

## CMS CALORIMETER LEVEL 1 TRIGGER CONCEPTUAL DESIGN

J. Lackey, S. Dasu, M. Jaworski, D. Panescu,  
W. H. Smith, W. Temple  
*University of Wisconsin*

### 1. Introduction

#### 1.1 CMS Trigger and Data Acquisition System

The CMS trigger and data acquisition system [1] is designed to operate at the nominal LHC design luminosity of  $10^{34}$ , where an average of 20 inelastic events occur at the beam crossing frequency of 25 nsec. This input rate of  $10^9$  interactions every second must be reduced by a factor of at least  $10^7$  to 100 Hz, the maximum rate that can be archived by the on-line computer farm. CMS has chosen to reduce this rate in two steps. The first level stores all data for approximately 3 microseconds, after which no more than a 100 kHz rate of the stored events is forwarded to the second level. This must be done for all channels without dead time. The second level is provided by a subset of the on-line processor farm, and passes a fraction of these events for more complete processing by the remainder of the on-line farm. During the 3 microseconds of the level 1 trigger processing time, trigger decisions must be developed that discard a large fraction of the data while retaining the small portion coming from interactions of interest.

#### 1.2 Level 1 Calorimeter Trigger System Function

The CMS level 1 trigger decision is based in part upon local information from the level 1 calorimeter trigger about the presence of objects such as photons, electrons, and jets, as well as global sums of  $E_t$  and missing  $E_t$  (to find neutrinos). Each of the objects, such as electrons, muons and jets, are required to pass a series of  $p_t$  or  $E_t$  thresholds, which are used in making the Level 1 Trigger Decision.

#### 1.3 Electron/Photon Trigger

The electron/photon trigger is based on the recognition of a large and isolated energy deposit in the electromagnetic calorimeter. The isolation requirements ask for a small amount of energy deposited in the calorimeter trigger towers surrounding the ECAL energy cluster, as well as a small hadronic energy deposited in the HCAL in the cluster region. There are different thresholds for inclusive electrons/photons, dileptons, and for very high  $E_t$  electrons. The isolation cuts are relaxed and finally eliminated for

triggers with increasing  $E_t$  thresholds. The tightest set of cuts is used only to push the electron and photon thresholds to the lowest possible  $E_t$  values.

The algorithm implemented in the hardware design involves two separate cuts on the longitudinal and transverse isolation of the ECAL energy deposit. The first cut involves the hit tower HCAL to ECAL energy ratio,  $H/E < 0.05$ . A second cut requires transverse isolation, i.e. a cut on a sum of HCAL transverse energies in the nearest eight towers surrounding the hit tower,  $H1 < 2.0$  GeV. In order to reduce the number of bits of information exchanged between electronics cards we limit the dynamic range of neighboring tower HCAL information to 2-bits. Overflows of both the 8-bit scale for ECAL and central HCAL towers, and the 2-bit scale for neighboring HCAL towers are treated as maxima. The electron/photon trigger algorithm is illustrated in Figure 1.

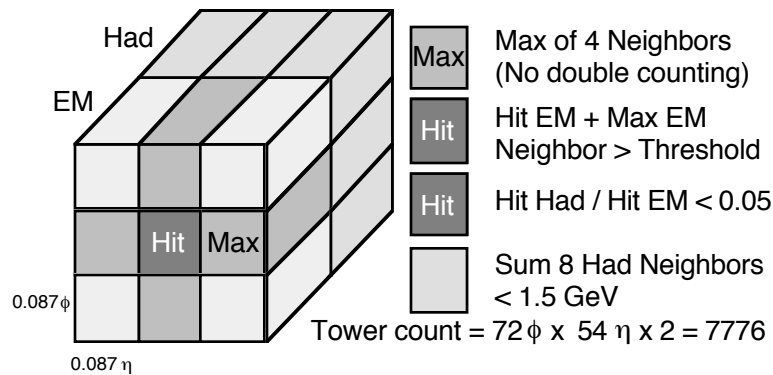


Figure 1. Level 1 Electron Trigger Algorithm.

#### 1.4 Jet and Isolated Hadron Triggers

The jet triggers are based on sums of  $0.35 \times 0.35$  eta-phi regions using a dynamic range that covers energies up to about 250 GeV. A second category of jets that exhibit the presence of isolated hadrons can be found by testing the high  $E_t$   $0.35 \times 0.35$  regions to determine if a high fraction of the total observed energy is contained in a single trigger tower. Such regions are candidates for isolated hadrons. These isolated hadron jet triggers are used in detecting tau's. Jet triggers will be used in combination with electron/photon, muon and other jet triggers. For these combinations of triggers, lower  $E_t$  jet thresholds will be employed.

#### 1.5 Missing $E_t$ and Sum $E_t$ Triggers

Neutrino identification consists of calculating the event missing  $E_t$  vector and testing it against a threshold. The calorimeter trigger calculates both sums of  $E_t$  and missing  $E_t$ . The transverse energy vector components are calculated from the 8-bit compressed-scale digitized HCAL and ECAL pulse

heights converted to a linear scale with a 10-bit dynamic range, and multiplied by entries in lookup tables containing the tower angular coordinates. The HCAL and ECAL sums are then combined into single tower sums. The tower sums over threshold are routed through the digital summing networks. The  $E_t$  trigger is not used as a standalone trigger, but only in combination with other triggers. When pre-scaled by factors of 1000 or more for  $E_t$  above 100 GeV, it is also used to check trigger efficiency and measure the  $E_t$  spectrum.

## 1.6 Luminosity

Operating the LHC requires continual non-intrusive measurement of relative luminosity and periodic intrusive measurement of the absolute luminosity using an absolute measurement. The non-intrusive measurement of relative luminosity can be done with signals accumulated in the level 1 calorimeter trigger system. The level 1 calorimeter trigger system provides four thresholds on the energy in each  $0.7 \text{ delta-eta} \times 0.7 \text{ delta-phi}$  calorimeter tower. The lowest threshold is set to have a rate 100 kHz. The granularity is set by the need to have a significant rate and the ability to observe localized "hot spots" in the calorimeter. The thresholds are determined by the nominal precision of 1% every 1/10 second, implying a count rate of 100 kHz. In order to fulfill this requirement, the following data requirements are set:

Provided every 0.1 sec:

Total number of  $0.7 \times 0.7$  towers with energies above each of 4 individually programmable thresholds. This consists of 4 16-bit numbers.

Time interval in crossings over which the number of towers above were accumulated.

Provided every 5 minutes:

Number of times each of 2,835 crossings has a  $0.7 \times 0.7$  tower over a threshold, independently programmable from the four listed above. If a particular crossing has  $N$   $0.7 \times 0.7$  towers that exceed this threshold in a single accelerator cycle, then the sum for this crossing is incremented by  $N$ . This enables the luminosity per crossing to be accumulated for analysis of occupancy and deadtime effects.

Number of times each  $0.7 \times 0.7$  tower region exceeded the single threshold listed above.

Time interval in complete accelerator cycles over which the number above was accumulated.

This information is provided over the DAQ data transmission network from the DAQ processors in the level 1 calorimeter trigger crates. The rate of transmission of this data will impose a negligible burden on the DAQ system.

## 1.7 CMS Level 1 Trigger System

The CMS Level 1 Trigger system is a deadtimeless processor that operates on a subset of the full data collected from each LHC beam crossing every 25 nsec. The Level 1 Trigger System analyzes the subset of data within a fixed amount of time to determine if the data collected from the beam crossing is to be discarded or held for further processing by the higher level trigger systems.

The Level 1 Trigger System is comprised of the front end electronics that generate trigger primitives on the detector, and the Level 1 processing logic in the electronics barracks. The Level 1 logic and front end electronics are interconnected both electrically and optically to form multiple pipelined data streams, which come together at the Level 1 Global Logic to generate a Level 1 Accept/Reject for the detector.

## 1.8 Input Data to the Level 1 Calorimeter Trigger

The front end electronics will operate and sample the input calorimeter pulse height at 80 MHz and use this information to produce a trigger tower sums at 40 MHz. These trigger sums will emerge in 13 (25 nsec) crossings after the input pulse height reaches the front end electronics. The front end electronics for the baseline design is the fully digital FERMI readout system [2].

The calorimeter level 1 trigger system receives digital trigger sums via optical fibers from the front end electronics system, which transmits energy on an eight bit compressed scale. The data for two HCAL or ECAL trigger towers for the same crossing will be sent on a single fiber in eight bits apiece accompanied by five bits of error detection code. The compressed scale of the trigger tower data will be derived from memory lookup tables within the front end system. Programmable tables will enable full flexibility to modify the compression algorithm as required. The Trigger system will use memory lookup tables to decode the 8-bit compressed scale data into a linear scale with a 10 bit dynamic range in the adder tree.

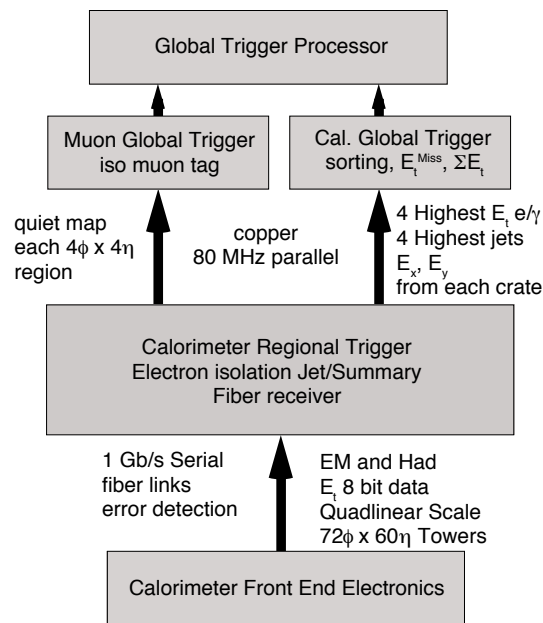
There is a 1:1 correspondence between HCAL and ECAL trigger towers. The trigger tower size is equivalent to the HCAL physical towers,  $.0873 \times .0873$  in  $\eta \times \phi$ . The  $\phi$  size remains constant in  $\Delta\phi$  and the  $\eta$  size remains constant in  $\Delta\eta$  down to an  $\eta$  of 2.1, beyond which the  $\eta$  size doubles. There are 3888 total ECAL trigger towers and 3888 total HCAL trigger towers from  $\eta = -2.6$  to  $\eta = 2.6$  ( $54 \times 72$   $\eta$ - $\phi$  divisions).

The very forward calorimeter (VFCAL) region will sum electromagnetic and hadronic energy for one tower into a single trigger sum. The trigger towers

will be divided more coarsely in phi and eta, possibly in 18 phi bins by 4 eta bins from  $|\eta| = 2.6$  to 4.0 and then 9 phi bins by a single eta bin to  $|\eta| = 5.0$ . Beyond  $|\eta| = 3.0$ , where the HCAL ends, the front end calorimeter electronics will have a different interface to the different calorimeter technology for the VFCAL, but will still provide the same type of single tower outputs. Therefore we estimate at most 3888 ECAL, 3888 HCAL and 162 VFCAL trigger channels for a total of 7938 trigger channels carried on 4050 fibers.

## 1.9 Overview of Calorimeter Level 1 Trigger Design

This document outlines a conceptual design for the level 1 calorimeter trigger. A block diagram of the CMS Level 1 Calorimeter trigger system is shown in Figure 2. General considerations that have been emphasized in this conceptual design include access to components, power, space and cooling requirements, diagnostics, efficiency and performance information, backplane traffic and timing, DAQ and timing interface, and I/O connections. The design is implemented using off-the-shelf technology wherever feasible, e.g., 10KH and ECLiPS ECL, TTL, and CMOS components. ASICs are used only where fully justified. The logic design maximizes flexibility and programmability by using memory lookup tables.



**Figure 2.** Level 1 Calorimeter Trigger Overview and details of one crate.

## 2. Calorimeter Regional Trigger Crates

There will be 19 calorimeter processor crates covering the full detector. Eighteen crates are dedicated to the barrel and two endcaps. These crates are filled out to an eta of 2.6, with partial utilization between 2.6 and 3.0. The remaining crate covers both Very Forward Calorimeters. Our intention at

present is to locate two crates each in Aleph-style racks. The remaining rack front panel space will be occupied by fans, heat exchangers, and crate power supplies. Front panels will be used at all card locations to provide an enclosed environment for the chilled air.

The crate, shown schematically in Figure 3, is based on standard Eurocard hardware. It is likely, however, that construction will require some custom fittings. The height is 9U and the depth approximately 700mm, as determined by the front and rear card insertion. The Aleph rack dimension (900mm deep) can handle the crate depth with some reserve for cabling, plumbing, and other services. The backplane is completely custom with a full 9U height. The top 3U is reserved for a 32 bit VME interface. The remaining 6U is used for the high speed data paths between individual cards. The front section of the crate is designed to accommodate 280mm deep cards, leaving the major portion of the volume for 400mm deep rear mounted cards.




Figure 3. Schematic view of a typical Calorimeter Level 1 Regional crate.

Power supplies will be mounted in a separate chassis either above or below each crate. They will be located in the forward 280mm of the volume in consideration of the lower heat load (per unit area) of the forward cards. It might be desirable to put the supplies above the crates just to keep the cards closer to the heat exchangers. Testing is required to determine whether this up and down arrangement will be successful. On board DC - DC converters are also being considered as possible solution to power distribution.

## 2.1 Enumeration of Card Types

The majority of cards in the Calorimeter Level 1 Regional Processor Crates, encompassing three custom board designs, are dedicated to receiving and processing data from the calorimeter. There are eight rear mounted Receiver cards, eight front mounted Electron Isolation cards, and one front mounted Jet Summary card for a total of 17 cards per crate dedicated to processing data from the calorimeter.

In addition, there are several support cards. The first of these is a readout crate controller and communication module (ROC) provided by the DAQ group. The second is a crate environment monitor (CEM) which may or may not be a commercial board. The decision to purchase or create a custom design will depend on the final requirements of the environmental monitoring system. Finally, there will be a card dedicated entirely to clock distribution and logging status for the cards in the data processing path. This card, the Local Timing, Trigger and Control card (LTTC) will interface to the detector Timing, Trigger and Control (TTC) distribution system via the timing receiver ASIC produced by the TTC group [3].

## 2.2 Crate Backplane

As mentioned above, the backplane is a monolithic, custom, 9U high printed circuit board with front and back card connectors. The top 3U of the backplane holds 4 row (128 pin) DIN connectors, capable of full 32 bit VME. The first front slot of the backplane will, however, use three row (96 pin) DIN connectors in the P1 and P2 positions with the standard VME pinout. Thus, a standard VME system module can be inserted in the first front station with a form factor conversion between the first slot and remaining slots performed on the custom backplane. The second front card position may also be designed with standard 96 pin DIN connectors if the crate monitor requires them. VME terminations will reside on the backplane.

The bottom 6U of the backplane, in the data processing section of the crate, utilizes a single high speed, controlled impedance, connector for both front and rear insertion at each card position. The design is based around a 340 pin connector, by AMP Inc, to handle the high volume of data transmitted from the Receiver cards to the Electron Isolation and Jet Summary Cards.

The AMP 340 pin connectors have a powerplane between every four signal pins, producing a stripline impedance of  $50\Omega$ . Separate high current contacts, provided for the powerplane connections, will be used to transmit power to the boards. The connectors are housed in a cast aluminum shell that doubles as a board stiffener. The electrical characteristics of the connectors are excellent, allowing rise times in the sub-nanosecond region with very low crosstalk. Press fit contacts are available for use on the backplane. The mating force required for the 340 pin AMP connector is approximately 190N. The four row DIN connector requires an additional 109N. For comparison, a standard 9U, three connector system requires a total of 244N. Therefore, the mating force is on the order of what is encountered in a standard 9U Eurocard system.

### 2.2.1 Number of Slots

The front and rear insertion of cards in the data processing section of the crate was chosen to allow greater separation between cards and to provide a more protected environment for the fibers connected to the rear mounted Receiver

cards. The increased separation will promote better cooling of the cards, and will enable a wider selection of front panel components. The staggering of the slots between front and rear cards, shown in Figure 4, is as much a result of the style of connector selected as the fact that piggybacking of connectors is inappropriate in this situation. Almost half of the signals entering the Electron Isolation board come from neighboring Receiver cards.

The spacing between cards, in the data processing area, is 3.05cm front or rear, with a 1.52cm stagger front to rear. Therefore, the eight Receiver cards, eight Electron Isolation cards, Jet Summary card, and Clock distribution card occupy 18 slots with a span of only 27.43cm across the front of the crate. The remaining 13.21cm will be allocated to a DAQ Readout card (ROC) (4.06cm), Crate Environment Monitor (CEM) (2.03cm), and a transition region for the change in form factor between the standard and non-standard VME. This is a total of 20 slots. The unused space amounts to about 5.08cm after allowing for the transition region. This represents a dedicated contingency for future additions or changes to the design.




Figure 4. Top view of region crate backplane.

### 2.2.2 Card Order

As shown in Figure 3, the first two or three slots in the crate will house the Crate Readout Card (ROC) and Crate Environment Monitor (CEM). The remaining crate volume is used for the data processing cards. The first card after the transition region will be the Local Timing Trigger and Control card (LTTC). This is followed by the first four Electron Isolation cards (EI) with their corresponding Receiver cards (RC) staggered 1.52cm to the right in the rear of the crate. The next card is the Jet Summary card (JS), with no matching Receiver Card. Finally the four remaining Electron Isolation cards fill out the crate. The corresponding Receiver cards are plugged in at the rear, staggered 1.52cm to the left. The figure is not to scale in that it implies




substantially more free space at the right hand end of the crate than actually exists.

### 2.2.3 Backplane construction

The backplane will be a 41.91cm x 39.70cm multilayer printed circuit board, between .32cm and .36cm thick. The final thickness will depend upon the tradeoffs necessary to achieve the required impedance and proper mechanical strength with standard core stock and pre-preg. There will be on board VME terminations, multiple studs for power supply connections, bypass capacitors, and mounting holes.

The design impedance will be  $50\Omega$  for the lower 2/3 of the backplane which contains all the trigger data paths. This impedance matches the AMP connector impedance and the impedance of the individual boards. The wiring in this section is all point to point, making for a fairly straightforward transmission line. Terminations for these lines will be on the receiving cards. The multi-drop transmission line of the VME backplane in the upper 1/3 of the backplane will have an impedance of  $100\Omega$  – before holes and connectors are added. The effective impedance of this section will drop to a little less than  $50\Omega$  after connectors and trace stubs on the individual boards are taken into account.

The VME specification is written for crates containing no more than 21 cards. In addition, the specification requires that VME signal stubs on individual boards be no greater than 5.08cm in length. The present design for the trigger processor crates contains 20 cards with a VME interface. Reducing the VME interface from two connectors to one has increased the difficulty of staying within the 5.08cm requirement. Special care will be taken throughout the design process to stay within the VME standards.



**Figure 5.** Schematic representation of backplane showing stack-up for VME and Trigger Data sections.

Preliminary calculations indicate it should be possible to route the high density wiring in the trigger data portion of the backplane in three signal layers. In order to ensure there are sufficient wiring channels and to maintain symmetry, four buried layers will be used in this section. A picture of the board stack-up is given in Figure 5. The organization of the layers is reflected about an extra thick core. An alternating power plane/signal plane structure produces a buried stripline design and surface micro-strip traces. The stripline structure provides good control over crosstalk and general noise immunity. The outside signal layers will be used mostly in the VME section of the backplane. Signal lines in that section will have ground traces running between each pair of lines to reduce crosstalk. Signals in the trigger data portion will be constrained, but not strictly limited to, the buried stripline layers.

#### **2.2.4 Data Rates**

All signals in the trigger data portion of the backplane will be transmitted at 160 MHz. This data rate was chosen because it offers the opportunity to compress the number of data lines on the backplane by a factor of four and because it should be realizable by available technology. More testing should be done, at these rates, to produce a design that maintains the necessary signal fidelity and clock to data skews.

All signals in this section, including clocks, are transmitted point to point. In the present design, all signals are differential. It may be that, with the stripline construction and the more than 160 backplane to board power connections afforded by the AMP stripline connectors, the data signals can be transmitted single-ended. The option to use single-ended transmission awaits the result of further testing and simulation.


### **3. Detector to Crate Mapping**

A preliminary mapping of the calorimeter to the Level 1 regional trigger crates is presented here. Trigger tower geometry has yet to be optimized, and changes in the calorimeter tower layout will result in changes to this mapping. The efficiency with which crates are packed will also be affected by changes in the calorimeter design.

#### **3.1 Eta - Phi Space**

The current proposal is for trigger towers to have a size of .087 in both phi and eta for all phi between eta of 0.0 and 2.1. This granularity produces 72 towers in phi and 24 towers in eta for each half of the calorimeter. Between eta of 2.1 and 2.6 the tower size changes to .174 in eta, with the phi granularity unchanged. This produces another three towers in eta. Therefore, between eta of 0.0 and 2.6, there are 1944 Electromagnetic and Hadronic trigger towers in the half detector. The Very Forward Calorimeter will have 18 trigger towers in phi and 4 towers in eta for eta between 2.6 and 4.0. The last region,

between eta of 4.0 and 5.0 will be divided into 9 towers in phi and 1 tower in eta. The Electromagnetic and Hadronic towers are combined on the detector, in the region of  $2.6 < \eta \leq 5.0$ , producing a single trigger tower for each pair. The grand total of trigger towers is 7938 for the full calorimeter.



**Figure 6.** Mapping of Calorimeter Barrel and Endcap onto Trigger crates.

Figure 6 illustrates a mapping of the calorimeter onto the Regional Trigger Processor crates for the region between eta of -3.0 to 3.0. The 18 crates shown are required for the barrel and both endcaps. Another crate (not shown) is included for the Very Forward Calorimeter. Although this mapping does not completely minimize the intercrate cabling, it simplifies the backplane and produces a rack layout that minimizes the cable lengths between racks.

The Very Forward Calorimeter covers the region  $3.0 < \eta \leq 5.0$  and is not used for the isolated electron triggers. The electromagnetic and hadronic information is combined on the detector and sent to the trigger processing hardware as a single sum. This region can be covered with a single trigger crate by using the same input pins between the 8 bit electromagnetic data and 4 bit hadronic data inputs of the Electron Isolation card and the  $E_T$  inputs of the Jet Summary card. With this overlap, along with consistency in power/ground and clocking inputs, it is possible to substitute a Jet Summary card for an Electron Isolation card. The Jet Summary card forms the  $E_T$  sums and performs the trigger cuts on the data from up to eight sub-regions of the Very Forward Calorimeter.

### 3.2 Backplane Layout

The map of the calorimeter onto crates is used to derive a card to slot mapping within the crate itself. The layout of cards within the crate is setup to minimize the length of traces on the backplane. Figure 7 illustrates a layout which attempts to satisfy this condition. For reference the schematic mapping of the calorimeter onto a single crate is repeated with a slot numbering (left to right) superimposed on the eight Receiver cards. This numbering is relative within the trigger data region of the crate and is not meant to represent a specific card slot.




Figure 7. Card layout in the trigger data portion of the crate.

The sharing shown in the figure is representative of an "average" card and does not show the intercrate connections. Card 3, for instance, would share with cards 0, 1, 2, 4, and 5, within the crate, as well as with cards 0, 2, and 4 in the adjoining crate. Every card, except those at the extreme rim of the barrel, shares with eight other cards. The four "corner" neighboring cards only transmit 2 bits of hadronic information to central card.

## 4. Receiver Card

The Receiver card is the largest board in the crate. It is 9U by 400mm. Both sides of the card are fully utilized to receive and process calorimeter trigger data. The rear side of the card receives the calorimeter data from optical fibers, translates from fibre to copper, and converts from serial to parallel format. The leading edge of the board will provide strain relief and some space for a small coil of fiber. The front side of the card contains circuitry to synchronize the incoming data with the local clock, and check for data transmission errors. There are also lookup tables and adder blocks on the

front. The lookup tables translate the incoming information to transverse energy on several scales. They are also used to test for Quiet and Minimum Ionization thresholds for each trigger tower. The memories were placed on this card, rather than the processing cards, because the sharing of data between processing cards would have required nearly twice as many memory chips. The energy summation tree begins on these cards in order to reduce the amount of data forwarded on the backplane to the Jet Summary card. Separate cable connectors and buffering are also provided for intercrate sharing.

#### 4.1 Inputs

Each card is designed to receive 32 fibres from the FERMI [2] front-end readout system on the calorimeter. Each fibre transmits either two towers of hadronic or electromagnetic information per crossing. The data is transmitted on the fiber in a compressed eight bit format, providing a total of 16 data bits. The present design uses transmitter and receiver links capable of handling 21 bits of information in 25 nsec with a baud rate of 960 Mbaud. The additional 5 bits, not used for data, contain a single Hamming code generated from the 16 bits of data. This code is sufficient to detect all single and double bit errors as well as many multiple bit errors. The rear of the Receiver card, containing the fibre receiving circuitry is shown in Figure 8.




Figure 8. Rear view of the Receiver card and the fibre receiver circuitry.

The mapping presented in section 3.1 shows two Receiver cards in each crate that are not fully utilized. The first 12 towers connected to each of these two cards form a partially filled 4 x 4 region which includes separate electromagnetic information for electron isolation. The front end electronics, reading out the region between eta of 2.6 to 5.0, sums the electromagnetic and hadronic towers on the detector before sending the data to the Level 1 regional Trigger crates. The Endcap calorimeter ends at an eta of 3.0. With a

subdivision in  $\phi$  of 18 we have a resolution, between  $\eta$  of 2.6 to 3.0, equivalent to approximately one  $4 \times 4$  sub-region for each  $\phi$  division. We plan to complete the coverage of the Endcap calorimeter by including one tower ( $2.6 < \eta \leq 3.0$ ) in each of the partially filled Receiver cards. This uses a 13th input channel on the card. Whereas the other front end electronic channels transmit two towers of information per fibre per crossing, this channel must be limited to one tower per fibre in order to limit the data on the Receiver card to the usual two  $4 \times 4$  regions transmitted to the Jet Summary card.

In order to minimize the cost of under-utilizing Receiver cards, those cards in the region of  $2.6 < \eta \leq 3.0$  will not have fibre receiving circuits and Synchronizing ASICs installed in the 19 unused channels. Since those circuits are likely to be the most expensive components on the board, partial assembly should achieve a major savings over the standard Receiver board cost.

The fibre optic link is also used to transmit synchronization information either on a regular basis or on demand. The re-synchronization philosophy will be discussed in greater detail in section 4.3.

## 4.2 Fibre Receiver

In the present design, the fibre receiver is made up of two separate components: The fibre to copper converter and the serial to parallel converter. As optical communication technology advances it is possible that integrated receivers (fibre to copper/serial to parallel), suitable to our application, will become available. For the present, we will concentrate on solutions that are available today in the market place and continue to examine any new offerings. The components used on the Receiver card will have to match those used on the calorimeter front end electronics to transmit the information.

Two separate fibre receivers have been considered for this design. The first device, made by BT&D, is one already in use at University of Wisconsin High Energy Physics. It is a high quality component in a 24 pin package, with long laser diode lifetime (transmitter), and a fiber pigtail. The pigtail is important in reducing the impact on the laser diode cavity by reflections from the first optical connection. This device is meant to operate at higher bandwidths and longer distances than we require and has a price of approximately \$700(US) per end. BT&D have indicated they will produce a less expensive device aimed at a different market niche sometime in the future.

The second converter is a more recent offering by Finisar. It is designed to handle data rates over the range of 100 Mbaud to 1.5 Gbaud and drive fibre lengths up to 500 M. The device is powered by a single +5V supply and is packaged in a .9" by 2.5" by .5" high case. The long dimension includes the fibre connector. The output from the receiver is differential ECL, but not at

true ECL levels. Capacitive coupling to the serial to parallel converter is required. The lasing diode lifetime, at 210,000 hours, is about half that of the BT&D device, but without an understanding of the statistical distribution it is difficult to determine the impact this will have on a 4000 link system. The specifications list the Bit Error Rate (BER) of this receiver at  $< 10^{-12}$  for an input dynamic range of -13 to +2 dBm. Finisar gives the typical value for the BER as  $< 10^{-15}$ . These BERs are evaluated using a  $2^7 - 1$  pseudo random bit sequence. The price for the Finisar device is presently \$250(US)/pair in lots of 1000.

Any transmitter/receiver pair considered for this application should be tested to determine the extent to which fibre length, baud rate, or environmental conditions degrade the BER.

The second component of the receiver duo is the serial to parallel converter. For the present design we have considered the Hewlett Packard HDMP-1014. It has several advantages over other offerings. The receiver and transmitter are packaged separately in 80 pin surface mount quad packages. Separate packaging of the receiver/transmitter pair leads to correspondingly lower power consumption at each end, about 2 watts, and a more compact package. A pair can transmit and receive 21 bit data frames and does not require 8/10 encoding. A single bit control switches between command and data transmission "on the fly". This last feature is particularly important to the synchronization scheme we are proposing. Finally, the parallel outputs are true ECL levels, providing a good match to the high speed circuitry downstream from the receivers. These devices are presently priced at \$100(US)/pair in lots of 1000.

The ground and power planes in the region of the fiber to copper converters will be split to minimize channel to channel crosstalk and will not be connected through the board to the digital power planes. We expect to use a separate voltage source for the optical converters. Filters will be used on the power inputs of the HP chips to reduce the possibility of coupling between the individual Phase Lock Loops. From a routing standpoint, there are few traces on the rear of the board up to the region containing the HP receivers. The outputs of the HP receivers will generate 32 x 21 lines of data. This set of lines will be referenced to the digital power planes and routed within them, in stripline fashion, to the Synchronization ASICs. Presently, the design has the fiber receivers at the back of the board (farthest from the backplane). We are exploring the interchange of positions of the HP converters and the fiber receivers to minimize the routing of the parallel outputs to the Synchronizers.

### 4.3 Synchronization


The synchronization circuitry is contained on the front side of the Receiver card along with the memory look up tables and adder tree. Figure 9 shows the organization of these components on the board. The outputs of the HP

receivers are not only unsynchronized with the local clock but are also not aligned to the same bunch crossing. The method of synchronization we are using was originally proposed to handle the SDC fiber link word synchronization on the SSC [4]. We offer here a short description of the technique followed by a detailed description of an ASIC proposed as a local component of the global synchronization method.

In order to achieve synchronization with this method several conditions must be met.

1. A FIFO several crossings deeper than the difference in time between the shortest and longest link is required at the receiving end of each link.
2. When a "reset" or "realign" of the link is released at the sending end, the front end data entering each link must be associated with the same bunch crossing.
3. One link is dedicated to control only. This control link is slightly longer, in transmission time, than any of the other links in the system and is used to enable the outputs of all the Synchronization ASICs.

If the above conditions are met, the circuit described below should successfully synchronize any individual link with all the rest.



**Figure 9.** Front view of the Receiver card with Synchronization and Adder tree circuitry.

### 4.3.1 Synchronization ASIC

The synchronization ASIC is designed to receive four channels of parallel data (84 bits) from the serial/parallel converters. A block level diagram of the ASIC is shown in Figure 10. Data is transmitted from each converter along with a strobe,  $CLK_N$ , and a command/data control bit,  $c/d_N$ . These two bits



are in synchronization with the incoming data, but not with the local trigger processing system clock.  $CLK_N$  is used to clock the data into a FIFO within the Synchronization ASIC. The command bit,  $c/d_N$ , is used to switch the ASIC out of data transfer mode into synchronization mode.

When a channel is in synchronization mode it is receiving a sequence of commands designed to bring the data into the FIFO in proper alignment. Each channel has a small controller which decodes the command information into STOP, CLEAR, START, and LAST. The STOP command blocks input data to the FIFO. The CLEAR command clears the FIFO. The START command enables the FIFO input and sets up the output to be enabled synchronously with the local system clock when the enable from the control link, Enable Out, reaches the FIFO. The LAST command is used as filler to ensure there will be at least two words in the FIFO of the longest data link when the Enable Out arrives.




Figure 10. Block diagram of Synchronization ASIC.

The Enable Out, arriving from the control link, has been synchronized with the local system clock by the internal logic of control link's Synchronization ASIC. The Enable Out is deskewed with respect to the individual crate clocks, fed to the data links in parallel, and used to enable each synchronizer's FIFO output. All the data in the separate FIFOs is clocked out on the same crossing and, as a result of the conditions given in section 4.3, the data at the head of each FIFO is from the same crossing.

In order to achieve maximum utilization of board space, all the logic following the Synchronization ASIC is run at 160 MHz. This high data rate provides a potential savings of a factor of 4 in component volume. A significant savings is also realized by placing the multiplexing circuitry, necessary to convert the 40 MHz data flow into 160 MHz, at the output stage of the Synchronization ASIC. The four input channels, handling a total of eight pieces of 8 bit information per crossing, are placed on two output channels of 8 bits each. This 4:1 compression requires a corresponding increase in frequency, from 40 MHz to 160 MHz, to keep up with the incoming data flow.

There are also four Error Correcting Codes (ECCs) associated with the four input channels. After synchronization, each is checked against the data. If an error is detected a single bit is set, one for each incoming channel, and appended to the original ECC code. These four ECCs, along with their error bits, are transmitted out of the ASIC at 160 MHz on a single set of 6 pins. The full ECC information is made available off chip in case further error analysis is required.

The technology has not yet been chosen for this ASIC. The FIFO structure implies some static ram support. The operating frequency suggests GaAs or BiCMOS. The number of pins is on the order of 116, so the design will not be pin-limited. The final design of this ASIC will commence when the full set of synchronization protocols is determined for the CMS trigger system.

#### 4.4 Lookup Tables

Lookup tables are required to translate the information coming from the calorimeter front end electronics, in compressed format, onto the several different scales used by the energy adder tree and the Electron Isolation logic. The Hadronic and Electromagnetic energies are individually translated into eight bits of linear  $E_T$  with a resolution of approximately one GeV. These values are summed to provide total energy in 4 x 4 trigger tower regions of the calorimeter. They are forwarded to the Jet Summary card for further combination in the total transverse energy calculation and used as the basis for  $E_X$  and  $E_Y$  missing transverse energy calculation.

The isolation processor on the Electron Isolation card requires electromagnetic and hadronic transverse energy with several dynamic ranges. The reference tower needs 8 bits of electromagnetic transverse energy with a resolution of .5 GeV and 4 bits of hadronic transverse energy with a resolution of .5 GeV. The neighboring towers need 6 bits of electromagnetic energy with a resolution of 2 GeV and 2 bits of hadronic energy with a resolution of .5 GeV. The energies for the neighboring towers are easily derived from the values provided by the lookup tables for the reference tower by truncating top or bottom bits of each value. To guard against wrap-around when the 2 bit hadronic values are generated, any value greater than the maximum is set to the maximum.

With a 6.25 nsec period the memories must have access times less than or equal to 3.0 nsec in order to handle the usual board propagation and setup and hold times. Both Cypress and IDT make memories capable of operating at the required speed. The major drawback to these memories is power dissipation. A large fraction of the power required for the Receiver card will be consumed by these memories.

Data is downloaded to the memories and read back through the VME interface. This requires support circuitry located in the area of the memory chips. The data inputs to the memories can be tied in parallel for writing the chips, but all the address lines (224) need to be individually buffered. The buffering is located near the memory in order to maintain short board traces in the high speed section of the logic. Reading the data out of the memories back to the VME interface requires buffering for all the data output lines. This buffering is provided by an 8:1 multiplexer within each Adder ASIC.

#### 4.5 Energy Summing

The beginning of the energy summation tree is on the Receiver card. The transverse energy for each of the two 4 x 4 trigger tower regions is independently summed and forwarded to the Jet Summary card. On the Jet Summary card these  $E_T$  sums are used to continue the energy summation tree and also compared against a threshold to determine whether any sub-region contained jets. The  $E_T$  sums are applied to a set of lookup tables to generate  $E_X$  and  $E_Y$  for each 4 x 4 region. A separate adder tree is used to sum up  $E_X$  and  $E_Y$  from the regional values.

Though the input values at the top of the adder tree have only 8 bits of range, the adder tree has been designed to handle a dynamic range of 10 bits for either positive or negative values. This implies an overflow at approximately 1000 GeV, using the compressed scale described in the CMS Level 1 Calorimeter Trigger Performance Studies [5]. The exact value will depend on the resolution required of the input transverse energy for other trigger functions. Any overflow (or underflow) generated as the result of an arithmetic operation ( $A_{OV}$ ) will stay in time with the data and be ORed with any other overflow that might have occurred in the same crossing. All values are handled as 11 bit *2's-complement* numbers.

A second overflow condition can also occur. The value sent to the trigger processor from the calorimeter front end electronics may be at the highest possible count. The lookup tables will be programmed to output  $3FF_H$  for this particular input. This output code is recognized, at the top level of the adder tree, as indicating an input data trigger tower overflow ( $T_{OV}$ ). It is handled much like the arithmetic overflow in that it is ORed with other  $T_{OV}$  overflows and passed down the tree in time with the data that generated it. If the overflow is caused by a hardware failure, the lookup table can be re-

written to zero out the affected channels. The arithmetic and tower overflow bits are handled separately through to the bottom of the adder tree.

#### 4.5.1 Adder ASIC

The adder ASIC is implemented as a 4-stage pipeline with eight input operands and 1 output operand. There are only three stages of adder tree, but an extra level of storage has been added to ensure chip processing is isolated from the I/O. We have determined that the ASIC must work reliably at a clock period of 5.0 nsec in order to ensure safe operation at an in-circuit period of 6.25 nsec.

This ASIC uses 4 bit adder macro cells to implement twelve bit wide adders. Eleven bits are wired, left justified, to each operand of an adder. The LSB of each adder will be internally set to ZERO. The MSB is treated as a sign bit. Therefore, although the adder tree may be constructed from three 4 bit adders, the width of the operand data paths has been limited to eleven bits.

An Adder ASIC chip is designated as 'master' if it is in the top rank of the adder tree and as 'slave' if it is further down. Masters can generate Tower overflow ( $T_{OV}$ ), but slaves can only propagate  $T_{OV}$ . Both masters and slaves can generate and propagate arithmetic overflow/underflow ( $A_{OV}$ ). These bits are appended to each input and output operand, making all operands 13 bits wide.  $T_{OV}$  becomes the twelfth bit of the output result and  $A_{OV}$  the thirteenth bit. The data outputs of the chip are forced to  $3FF_H$  when either an overflow or underflow occurs.

##### 4.5.1.1 Detailed Description

The top of the adder tree is composed of four 12 bit adders and includes the logic required to detect and propagate  $T_{OV}$  and  $A_{OV}$ . The  $T_{OV}$  generate circuitry is a filter designed to detect the input code  $3FF_H$ . The  $A_{OV}$  generate circuitry examines the sign bits of the input operands and the results operand, together with the carry out, to determine whether or not an overflow or under flow has occurred. All eight of the  $T_{OV}$  bits are ORed together and all four of the  $A_{OV}$  bits are ORed together to form two separate overflow bits that are forwarded with the data in the pipeline. Edge triggered registers are used to store the results for the next stage of the adder tree. A block diagram of the Adder ASIC is shown in Figure 11.

The second stage contains two more 12 bit adders and includes the logic needed to propagate  $T_{OV}$  and to detect and propagate  $A_{OV}$ . From this point on,  $T_{OV}$  is forwarded down the pipeline from register to register.  $A_{OV}$  is generated in the same manner as in the first stage and the resulting two bits are ORed with the  $A_{OV}$  from the previous stage. Edge triggered registers are used to store the results for the next stage of the adder tree.

The third stage contains the final adder as well as a continuation of the  $T_{OV}/A_{OV}$  circuitry. The register at this level is the last storage element before the ASIC output. If either  $T_{OV}$  or  $A_{OV}$  have been detected, the output operand stored in this register has the value  $3FF_H$ .  $T_{OV}$  and  $A_{OV}$  are stored along with the operand. Adder Tree ASICs further down in the tree are designated "slaves" and are blocked from using the operand  $3FF_H$  to generate  $T_{OV}$ . Thus we retain the identity of the tower overflow bits through the entire tree.




Figure 11. Block diagram of the Adder ASIC.

The last register is presented to one input of a 2:1 multiplexer before leaving the chip through the boundary scan cells and pads. The other side of the multiplexer is fed by an 8:1 multiplexer which passes any one of the eight input operands, less the two overflow bits, to the output of the ASIC. This feature was provided to minimize the external logic needed to read back the values of the lookup tables that feed the first stage of the adder tree logic.

The equations below provide definitions for  $A_{OV}$ ,  $T_{OV}$ , and the output operand. Clock cycle  $n+1$  is the time at which these three values are present at the outputs of the ASIC.

$$\left. \begin{aligned} T_{OV}(n+1) &= 1 \text{ if any input operand } OP_{IN}^{1 \rightarrow 8}(n-3) = 3FF_H \text{ for the masters.} \\ T_{OV}(n+1) &= T_{OV\_IN}^1(n-3) + \dots + T_{OV\_IN}^8(n-3) \text{ for the slaves.} \end{aligned} \right\} (1)$$

$$\left. \begin{aligned} Out_{0 \rightarrow 10}(n+1) &= S_{1 \rightarrow 11}^7(n), \text{ if } A_{OV}(n) \text{ and } T_{OV}(n) \text{ both equal ZERO.} \\ Out_{0 \rightarrow 10}(n+1) &= 3FF_H, \text{ if either } A_{OV}(n) \text{ or } T_{OV}(n) \text{ equal ONE.} \end{aligned} \right\} (2)$$

$$\left. \begin{aligned} A_{OV}(n+1) &= A_{OV}(n) + A_{OV}(n-1) + A_{OV}(n-2) + A_{OV}(n-3) \\ &\text{where} \\ A_{OV}(n-3) &= \bigcup_{p=1}^{p=8} A_{OV\_IN}^p(n-3), \\ A_{OV}(n-2) &= \bigcup_{r=1}^{r=4} \left[ \left( \overline{A_{11}^r(n-2) \oplus B_{11}^r(n-2)} \right) * \left( S_{11}^r(n-2) \oplus A_{11}^r(n-2) \right) \right], \\ A_{OV}(n-1) &= \bigcup_{r=5}^{r=6} \left[ \left( \overline{A_{11}^r(n-1) \oplus B_{11}^r(n-1)} \right) * \left( S_{11}^r(n-1) \oplus A_{11}^r(n-1) \right) \right], \\ &\text{and} \\ A_{OV}(n) &= \left( \overline{A_{11}^7(n) \oplus B_{11}^7(n)} \right) * \left( S_{11}^7(n) \oplus A_{11}^7(n) \right). \end{aligned} \right\} (3)$$

$S_{11}^x$  is the sign bit of the xth sum.  $A_{OV\_IN}^p$  is the overflow input for the pth operand.  $A_{11}^r$  and  $B_{11}^r$  are the sign bits of the A and B operands of the rth adder.  $T_{OV\_IN}^x$  is the tower overflow bit for the xth operand. Finally, n is the nth clock cycle in the pipeline.

#### 4.5.1.2 ASIC Boundary Scan

The chip also contains boundary scan support. The ASIC boundary scan implementation, along with a proper board level implementation should provide full testing capability of the ASIC while it is in circuit. The boundary cells can also be used to verify circuit integrity (shorts, opens, and stuck at one/zero) at the board level. IEEE standard 1149.1 has been strictly adhered to in order to ensure compatibility with other ASIC and board level boundary scan controllers. The full JTAG controller and a major subset of the commands has been implemented. All inputs and outputs, with the exception of the five boundary scan control signals, have scan I/O cells.

The following instructions are recognized by the JTAG controller: BYPASS, IDCODE, SAMPLE/PRELOAD, EXTEST, and INTEST. BYPASS collapses the scan loop to a single bit, thus bypassing the ASIC as seen by the board level scan loop, providing faster testing of other chips on the board. IDCODE serially outputs an identification code indicating the manufacturer, part

number, and revision level of the ASIC. The four least significant bits are used to provide an address for the chip. This address will be set at the board level by hardwiring four input pins. SAMPLE/PRELOAD samples and preloads the boundary scan cells on two separate phases of the instruction. EXTEST tests the communication with the exterior by setting up values on the ASIC outputs and sampling values on its inputs. INTEST is used to test the internal circuitry. It accomplishes this by setting up known vectors on the input scan cells and capturing the results, after single stepping the core logic, in the output scan cells.

#### **4.6 Staging Circuitry for the Electron Isolation Board**

Thirty-two towers, in a 4 x 8 array, are processed on each Electron Isolation board. Data from twenty-eight neighboring towers is required to determine isolation for towers on the edge of the 4 x 8 region. Data is transferred between the Receiver card and the Electron Isolation card at 160 MHz. In order to retain point to point transmission, data going to a neighboring Electron Isolation board must be transmitted through separate drivers on separate backplane lines.

The order in which the data is transmitted to the Electron Isolation card is important. Since connector pins are at a premium, it is necessary to ensure that all lines have useful data for each of the four 6.25 nsec cycles making up a 25 nsec crossing. Board space is also at a premium, so it is necessary to limit the amount of circuitry committed to staging the data for the backplane. The most efficient way to satisfy these goals simultaneously is to pay careful attention to the order in which the fibres are connected to the input channels of the Receiver card. This, in conjunction with the multiplexing of data in the Synchronization ASIC, handles most of the data staging. The remaining staging is handled by a minimum amount of extra buffering on both the Receiver and Electron Isolation cards. More detail is provided below in the section describing the Electron Isolation processing.

#### **4.7 Outputs**

All outputs in the trigger data path are driven by ECL registers with differential outputs. Individual outputs are pulled down to  $V_{TT}$  through 390 $\Omega$ . The pulldown is used to provide signal levels on the board outputs or backplane when the receiving card (Electron Isolation or Jet Summary) is not plugged into the backplane. In normal operation the receiving card will provide a 50 $\Omega$  termination to the transmission line.

Even at 160 MHz there are a considerable number of outputs to deal with. The Electron Isolation card directly associated with a given Receiver card must have 32 towers of 8 bit electromagnetic and 4 bit hadronic information per crossing. Transmission of this information requires 192 pins in differential mode. In addition the neighboring cards to the left, right, top, and bottom, as shown in the top half of Figure 7, need 24 towers of 6 bit

electromagnetic and 2 bit hadronic information per crossing. These neighbors require an additional 96 pins. Finally the corner neighbors must have 2 bits of hadronic information transmitted each crossing. Since the corner information must go to four separate cards it is necessary to send it at 40 MHz, using 16 pins in differential mode. These numbers total to 304 pins required to transmit data from one Receiver card to several Electron Isolation cards.

The adder tree information is sent to the Jet Summary card. There are two totals, one for each of the 4 x 4 trigger tower regions covered by a Receiver card. These two 13 bit numbers will be multiplexed onto a single set of 13 differential pairs at 160 MHz. One of the remaining two 6.25 nsec data frames is used to transmit the quiet and minimum ionization information (4 bits), determined for the two 4 x 4 trigger tower regions, to the Jet Summary card. The remaining bits in that data frame and all the bits in the fourth 6.25 nsec data frame are presently undefined, but may be used in support of error logging and/or the tau trigger.

Every Receiver card shares its data with at most 6 Electron Isolation cards within the same crate. In addition each Receiver card sends some of its data off crate at 40 MHz to two or three neighboring crates. Crate to crate communication is handled by special cables running between the Receiver cards. This distributes the inter-crate buffering among the eight Receiver cards in a crate rather than attempting to put it all on one or two special cards at the ends of each crate. The most amount of information shared between two Receiver cards in different crates is carried on 204 twisted pair (102 in each direction) at 40 MHz.

#### **4.8 Boundary Scan**

The board level boundary scan is still under design. The boundary scan commands supported by the Adder ASIC have already been described, and the board level controller will support them as well. The following is an outline of functions under consideration.

To the extent it is possible, all registers on the outputs of the board, driving backplane lines, and all registers immediately following the Sync ASIC, will be fashioned after a boundary scan cell. These scan registers exist so that, through a combination of shifting and single stepping, we will be able to set up or capture data on the board boundaries. We do not expect to have enough room to implement the extra level of storage necessary to hold all inputs or outputs at a defined state while testing the board logic.

The board level boundary scan controller will have a hardwired program that is entered on power-up or by command through the VME interface. This program will use vectors stored in a PROM on the board to perform a minimal test on intra-board circuit integrity and make simple tests of data paths within the ASICs. The vectors will check continuity, stuck-on ones and stuck-on zeros, and shorts for those nets in the data processing path. The



results of the applied vectors will be read back into the boundary scan controller via the scan loop and compared with the expected results, also stored in PROM.

The boundary scan controller interprets commands sent through the VME controller to perform the individual JTAG instructions implemented in the ASICs, captures the resulting data, and sends it back to the crate controller on request. Several boundary scan loops will be implemented rather than a single long loop in order to minimize the time required for tests of a specific class.

#### **4.9 VME Interface**

The VME interface is a full implementation of the VME specifications, excluding READ/MODIFY/WRITE. Interrupts will be supported as these are useful during the boundary scan operations. At the moment, the implementation of Block Transfers is under consideration. While they may be useful for the memory operations, the extra logic may not be justified.

#### **4.10 Clock Distribution**

The on-board clock distribution is under design. The initial concept is discussed here. We will continue this development through board level simulation and test set-ups.

Synchronization of the incoming data with respect to the board clock is effectively handled by the Sync ASIC. The only requirement is that the global enable used by the synchronization chips must be in proper phase relationship with the board clock. Therefore, the board clock need not be adjusted to coincide with any particular phase of the incoming data.

The most successful clock distribution scheme is likely to be one in which the clock travels parallel to the data path towards the outputs on the backplane. All of the trigger data flow on the board is in the direction from the Sync ASIC to board output. Also, there are no re-entrant adds within the adder tree. As long as both the clock and the data take roughly the same path, it can be assumed that the extra delay of clock to data out, plus any combinational logic in the data path, is sufficient to ensure the clock will always arrive far enough ahead of the data at each successive storage element to guarantee hold times at that element. Setup times must be carefully determined since they are easily overrun by too much logic between registers.

Locating the clock distribution circuit near the center of the board could lead to the clock being distributed both parallel and anti-parallel to the flow of the data. Consistent distribution of the clock anti-parallel to the data will work, and could be safer, if there is known to be sufficient time between all levels of storage to allow for the small loss in cycle time due to the propagation of the clock on the board.

## 5. Electron Isolation Card

The Electron Isolation Card, shown in figure 12, receives data at 160 MHz in a staged fashion from at most nine Receiver Cards and performs the isolated electron algorithm described in the introduction. Some of the data originates from Receiver cards in neighboring crates, but is transmitted through the local Receiver cards. The Electron Isolation card is 9U x 280mm and resides in the front of the crate, offset from the Receiver cards by 1.52cm. The electron isolation algorithm is performed on this card and the final results sorted to identify the 4 highest rank electron isolation candidates.




Figure 12. The Electron Isolation Card.

The Electron isolation algorithm will be implemented in a custom ASIC. The results from the electron identification ASIC are sorted in a second ASIC (Sort ASIC) and the top four candidates from the 4 x 8 trigger tower region are transferred to the Jet Summary card. The Jet Summary card does a further sort using the same Sort ASIC to output the top four electron, jet, and isolated hadron candidates to the global trigger processor crate. The total  $E_t$ ,  $E_x$  and  $E_y$  information from the crate region is also forwarded to the global trigger processor.

### 5.1 Inputs

All data received by the Electron Isolation card comes from local Receiver cards in differential mode. Terminations for the lines will be on the cards rather than the backplane. Each card processes 32 towers of electromagnetic and hadronic information organized in a 4 x 8 array. Data is also required from the neighboring 28 towers to determine isolation on the boundaries of

the 4 x 8 region. Section 4.7 itemizes the 304 lines required by the Electron Isolation card. A 340 pin AMP stripline connector will be used as the main data connector. The additional pins beyond 304 will be used to forward the results to the Jet Summary card and to input clock and control information. As in the case of the Receiver card, the top part of Electron Isolation card will use a 128 pin DIN connector to interface to a 32 bit VME bus.

## 5.2 Isolation

The algorithm used to determine isolation examines the individual sums of a reference tower with its four nearest neighbors. The maximum sum is chosen and two cuts are applied on the longitudinal and transverse isolation of the ECAL energy deposit. The first cut requires the central tower HCAL to ECAL energy ratio to be  $< 0.05$ . The cut on transverse isolation requires the sum of HCAL transverse energy in the eight nearest neighbors to be  $\leq 2.0$  GeV. In order to reduce the number of bits exchanged between cards, we limit the dynamic range of neighboring HCAL information to 2 bits. Overflows of any of the energy ranges are treated as maximum values. An option is being considered to include electromagnetic isolation by requiring the energy in one of 4 corner-centered combinations of 5 ECAL towers to be less than 2.0 GeV.

The input data is staged on the Receiver card to arrive in a particular order in time at the Electron Isolation card. Careful data organization makes it possible to process the data as it is being received, rather than storing it for a full crossing to use in parallel on the card. The Isolation ASIC is designed to shift in the data for 16 towers, 4 towers at a time, over a single bunch crossing time. The data for the 20 neighboring towers is also be loaded in the same time period. The entire 4 x 8 region is processed by two ASICs in four 160 MHz cycles. A detailed description of the order of the data is given in section 5.2.1.

The output of the ASIC is four 2-tower sums, with eight bits of dynamic range, and four 1-bit results indicating whether the eight nearest neighbor hadronic sums are  $\leq 2.0$  GeV. These results are produced every 6.25 nsec. In parallel with the Isolation ASIC are a set of lookup tables which use the same input information to determine H/E for each of the reference towers. The four 2-tower sums from the Isolation ASIC are combined with the two single bit results from the hadronic sum and H/E cuts and presented to a programmable look-up table which ranks the electron trigger results.

These ranks, one per reference tower, are presented in parallel to a single Sort ASIC. The Sort ASIC receives all 32 ranks from the two Isolation ASICs in one crossing and appends 5 bits of location information to each input. The 5-bit location follows each datum through the Sort ASIC and uniquely identifies the four largest. The resultant from the ASIC is the four highest ranks from the 4 x 8 trigger tower region. There is a four crossing latency for the result, but the pipeline architecture ensures that, once filled with data, a new result will appear every crossing.

### 5.2.1 Isolation ASIC

The Isolation ASIC, shown in Figure 13, handles four electromagnetic energies ( $A_{in}$ ,  $B_{in}$ ,  $C_{in}$ ,  $D_{in}$ ) on an 8 bit scale, every 6.25 nsec. Nearest neighbors are also included in the data flow. During the first cycle of every crossing the four neighboring energies ( $TE_A$ ,  $TE_B$ ,  $TE_C$ ,  $TE_D$ ) from the adjacent 4 x 4 region (top) are also be strobed into the ASIC. The neighbors along either edge of the 4 x 4 region ( $LE_{in}$ ,  $RE_{in}$ ) are also included, two at a time (left and right edges), during each 6.25 nsec period. Finally, the last cycle strobes in the four neighboring towers of the bottom edge ( $BE_A$ ,  $BE_B$ ,  $BE_C$ ,  $BE_D$ ). The nearest neighbor corners are not required by the algorithm. Thus, in one bunch crossing time, a total of 36 towers are clocked into the Isolation ASIC. All the neighboring electromagnetic energies are provided as the top 6 bits of the 8 bit range used for the reference towers.




Figure 13. Block diagram of the Electron Isolation ASIC.

Separate inputs are used to clock in the top and bottom neighboring towers in order to avoid unfavorable routing or extra components on the board due to board level multiplexing. The top and bottom neighboring edges require a total of  $2 \times 4 \times 6 = 48$  input pins. The central region needs  $4 \times 8 = 32$  input pins, and the left and right edges need  $2 \times 6 = 12$  pins. The total number of data inputs is 92 pins, for a total of 36 towers of information. In addition pins

are allocated for reset and clock inputs. All signal I/O, with the exception of the clock, is single-ended.

The Isolation ASIC also handles the hadronic tower data, as discussed in section 5.2.1.4, necessary to check for hadronic isolation around individual towers. All neighboring hadronic information is on a 2 bit scale. Again, the data is shifted into the ASIC on 6.25 nsec cycles. The central 4 x 4 region requires 4 x 2 bits, or 8 pins, for each cycle. The top and bottom edges are entered on separate inputs, requiring a total of 2 x 6 x 2 bits, or 24 pins. The factor of six is necessary because the nearest neighbor corners are included in the determination of hadronic isolation. The left and right edges need an additional 2 x 2 bits, or 4 pins, to complete the count. The total number of pins required for the hadronic data is 36.

The only outputs are the four electromagnetic and four hadronic results. Four two tower sums, on an 8 bit scale, and four corresponding 1 bit hadronic isolation flags are produced each 6.25 nsec. One crossing is required to output the results for the entire 4 x 4 array. The total number of output pins is 36. The combined total of inputs and outputs, including control signals, is 170 pins. Since the ASIC is essentially a dataflow device, there is very little in the way of control logic or control lines needed for the operation of the chip.

### 5.2.1.1 Input Staging

The Isolation ASIC processes the data through three separate blocks. The first of these, the Input Staging, is illustrated in Figure 14. The purpose of this block is to receive the data at the time when it is available and change the time relationship to be suitable for the processing that follows. Only one column of input data is represented in the figure.

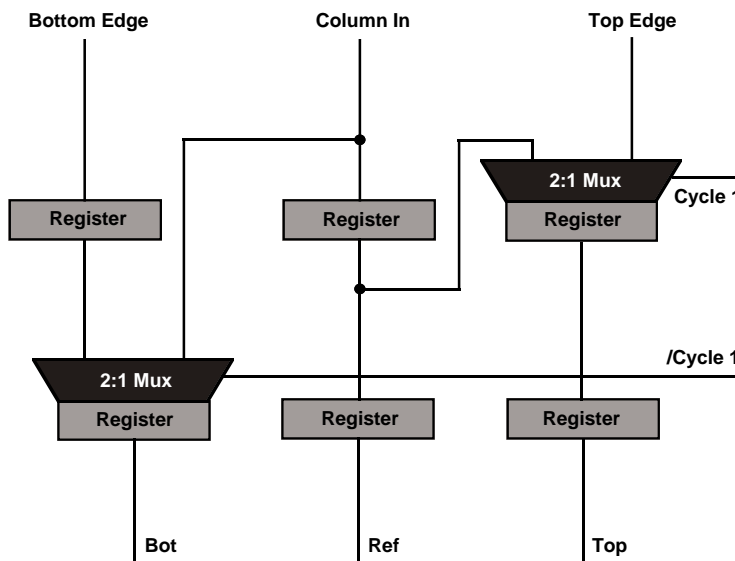


Figure 14. Isolation ASIC input staging.

At the beginning of a crossing, the first row of the 4 x 4 array is available, along with the top edge. The signal **Cycle 1** selects the **Top Edge** input on the right hand multiplexer. After the first 6.25 nsec clock, the first rank of registers contain one of the towers in the 4 x 4 array (a reference tower) along with its top neighbor. The left-most register in the top rank is undefined at the beginning of the sequence. After a second clock cycle, the reference tower is in the middle register of the bottom rank of registers and its top neighbor is in the right hand register. The left-most register in the bottom rank contains the next successive reference tower, as does the middle register in the top rank. This value in these registers is the bottom nearest neighbor for the first reference tower. The sequence continues to the cycle where the last reference tower in a column of 4 towers is clocked into the middle register in the bottom rank. During the same cycle the **Bottom Edge** data is available from the neighboring card. It is clocked into the bottom left register during **Cycle 1** at the beginning of the next sequence.

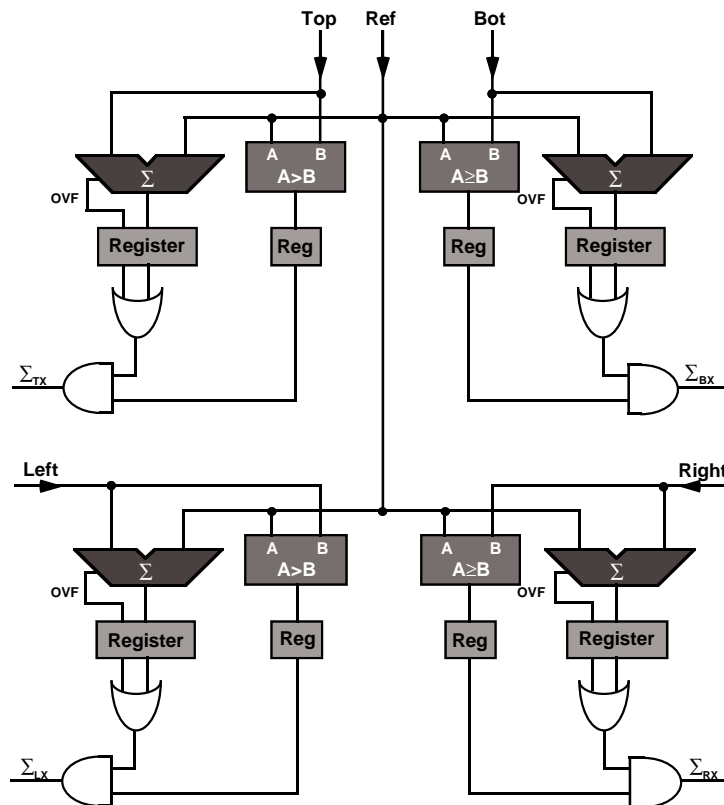


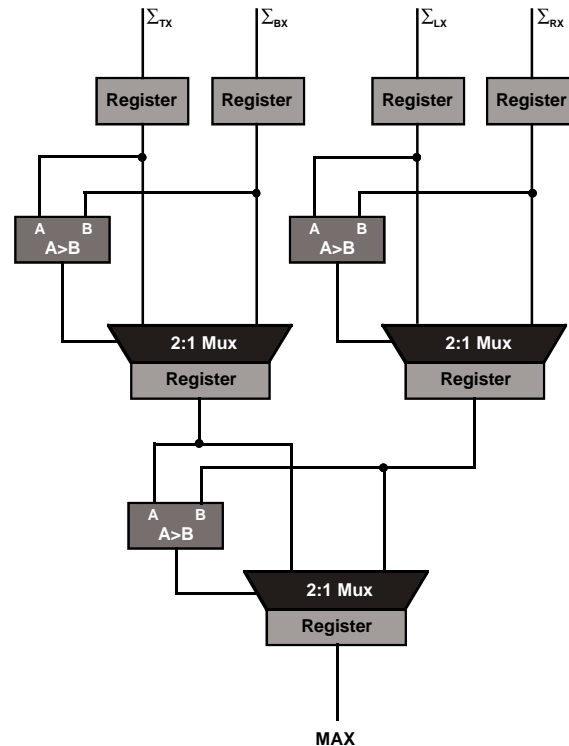
Figure 15. Electron Isolation ASIC Add/Compare block.

Once the pipeline has been filled, data will continue to be output from the Input Staging block four towers at a time each with their corresponding top and bottom neighbor. The left and right neighbors are either the adjacent reference towers in the 4 x 4 array or the left and right columns of data from neighboring boards.

### 5.2.1.2 Add/Compare

The Input Staging block places each tower and its neighbors in the same time frame. The remaining blocks in the chip can now handle the processing in parallel. The function of the Add/Compare block, shown in Figure 15, is to form four sums between a reference tower and its top, bottom, left and right neighbors. At the same time the sums are being formed, four compares are made to determine for each pair of towers whether the reference tower is larger than or equal to its neighbor. In the case of comparison with the top tower, the reference tower must be greater than the top. When compared with the bottom tower the reference tower must be greater than **or** equal to the bottom. The different conditions used to check for inequality with respect to the top and bottom towers is imposed to remove the possibility of double counting when two towers, one above the other, are equal. This does build a bias into the circuit in that the bottom tower, of an identical pair, is always taken as the reference tower. The same process is used to guard against double counting in the left to right direction.

When a tower pair satisfies the equality check the sum from the adder is enabled to the Find Max block. When the sum is disabled, a value of zero is passed on to the next block. If the adder has overflowed (carry out equals one) the results of the addition are set to FF<sub>H</sub>.



**Figure 16.** Electron Isolation ASIC block which determines the maximum pair.

### 5.2.1.3 Find Max

The last stage in processing the electromagnetic information, the Find Max block, is shown in Figure 16. The four sums ( $\Sigma_{TX}$ ,  $\Sigma_{BX}$ ,  $\Sigma_{LX}$ , and  $\Sigma_{RX}$ ) are presented, in parallel, to two comparators. The outputs of these comparators are used to select the maximum of each pair which are placed in intermediate storage. These two maxima are presented to a single comparator during the next clock cycle. The output of the comparator is used to select the maximum. The single maximum from the original four values is stored in the bottom register shown in Figure 16.

The total latency for the electromagnetic data path is  $7 \times 6.25$  nsec or 1.75 bunch crossing times. It is likely that detailed simulation will show the register at the input of the Find Max block can be removed. This would reduce the latency to 1.5 bunch crossings.

### 5.2.1.4 Nearest Neighbor Hadronic Sums

The hadronic information enters the ASIC in the same time frame as the electromagnetic information. Some staging is necessary on the board to put the four corner towers from neighboring boards into the proper time sequence. Figure 17 illustrates the section of the chip that checks the hadronic isolation.

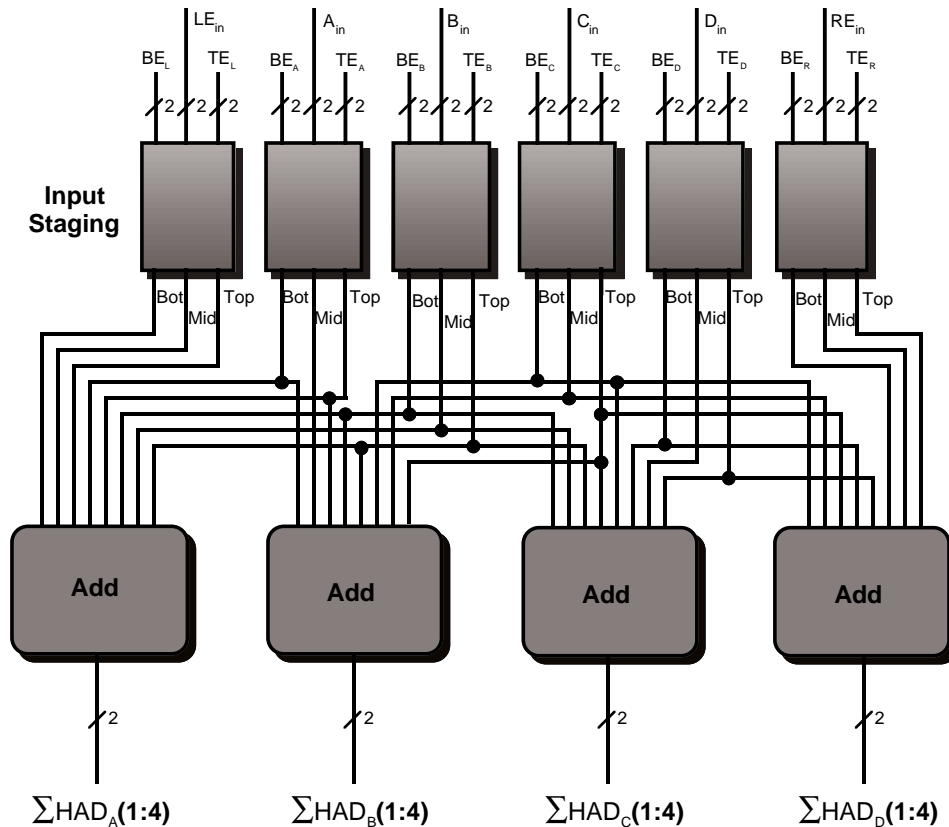


Figure 17. Electron Isolation ASIC Hadronic isolation block.



The same Input Staging design as that used in the electromagnetic processing section is used to put reference towers and nearest neighbors into the same clock cycle. Two cycles (12.5 nsec) after entering the ASIC, the data is presented in parallel to four 8-operand adders. Each operand has a resolution of 2 bits. Details of the adder block are not shown. The adder blocks are implemented as binary trees, each reducing the original 8 operands to one in three 6.25 nsec cycles. The result from each adder block will be checked for  $\leq 2.0$  GeV and a single bit set to indicate the sum has passed the cut. The output of this bit will be timed in to coincide with the output from the electromagnetic processing section of the ASIC.

#### 5.2.1.5 Boundary Scan

It is our intention to implement the same controller and instruction decoder for boundary scan as used in the Adder ASIC. All ASICs in the Level 1 trigger processor will have compatible boundary scan circuits.

### 5.3 Lookup Tables

Lookup tables are used for two separate functions. The first is to calculate the H/E cut. This operation is performed in parallel with the processing in the Isolation ASIC. Eight memories, cycling at 160 MHz, receive the same data as that shifted in to the reference tower locations in the Isolation ASIC. Only the top 4 bits of the 8-bit electromagnetic data are used. Simulation studies indicate that full resolution on  $E_T$  is not required [5]. The four bits of the hadronic information is combined with the four bits of electromagnetic information, producing an 8-bit address into the H/E lookup table. A single output bit is used to indicate, when true, that  $H/E < 0.05$ . The full  $4 \times 8$  trigger tower region is processed in one crossing, which is slightly less than the time required by the Isolation ASIC. The H/E bits will be delayed a sufficient number of cycles to ensure they are available at the same time as the corresponding sums and hadronic isolation bits are available from the Isolation ASIC.

A second set of lookup tables are used to produce a ranking of the electron trigger information based on the 8-bit sums, hadronic isolation from the Isolation ASIC, and the single bit result from the H/E cut. These 10 bits are used as an address into the lookup tables and will produce a value of 6 bits in length which represents a ranking of the different trigger categories resulting from the data. The memories will be cycled at 160 MHz. Therefore, one bunch crossing is required to produce the ranking information for the  $4 \times 8$  region.

### 5.4 Sort

In order to limit the amount of information transmitted from each Electron Isolation card and each regional trigger crate, we have chosen to send only the four highest ranked trigger categories from each Electron Isolation card to the

Jet Summary card. These four values have not been put in any particular order. Jet Summary card receives a total of 32 values from 8 Electron Isolation cards and sorts these again to determine the four largest values. The  $E_T$  information sent to the Jet Summary card from the Receiver card is also sorted to determine the four  $4 \times 4$  subregions with the largest energy deposition.

### 5.4.1 Sort ASIC

The Sort ASIC is designed to find the four largest of thirty-two 10-bit values. Ten bits is sufficient to handle both the  $E_T$  sums and the trigger category information. Figure 18 is an illustration of the major functional blocks that make up the ASIC. Rather than try to design an ASIC that will handle thirty-two 10-bit operands in parallel, it was decided to shift the data in, eight operands at a time, over four 6.25 nsec cycles. This matches the rate at which data is coming from the lookup tables and reduces the number of input pins from 320 to 80.

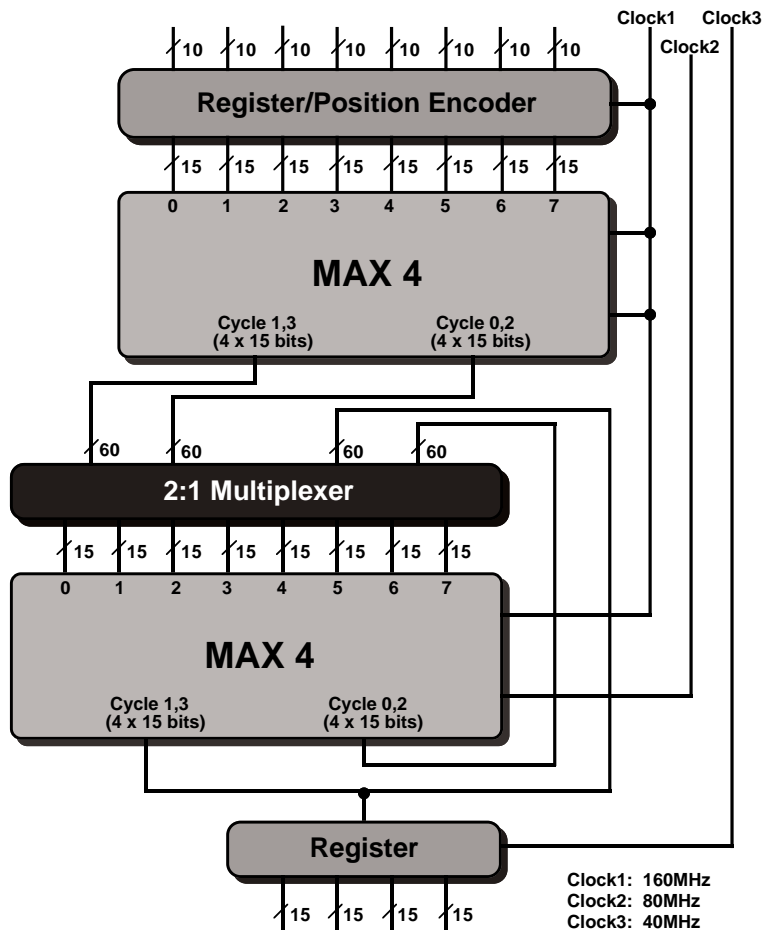



Figure 18. Sort ASIC block diagram.

The algorithm implemented within the Sort ASIC is based on a simple rotation of operands and is shown schematically in figure 19. The eight

operands are divided into two groups of four. The operands are compared in pairs between the two groups, with the larger of the two taking over the position of the left hand member of the pair. This comparison is performed in four stages with a rotation of compared pairs occurring between each stage. By the end of the fourth stage a sufficient number of comparisons have been made to ensure the four largest values are in the left-hand group. In order to save steps, and thus minimize the total latency, these four values are not placed in any rank order. The final four values, produced by the global trigger processor, are ordered during the final sort.

### 5.4.1.1 Register/Position Encoder

The first functional block of the Sort ASIC is the Register/Position Encoder. The main purpose of this block is to provide storage for the incoming data to isolate it from board propagation delays. It is also used to append 5 bits of positional information to each of the operands as they arrive at the chip. The position is based on the which input (1 of 8) and which cycle (1 of 4) the data coincides with. The cycle number is derived from a 2 bit counter which is initialized via a reset at startup time. The two values, position and time, are concatenated to produce the 5 bit value.



**Figure 19.** Sort ASIC compare pattern.

### 5.4.1.2 MAX4

The next block in the sequence is the MAX4 block. A more detailed picture is given in figure 20, which shows that the block is isolated from the previous block with a register. This level of register may not be required; it is included pending simulation results on its necessity. The rest of the MAX4 block is made up of four Compare/Select blocks. One of these blocks is illustrated in greater detail in Figure 21. Each of the four blocks is different from its neighbor in that each has one of the four stages of pairing, shown in Figure 19, for the compare hardwired in the circuit.

The Compare/Select block performs all four compares in parallel. Fifteen bits arrive and leave at the inputs and outputs. The top five bits are the positional information mentioned above. They are carried along with the data through the 2:1 multiplexers. The bottom 10 bits (raw data) are used by the compare to drive the select lines of the multiplexers. Each pair of multiplexers has the same operand wired to opposite inputs. The results from the compare force the left hand multiplexer to store the larger of the two values and the right hand multiplexer to store the smaller of the two. After one 6.25 nsec clock cycle the eight operands have been reorganized with the larger of each pair of values stored in the registers in the left hand column.




Figure 20. Sort ASIC MAX4 block diagram.

During the next three 6.25 nsec clock cycles the data is passed through the remaining three Compare/Select circuits. Only the left four operands are stored at the end as they contain the four maximum of the original eight operands. One clock cycle later the first results are stored in the register shown at the bottom right. The register to the left now contains the four largest values from the second set of eight operands. These two sets of four operands must be examined to select the four largest among them. This second level of selection occurs in the second MAX4 block shown near the bottom of figure 18. During the next two 6.25 nsec cycles the last two sets of four maxima are generated in the top MAX4 circuit and passed on, through

the 2:1 multiplexer, to the second MAX4 circuit. New data enters the second circuit every 12.5 nsec. This circuit is clocked at 160 MHz, but the data is interleaved with empty cycles.

After four more 6.25 nsec cycles, the four maxima from the first sixteen operands are at the Cycle 1,3 outputs of the second MAX4 circuit. This result is clocked into the Cycle 0,2 register of the circuit before it is overwritten by the empty cycle immediately following. The second register is clocked with an 80 MHz clock instead of the 160 MHz clock. Thus the first result is held for 12.5 nsec providing enough time for the empty cycle to clear and the maximum 4 from the second 16 operands to be clocked in. These eight values are presented back to the inputs of the second MAX4 circuit through the right side of the 2:1 multiplexer during one of the empty cycles. They are interleaved with new data from a later crossing. Thus, the four largest of the original 32 operands finally appear at the bottom of the second MAX4 circuit approximately 4 crossings (100 nsec) after the data first entered the Sort ASIC. This result is captured in the register as shown at the bottom of Figure 18. This register is clocked at 40 MHz making the results available at the outputs of the ASIC for a period of time equal to a full crossing (25 nsec). The latency for the full sort operation without any optimization is four crossing times, or 100 nsec.

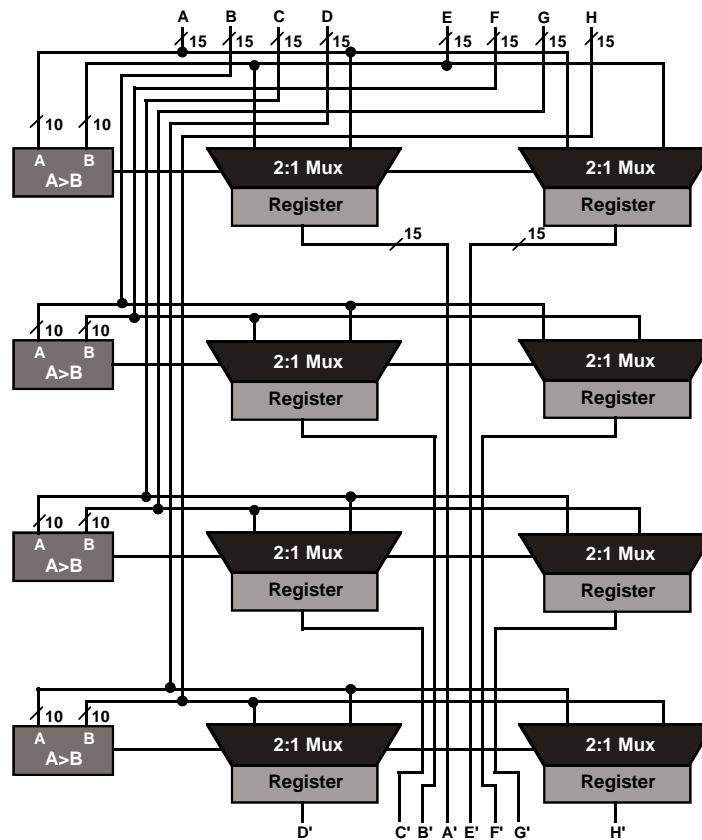


Figure 21. Sort ASIC Compare/Select block.

### 5.4.1.3 Boundary Scan

We intend to implement the same controller and instruction decoder for boundary scan as used in the Adder and Isolation ASICs. All ASICs in the Level 1 trigger processor will have compatible boundary scan circuits.

## 5.5 Outputs

Each Electron Isolation card produces four ranked trigger category values (6 bits in length) each combined with a single bit indicating which of the two 4 x 4 regions produced the value. These results are produced once every 25 nsec crossing. Every 6.25 nsec one of the values is placed on the backplanes using a set of 7 differential pairs. All four values are transmitted to the Jet Summary card during the period of one crossing. These 14 lines, together with the 304 input lines noted above, use 318 pins on the 340 pin connector. Additional pins will be used for clock, controls, and status information.

## 6. Jet Summary Card

The Jet Summary card sits near the middle of the trigger data processing section of each regional trigger processing crate. It is shown in Figure 3 situated in the front portion of the crate between the 4th and 5th Electron Isolation cards. It collects and summarizes data from both the Receiver cards and the Electron Isolation cards at 160 MHz. The Jet Summary card has the same form factor as the Electron Isolation card, 9U x 280mm. It has a VME interface via a 128 pin, 4 row DIN connector. The trigger data is received on an AMP 340 pin connector.

The electron trigger rank information from eight Electron Isolation cards (32 values) is sorted to produce the four highest ranked electron triggers. These values have a 4-bit address appended to them which indicates which 4 x 4 trigger tower region, covered by the crate, produced them. The  $E_T$  information from the Receiver cards is also on a 4 x 4 tower resolution. These values are summed by a pair of Adder ASICs to produce a total  $E_T$  sum. The same information is processed by a Sort ASIC to determine the four largest  $E_T$  values produced by the crate. This is used as the jet trigger. In addition, we test the  $E_T$  in 4 x 4 tower regions against a series of thresholds and encode the results in 4 bits. The  $E_T$  information from the Receiver card is also sent to memory lookup tables which generate  $E_X$  and  $E_Y$  for each 4 x 4 region. The  $E_X$  and  $E_Y$  values are summed individually in the Adder ASICs to produce total  $E_X$  and  $E_Y$  for the 256 tower region of the calorimeter covered by the crate.

The Jet Summary card also contains logic to search for tau candidates by testing for isolated hadrons. This logic will be based on the existing  $E_T$  sums

in  $0.35 \text{ eta} \times 0.35 \text{ phi}$  regions and the isolated electrons that are already present on the card. One of the tests that can be implemented is to determine if a high fraction of the total observed  $E_T$  in the  $0.35 \times 0.35$  region is contained in a single trigger tower. Such regions are candidates for isolated hadrons. Another test is an extension of the electron/photon algorithm and exploits the fact that the tau hadronic jet is characterized by a local e.m. energy deposit. Other algorithms are being studied and space for the implementation of such algorithms is reserved on the card.

Space has also been allocated for the luminosity algorithm described in section 1.6. This circuit may be moved to the global trigger processor, but until this design decision is determined we reserve the space on the Jet Summary card.

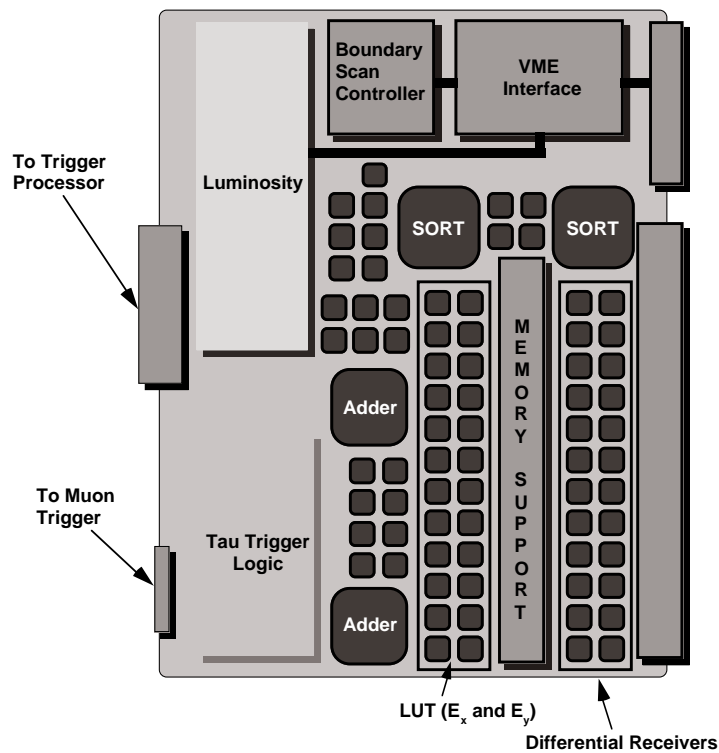


Figure 22. Jet Summary card.

## 6.1 Inputs

Each of the eight Electron Isolation cards transmits 6 bits of trigger rank every 6.25 nsec in differential mode. A single bit indicating which of the two  $4 \times 4$  trigger tower regions generated each rank is also included. Therefore, the Jet Summary card receives 7 differential pairs from each Electron Isolation card for a total of 112 input pins.

The Receiver cards send the 13 bits of  $E_T$  information for the two  $4 \times 4$  trigger tower regions differentially on 26 lines. This data is transmitted at 160 MHz, requiring two 6.25 nsec time frames. The third 6.25 nsec time frame contains

the Quiet and Minimum Ionization information for the two 4 x 4 regions. The remaining bits in the third and fourth 6.25 nsec time frames are at present reserved for future assignment. The Jet Summary card receives this data from the eight Receiver cards on 208 input pins. Therefore, the total number of signal pins dedicated to trigger data is 320.

## 6.2 Outputs

The outputs from the Jet Summary card appear on connectors at the front edge of the card. Transmission is via twisted pair cable at 40 MHz. Two connectors are shown in Figure 22, but the top connector to the global trigger processor is likely to be split into several separate connectors to segregate the individual quantities. The bottom connector is for the Muon Trigger and contains a Quiet region map of 16 bits (one bit per 4 x 4 region).

The global trigger processor receives 10 bits of  $E_T$ , 13 bits each of  $E_X$  and  $E_Y$ , 40 bits for the top 4 Jets, 24 bits of electron trigger rank, 4 bits of Jet  $E_T$  threshold information, 16 bits of location for the top 4 electrons, and 16 bits of information for the top 4 jets.

## 7. Latency

The overall latency of the level 1 calorimeter trigger regional logic design described in this document is approximately 13 crossings, or 325 nsec. This latency follows a latency of 4 crossings after the interaction for the physics signals to propagate to the calorimeter readout, 13 crossings for processing by the calorimeter front end electronics, 30 crossings to account for the maximum possible fiber optic cable length of 150 m for transmission of the data from the front end electronics to the calorimeter trigger logic in the electronics barracks and another crossing for synchronization of the data to the regional calorimeter trigger. We estimate that the transmission of the output of the regional logic information to the global calorimeter trigger logic will take another 4 crossings (including one crossing for synchronization). The result is that the calorimeter trigger information is provided at the input of the global calorimeter trigger approximately 79 crossings after the interaction occurred.

The 13 crossing latency of the regional calorimeter trigger logic is divided amongst several stages. Since the logic is mostly clocked at 160 MHz, the operations are composed of 6.25 nsec cycles, four of which add up to a single 25 nsec crossing time. The Receiver card uses 1.5 crossings after the synchronization stage divided into two 6.25 nsec cycles and a full crossing to accommodate crate to crate data transfers. The first 6.25 nsec cycle is used for the memory lookup table and the second is used to stage the information to the output connector. The data transmission across the backplane from the Receiver card to the Electron Isolation Card uses a single 6.25 nsec cycle.



The longest latency operation performed by the regional calorimeter trigger logic is that of the isolated electron identification and sort. Other data will be held in memory to wait to be aligned with these results. We therefore examine this function to determine the overall latency.

The Electron Isolation Card used 5.75 crossings to find and sort the isolated electrons. Most of this time is spent in the sort operation. The Isolation ASIC is fully pipelined and therefore is able to accept a new crossing's worth of data every 25 nsec. However, its overall latency is 1.5 crossings (37.5 nsec). The first 3 6.25 nsec cycles are used for the first 3 stages of the Isolation ASIC logic. The fourth cycle is used for the 4-way compare and the last 2 cycles are used to set up and place the results at the output stage for the Sort ASIC. The Sort ASIC then finds the 4 highest rank candidates of the 32 inputs in 4 crossings, accepting 32 new candidates every crossing. One more 6.25 nsec cycle is then used to stage the results to the output connector.

The transmission of the 4 highest rank electron candidates on the backplane from the Electron Isolation card to the Jet Summary card takes one 6.25 nsec cycle. The Jet summary card then sorts its total of 32 input electrons from 8 Electron Isolation cards through another Sort ASIC with a 4 crossing latency and uses another two 6.25 nsec cycles to stage and present the Crate's four highest rank electron candidates on the cable for transmission to the global calorimeter trigger. The total from Receiver card synchronization to output of the Jet Summary card is 12.25 crossings or slightly more than 306 nsec. We retain some additional time as contingency and give the Regional Calorimeter trigger latency as 325 nsec.

REFERENCES

---

- [1] The Compact Muon Solenoid Technical Proposal CERN/LHCC 94-38
- [2] RD-16 Status Report, CERN/DRDC 94-16
- [3] B.G. Taylor, RD-12, RD-27. Timing, Trigger and Control Distribution for LHC Detectors, ECP Division Working Document
- [4] M.A. Thompson. SDC Link Word Synchronization, SDC Note SDC-93-536
- [5] S. Dasu et al. CMS Level 1 Calorimeter Trigger Performance Studies, CMS TN/94-285 (1994)