TAB SWA FPGA receives η-overlap data from neighboring SWA FPGAs. The 36 trigger towers used by the algorithm from the A,B, and C cables are given a trigger tower index as shown in Table 4.

The TAB/GAB algorithm calculation uses serial logic, using data words with the serial format described in Table 5 along the path illustrated in Figures 5.

As the serial raw TTs are only 8-bits, the upper four bits are set to zero. The output of the first stage of the calculation is the index of the maximum threshold $\tau_{\max}$ object found at each $(\eta, \Delta\phi)$ point, arranged serially in ϕ, thus requiring only 4 serial words per algorithm or 12 words total.

The 12 data words from each of the 10 SWA chips is received by the TAB global chip (see Figure 5 and summed across three eta regions: North, Central, and South. The global chip calculates the 3-bit count $n(i)$ of objects exceeding threshold i for each region R and $\Delta\phi$ are arranged serially in threshold i. Four 3 algorithms, 3 regions, and 4 $\Delta\phi$ this leads to 36 output words per TAB.

The GAB LVDS receivers accumulate $36 \times 8 = 288$ serial words which are passed without changes to the GAB S30 chip. The GAB S30 chip produces the and/or terms based on calculation of the regional counts $n(i)$.

In addition to basic EM, Jet, and tau objects there are some additional data words passed through the trigger path. Each sliding windows chip produces an EM isolation bit based on EM outer ring and Hadron isolation requirements, for 16 bits total. This requires two additional 12-bit serial words, arranged in increasing ϕ then increasing η as

$$\text{word } n = floor(\Delta\eta/3), \tag{3}$$

$$\text{bit } n = 4 \times (\Delta\eta mod 3) + \Delta\phi. \tag{4}$$

At the TAB global chip, the EM isolation bits are ORed across three eta regions, using the first four bits of three words: one word for each η-region R, with bit $n = \Delta\phi$.

For E_t -related triggers, each SWA chip sums E_t across four $\Delta\eta$ for each $\Delta\phi$. On the global chip, these are summed across η and $\Delta\phi$ with an unsigned sine and cosine scale factor to obtain a single E_t , E_x , and E_y . On the GAB S30, the appropriate sign is ascertained based on the TABs ϕ coordinates for calculation of total E_t , E_x , and E_y .

The TAB SWA chips also count the number of *single* towers with EM+hadron E_t greater than three VME programmable thresholds. The sums saturate at 4-bits (15). The global chip sums the three 4-bit counts from all ten SWA chips to produce three 8-bit sums. These three sums are placed in the 8-bits of extra space from the EM isolation words described above. The GAB S30 chip sums the data from all eight TABs and fires a trigger if the count exceeds an additional count threshold.

stage	η range	ϕ range	serial word format											
			11	10	9	8	7	6	5	4	3	2	1	0
raw trigger towers	$\eta[0,39]$	$\Delta\phi[-2,6]$	0				EM/HD trigger tower							
TAB SWA output	$\eta[0,39]$	$\Delta\phi[0,4]$	$\tau_{\text{max}}^{\eta,\Delta\phi=3}$			$\tau_{\text{max}}^{\eta,\Delta\phi=2}$			$\tau_{\text{max}}^{\eta,\Delta\phi=1}$			$\tau_{\text{max}}^{\eta,\Delta\phi=0}$		
TAB global output	$R[\text{N},\text{C},\text{S}]$	$\Delta\phi[0,4]$	$n_7^{R,\Delta\phi}$		$n_6^{R,\Delta\phi}$		$n_5^{R,\Delta\phi}$		$n_4^{R,\Delta\phi}$		$n_3^{R,\Delta\phi}$		$n_2^{R,\Delta\phi}$	
GAB lvds recv	$R[\text{N},\text{C},\text{S}]$	$\phi[0,39]$	$n_7^{R,\phi}$		$n_6^{R,\phi}$		$n_5^{R,\phi}$		$n_4^{R,\phi}$		$n_3^{R,\phi}$		$n_2^{R,\phi}$	
GAB S30 recv	$R[\text{N},\text{C},\text{S}]$	$\phi[0,39]$	(same)											

Table 5: The raw trigger data consists of an EM and Hadron trigger-tower for each position (η,ϕ). The TAB sliding windows algorithm output $\tau_{max}^{\eta,\phi}$ is the 3-bit [0,7] index of the largest $E_{\rm t}$ threshold met by an object at position (η,ϕ). The first threshold has index 1 while an index of 0 means no threshold was exceeded. The TAB global output $n_i^{R,\phi}$ is the 2-bit [0,3] number of objects exceeding *only* the threshold with index i. For three or more such objects, $n(\mathrm{i}) = 3$. Note that information about the first threshold (i.e. $n(1)$) has been discarded.

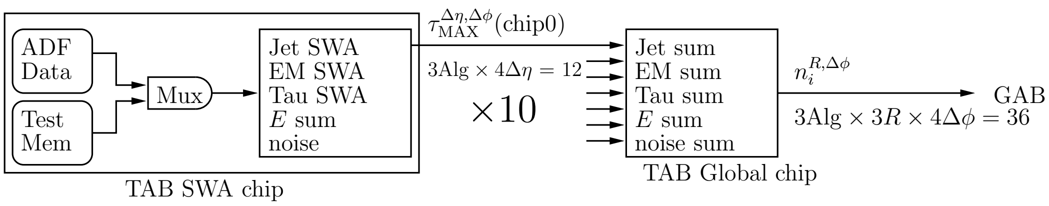

Figure 5: The TAB trigger path.

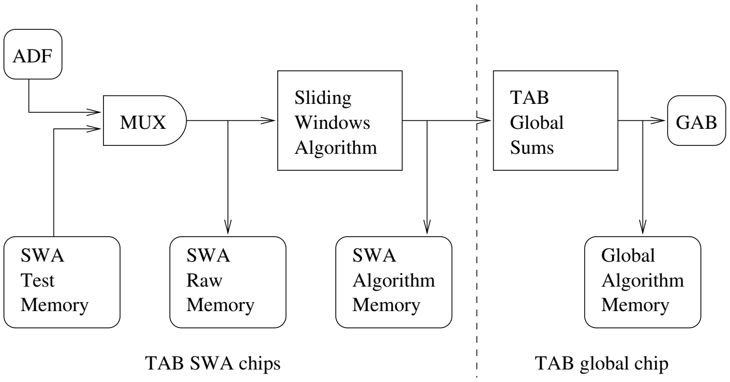

Figure 6: TAB diagnostic memories.

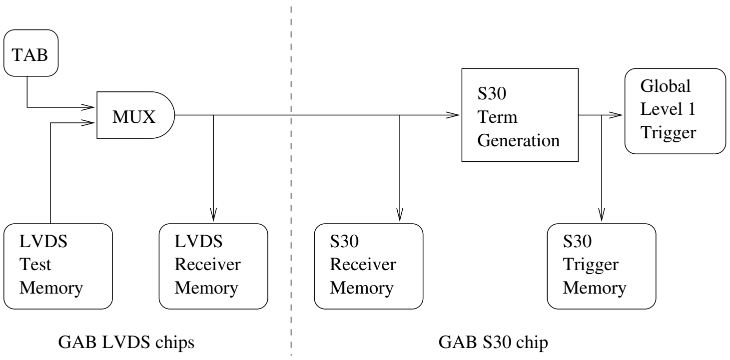

Figure 7: GAB diagnostic memories.

5 TAB/GAB diagnostic memories

The TAB and GAB modules have a series of VME-readable diagnostic memories designed to capture data from each step of the algorithm calculation. This is extremely useful for debugging and understanding the hardware performance, as well as long-term system monitoring. The memories are written whenever a diagnostic pulse signal is sent from the VME/SCL card. See Figure 6, Figure 7, and Table 6. The format of the diagnostic memory words is generally the same as the corresponding 12-bit serial word with 4 upper bits set to zero to make a 16-bit VME word.

Both the TAB and GAB have VME-writable test input memories which

Memory	Write	Depth	Contents
TAB SWA Raw	DP	32	EM/Hadron TTs
TAB SWA Algorithm	DP	32	E Sums, Jet/Em/Tau Objects ($\tau_{\text{MAX}}^{\Delta\eta,\Delta\phi}$) EM isolation, Noise
TAB Global Algorithm	DP	32	E Sums, Regional Sums ($n_i^{R,\Delta\phi}$) EM isolation, Noise
TAB Global L2/L3	L1A	1	E EM/Hadron TTs L2 object masks (unused)
GAB LVDS Receiver	DP	32	E Sums, Regional Sums ($n_i^{R,\phi}$) EM isolation, Noise
GAB S30 Receiver	DP	1	(copy of above)
GAB S30 Trigger Terms	DP	32	E Sums, Trigger Terms.
GAB S30 L2/L3	L1A	1	E Sums, Trigger Terms.

Table 6: Overview of TAB/GAB diagnostic memory. "DP" stands for diagnostic pulse, "L1A" is a Level-1 Accept. The variables under contents are described in the TAB/GAB trigger path section.

allow arbitrary patterns to be used in the place of the incoming data from the ADF or TAB cards.

For debugging the L2/L3 readout, both TAB and GAB cards save an exact copy of the most recent event of data sent to L2/L3.

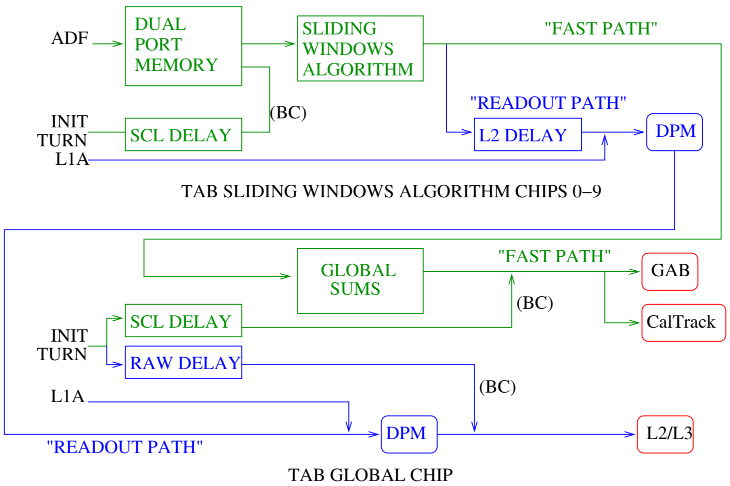

Figure 8: Timing and Readout

6 TAB/GAB Timing and Readout

The timing and readout of the TAB and GAB modules are interrelated. One can think of the synchronous trigger path as the *fast path*, which each event passing through the system with minimal latency. The *readout path* is also synchronous, but must hold a fixed length of time for a possible Level-1 Accept from the global Level-1 trigger decision.

The timing information throughout the system is containing in the turn number and bunch-crossing numbers. The turn number is reset after an SCL init and incremented every turn. The BC number is reset every turn and incremented every 132 ns bunch-crossing clock tick.

The BC number used to extract data from the ADF/TAB dual-port memory is delayed by an amount specified by the VME parameter “scl delay”. For minimal latency, this is adjusted so that data is extracted for a particular BC just after it has been written by the ADF. At the global chip, a similar “scl delay” parameter adjusts the BC associated with the data as it is processed. Status registers indicate whether or not the BC numbers at each stage agree.

On the TAB SWA chips the Level-2/Level-3 buffer is kept in a delay pipeline with a length determined by the parameter “L2 delay”. If a Level 1 accept is asserted as an event is leaving the pipeline it is immediately stored in dual-port memory for slow transport to the global chip. The size of the memory is actually several events long, but the additional buffers are not needed as the D0 framework guarantees the next Level 1 accept will not be sent until all the data has been received on VRBs.

When the global chip receives the level 1 accept it stores the value of the BC number delayed by a VME-parameter "raw delay". It then begins polling the SWA chips. Once all the SWA chip data has been recovered, the data is transmitted to the serializer and fiber-optic transmitter.

On the GAB, the mechanism is much simpler as all of the data is present on the one S30 chip. Both the LVDS chips and S30 chips have an "scl delay" parameter set to provide minimal latency between stages. The pipeline with length "L2 delay" is located on the S30 chip, with transmission to the serializer and fiber-optic transmitter commencing upon receipt of an Level 1 accept. Some additional latency has been added before transmission to make the GAB timing similar to the TAB timing.

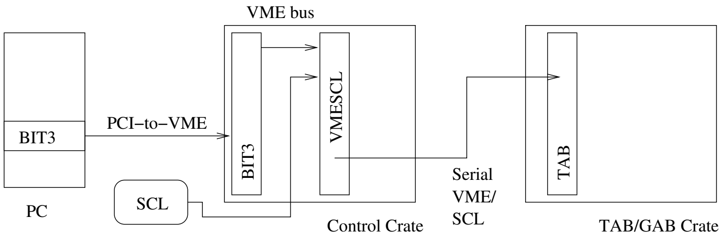

Figure 9: TAB/GAB VME/SCL control path

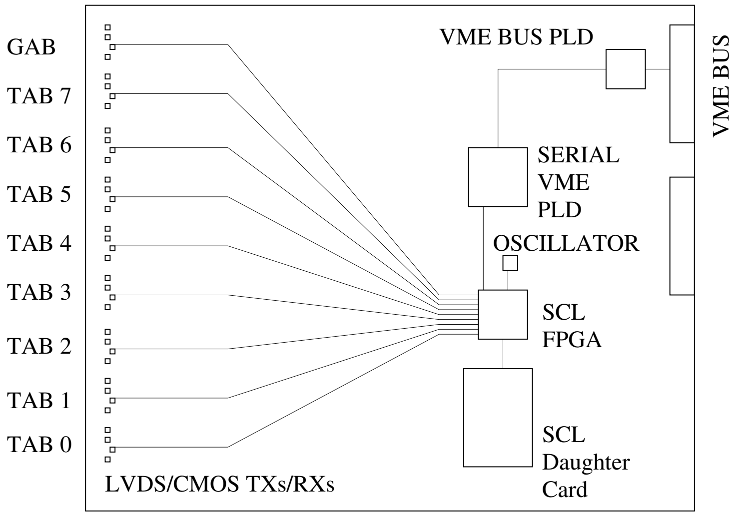

Figure 10: Schematic overview of the VMESCL.

7 VMESCL and TAB/GAB control path

Because of the high-density of inputs to the TAB and GAB modules, a custom control path (Figure 9)is provided by the VMESCL module (Figure 10). The VMESCL converts VME commands from a standard VME backplane into a serial format for transmission to the TAB and GAB modules. It also relays DØ serial command link (SCL) signals from an SCL daughter card to the TAB and GAB modules.

The D0 SCL used by the TAB and GAB modules are summarized in Table 8. The init signal resets the state to prepare for data taking. The 132 ns bunch-crossing clock is synchronous with the detector interactions. The

VMESCL Component Name	Vendor	Part Number
Configuration Device	Altera	EPC1441PC8
SCL FPGA	Altera	EP1K50QC208-2
VME Bus PLD	Altera	EPM3128ATC144-7
Serial VME PLD	Altera	EPM3256ATC144-7
Crystal Oscillator	Vectron	VCC1-B3A 53M000
Clock Multiplexer	IC SYSTEMS	ICS581G-011
LVDS-CMOS Driver	National	DS90LV047ATMTC
LVDS-CMOS Receiver	National	DS90LV048ATMTC

Table 7: Major VMESCL components

signal	description
init	initialize
turn	begin of turn (i.e. first of 159 bunch crossings)
clk	the 132 ns bunch-crossing clock.
accept	data from BC synchronous with this signal is accepted at Level 1.
BC num	derived count from 1 to 159 starting after turn.
run	derived high once init goes low and turn is seen.

Table 8: SCL signals and derived signals.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	Chip				Group				0	0	0	0	0	0	0

Serial Frame Word

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Address															

Serial Address Word

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Data															

Serial Data Word

Figure 11: Serial VME Protocol

turn signal marks the first of 159 bunch crossings. On the TAB and GAB modules, the bunch-crossing number is incremented every clock tick and reset by the turn signal. The run signal is used in the TAB/GAB synchronization scheme. It goes low on init, and high on the first turn after an init signal has gone low.

The serial VME protocol between the VMESCL and TAB/GAB modules is shown in Figure 11. The VME address space is broken into chip, a unique number for each FPGA on a module, and group, a general high-order grouping of addresses (e.g. into timing parameters). The leading one on the serial frame word signals the beginning of a serial word which lasts for 16-bits. The serial data word contains the data specific to the chip, group, and serial address.

<table>
<tr><th>RF Tick Number</th><th>11</th><th>11</th><th>11</th><th>11</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="15"></td><td>0</td></tr>
<tr><td>1</td><td>0</td><td colspan="7">Jet($\Delta\phi = 0$)</td><td>0</td><td colspan="7">EM($\Delta\phi = 0$)</td></tr>
<tr><td>2</td><td>0</td><td colspan="7">Jet($\Delta\phi = 1$)</td><td>0</td><td colspan="7">EM($\Delta\phi = 1$)</td></tr>
<tr><td>3</td><td>0</td><td colspan="7">Jet($\Delta\phi = 2$)</td><td>0</td><td colspan="7">EM($\Delta\phi = 2$)</td></tr>
<tr><td>4</td><td>0</td><td colspan="7">Jet($\Delta\phi = 3$)</td><td>0</td><td colspan="7">EM($\Delta\phi = 3$)</td></tr>
<tr><td>5</td><td colspan="15"></td><td>0</td></tr>
<tr><td>6</td><td colspan="16">Parity Word</td></tr>
</table>

Figure 12: Format of TAB global output to Caltrk.

8 TAB/GAB data to Caltrk

The Caltrk system matches calorimeter objects to tracks. The TAB global chip sends EM and Jet object data for each ϕ position to the Caltrk system. Each TAB sends three identical copies of the same data (for use as overlap regions) to the Caltrk system using three Muon Serial Link Daughter cards. The daughter cards serialize seven 16-bit words per crossing for transport to the Caltrk system.

The format of the Caltrk data sent by the TAB is shown in Figure 12. Each Jet or EM word is composed of 7 bits. A bit i is high provided a Jet or EM object exceeds threshold i for any allowed η at the specified $\Delta\phi$ position. The allowed eta range is $\eta[10:29]$ corresponding to physical $|\eta| < 2.0$. Note that this is an *inclusive* formulation with respect to thresholds.

The seventh word send to the Caltrk is a parity word calculated by the daughter card.

9 Performance

BER limits: coming soon. Latency: coming soon. Readout dead-time: coming soon.

References

[1] Agilent Technologies. *Low Cost Gigabit Rate Transmit/Receive Chip Set with TTL I/Os*, 1999.

[2] Altera. *Configuration Handbook, Volume 1*, 2006.

[3] Altera. *Cyclone Device Handbook, Volume 1*, 2006.

[4] Altera. *Stratix Device Handbook*, 2006.

[5] Cypress. *2.5V or 3.3V, 200-MHz, 1:12 Clock Distribution Buffer*, 2002. `38-07288 Rev. B`.

[6] D0. *I/O Specification for Serial Transmitter Daughter Board (PCB-0140-XMIT)*, 1999.

[7] Fairchild. *Low Power Hex TTL-to-ECL Translator*. `DS009878`.

[8] Gore. *Gore Z-PACK 2MM HM Cable Connector (2MMS02XX)*.

[9] National Semiconductor. *LVDS Owners Manual*, 2000.

[10] National Semiconductor. `DS90CR483` /`DS90CR484` *48-Bit LVDS Channel Link SER/DES – 33 - 112 MHz*, 2002.

[11] Stratos Lightwave. `M2T-25-4-1-L` *Optical Gigabit Ethernet/Fibre Channel*. `3059.04`.