

The Global Calorimeter Trigger for CMS

D.S. Bailey, G.P. Heath*, J. Lamblin, A. Mass, D.M. Newbold, U. Schäfer†, J.P. Scott
H.H. Wills Physics Laboratory, University of Bristol, Bristol BS8 1TL, UK

A.J. Maddox, V. Perera

CLRC Rutherford Appleton Laboratory, Chilton OX11 0QX, UK

Abstract: *The CMS calorimeter trigger logic is designed to search for objects, such as electron/photon or jet candidates, in local regions of the calorimetry. Objects found in different regions must then be compared with each other and the best candidates selected for onward transmission to the CMS trigger decision logic. This paper describes the design of the Global Calorimeter Trigger for CMS, which performs the selection. In particular we describe a low latency pipelined sort algorithm which has been developed, and preparations for the test of an ASIC implementation of the algorithm.*

1. Introduction

The Level 1 trigger logic for CMS is designed to search for specific signatures, which distinguish physics events of interest from the large background of hadronic events. In the calorimeter processing, the most important of these is the search for energetic electromagnetic showers from either electrons or photons. This part of the logic is known as the e/γ trigger. Other signatures of interest include high transverse energy jets and large missing transverse momentum.

The pattern recognition logic for the various types of object is housed close to the readout electronics for the calorimetry, in the Regional trigger processor crates[1–3]. Up to four different types of object will be searched for, including two separate processing chains for the identification of e/γ candidates. One of these will impose strict isolation requirements, with cuts optimised for the identification of electrons or photons from the decay of high mass objects such as the Higgs. The other is intended to search for electrons from the semileptonic decay of heavy quarks, which may be found relatively close to some hadronic activity in the calorimeters. The Regional trigger will also search for jet-type objects, and for isolated

single hadrons e.g. from τ decay. In addition, each Regional trigger crate will produce transverse energy sums. The Global Calorimeter Trigger (GCT) collects together all the data output from the Regional Trigger Crates. Its function is to extract the 4 highest ranked objects for each type and to complete the calculation of total and missing transverse energy, to be transferred to the Level 1 Trigger Processor.

This paper is organised as follows. In the following section we will introduce the components of the Global Calorimeter Trigger. Section 3 describes the link used to transfer data from the Regional crates to the GCT, as well as for inter module communication within the GCT crate. In Sections 4 and 5 we concentrate on the object sort processor and introduce a fast sort algorithm. Finally we report on the design of an ASIC implementation of the sort algorithm and a test system, which will be used to check the ASIC function at clock frequencies up to 200 MHz.

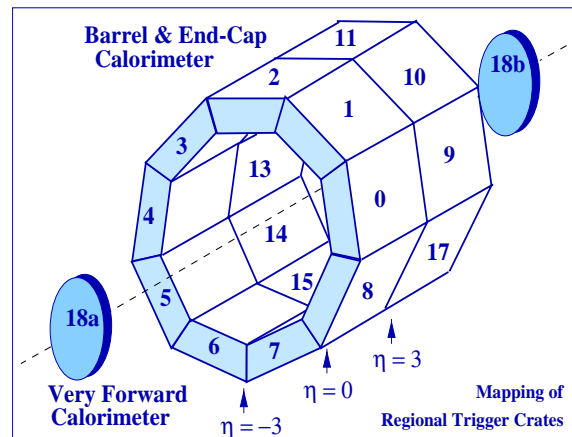


Figure 1. Mapping of the regional trigger crates onto the central calorimeter

*email: Greg.Heath@Bristol.ac.uk

†Now at: Institut für Physik, Johannes-Gutenberg-Universität Mainz, Germany

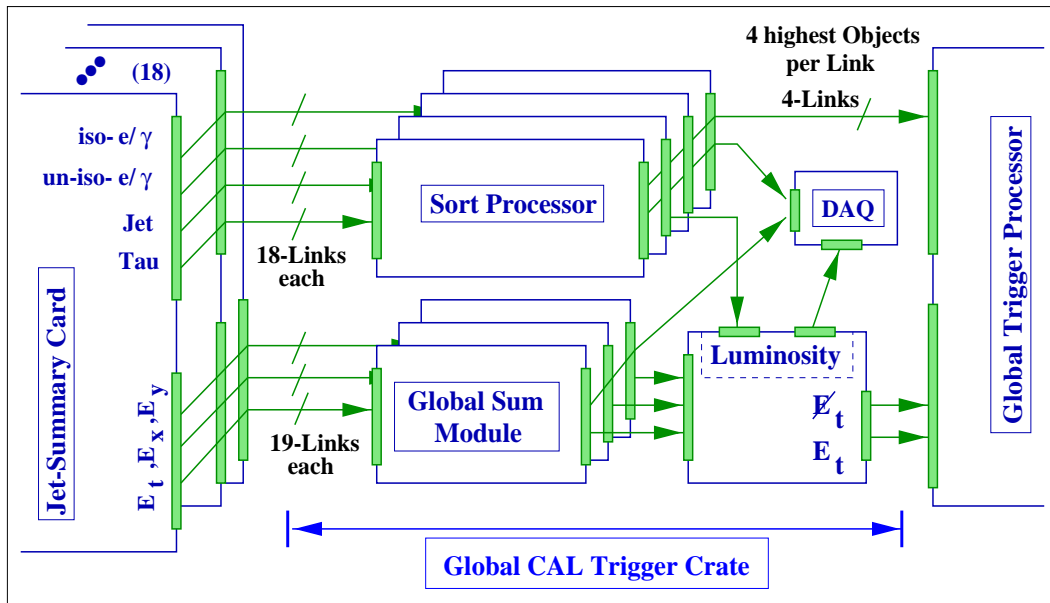


Figure 2. Block diagram of the CMS Global Calorimeter Trigger

2. Global Calorimeter Trigger

The input to the GCT is determined by the topology of the Regional Calorimeter Trigger electronics. The Regional trigger processor is distributed over 18 crates, with a mapping as depicted in figure 1. Data from the two very forward calorimeters, which contribute to the transverse energy sums only, are processed in a 19th crate.

Each regional crate receives data from the ECAL and HCAL front end electronics for (30×8) trigger towers (in the $\eta \times \phi$ -space). The pattern recognition logic identifies objects and assigns them to fixed, non-overlapping (4×4) trigger tower windows. This results in an initial map of (2×8) objects. The objects found in each Regional crate are collected on the Jet/Summary Card (JSC), the interface to the Global trigger. The object energy, plus additional information like isolation status or the fine grain bit, which identifies objects with electron type showers, are translated into a 6 bit rank using a lookup table. In order to reduce the number of interconnections 4 data words are time multiplexed on each link within the Regional crates. All object data processing in the Regional as well as in the Global trigger takes place at 160 MHz.

Within the JSC, the ranks are input to the first stage of a sort tree. The four highest ranks for each object type are identified and transferred, including a 4 bit location identifier to the GCT (see fig.2). This processor identifies the 4 highest ranked objects and sends them including a location identifier to the

global trigger. In addition, the number of jets above a threshold will be sent to an interface module in the GCT, to allow a fast online measurement of the luminosity for every bunch crossing separately.

Partial transverse energy sums in E_t , E_x and E_y are also calculated on the JSC. These data are transferred to a global sum module in the GCT, for E_t as a 12 bit unsigned and for E_x and E_y as a 12 bit signed number. Additional partial sums from the two very forward calorimeters (covering the range $3 < |\eta| < 5$) are received from a 19th crate. The total transverse energy values are calculated and forwarded to an interface module. The missing transverse energy is calculated using LUTs and data is scaled for the transfer to the global Level-1 Trigger.

All GCT modules store their input data for the level-1 latency time, and provide data for readout in the case of a level-1 accept. The Luminosity values are read out after an appropriate integration time.

3. Links between Calorimeter Trigger Modules

The GCT has to receive data for the sort processor from a large number of data sources running internally at 160 MHz. No commercial link technology is available to serialise data at this data rate. Hence parallel data transfer is unavoidable. The link length of up to 15 m, and the the expected spread of data arrival times, of ± 6 nsec, have led us to the decision to transfer data as differential ECL at the 40 MHz bunch crossing frequency. A compact con-

nector (CHAMP 0.8) has been chosen, which can transfer up to 68 differential pairs on a 4 cm high connector.

On the Sort Processor modules discrete ECL logic will be used to time multiplex data back to 160 MHz. Data transfer between modules within the GCT crate will use the same links to avoid the design of a custom backplane. Data for the transverse energy sum modules is only updated with the bunch crossing frequency. Since we need to utilise the full width of the chosen link, we have to reorder the partial sum data using a patch panel. Transverse energy data from up to 5 Jet Summary Cards will be combined into one link, resulting in 4 input connectors for the Global Sum Modules.

4. Sort Processor

The four highest ranked objects are extracted in two stages. Figure 3 depicts the general structure of the sort processor with three input sort modules and one final sort module. Each input sort module receives 24 ranks from six JSCs. The four highest ranked objects are determined and forwarded to the final sort module. This module identifies the absolute four highest amongst the remaining (3 × 4) objects and transfers them to the global trigger. The data link width is increased during the sort process by the numbers of bits needed to locate the sorted objects. This location identifier is 4 bits wide for data from the JSCs and 7 bits for data exchanged between input and final sort module.




Figure 3. Block Diagram of the Sort Processor

Figure 4 shows the Sort Module, which will be used in different configurations for the input and the final sort. In the input sort configuration six data links have to be received. In order to limit the amount of discrete ECL logic needed (order of 100 components), half of the link receivers are sourced out onto a rear Receiver Module. The data received at 40 MHz are time multiplexed onto a 160 MHz bus and transferred to a sort ASIC. The sorted output is sent to a mezzanine card, where it is either sent over a 52 bit wide link to the final sort module or on a 60 bit wide link to the Global Trigger.

All input data are recorded in a latency buffer for readout. In the case of a level-1 accept, data for three bunch crossing clock cycles are stored in a derandomizer buffer. ECL FIFOs are used for both buffer stages. Another mezzanine card is used for the interface to a Front-End Driver Module.

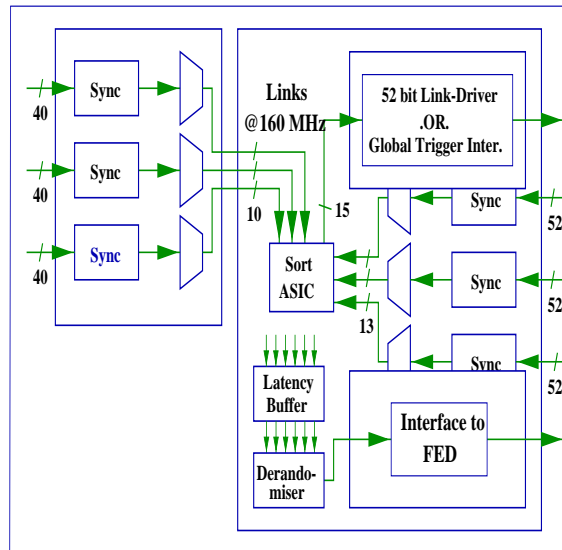


Figure 4. Block Diagram of the Sort Processor Module including the Receiver Module

5. Object Sort Algorithm

The object sort is accomplished in the calorimeter trigger in a 3 stage sort tree. A location identifier is created on every sort stage. A pipelined sort algorithm with 2 bunch crossing cycle latency has been developed. Its implementation on an ASIC will be employed on all three sort tree levels.

The number of input data links is six, corresponding to the number needed at the input sort module stage of the tree. Each input link receives four data words per bunch crossing cycle. The algorithm shown in figure 5 assumes that these four words arrive already sorted in descending order. The sort of the up to 24 data words is performed in two pipelined steps. First the highest ranked data words from the six input busses have to be compared with each other. The order is determined by counting how often a rank is higher than the others. The second step investigates the ranks of those objects which may be buried in the lower ranks of the input data. A multiplexer is used to select those data words which could contribute. The link with the highest rank may also contain data of the second, third and fourth highest ranked object. The link with the second highest object may contain data for third and fourth highest and so forth. The first sort stage already determines

the object with the highest rank. The other 9 objects remain to be sorted in the second step. This second sort stage can be implemented with 15 comparators, since the input data is already pre sorted and also only the four highest ranks have to be identified. A second multiplexer is used to select the four highest ranked objects and multiplex them in descending order on a 160 MHz output link. This algorithm requires a pre sort unit for the initial sort, where data is received unsorted.

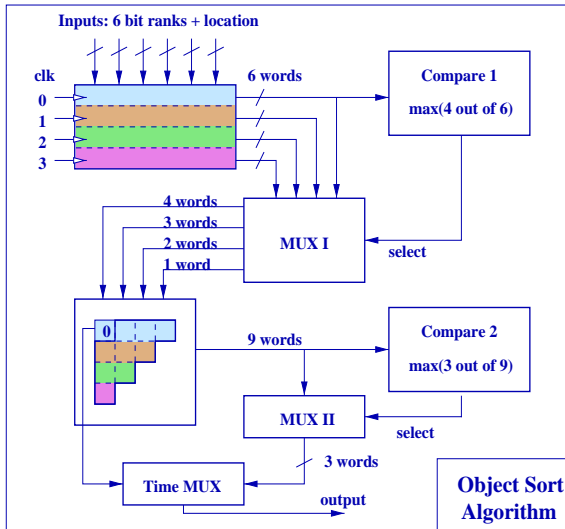


Figure 5. Object sort algorithm of Sort ASIC

The definition of the location identifier is closely related with the data origin. The JSC provides unsorted data for 16 ranked objects on 4 input links. The pre sort unit reorders the input data. The first 2 bits of the identifier specify the 160 MHz clock cycle of the initially received rank. The object sort algorithm appends always the number of the input link to the identifier. This means 2 plus 2 bits additional on the JSC. The input sort module provides data on 6 input links so another 3 bits are added and on the final sort module another 2. This adds up in total to a 9 bit identifier to be passed out to the Global trigger.

The ASIC implementation of the object sort algorithm will be configurable to work on all sort tree levels. The sort ASIC will include pre sort units for four input data links. All data I/O will be true single ended ECL. All I/O pads will be equipped with boundary scan registers.

6. Demonstrator Project

The goal of a demonstrator project is to show the ability to design the technically demanding components of the proposed system. In this case we need

to show that we are able to implement the sort algorithm as an ASIC and that it works at the required speed. We also have to demonstrate that we are able to handle a PCB design with 160 MHz clock and data lines for the Sort Module.




Figure 6. Prototype Sort ASIC

A prototype ASIC with the proposed object sort algorithm has been developed. The input data links are 11 bits wide. The data is composed of 8 bit ranks plus 3 bits of location identifier. The object sort circuitry was designed to receive 8 data links, instead of 6 as currently foreseen. However to avoid a pad limited design for the prototype, only the first four are connected to I/O pads. The unconnected data links are wired internally to a small, fixed rank value. The sort processor of the ASIC adds 3 bits to the location information for each output word to identify one of the eight input data links. No pre sort unit or configuration control is included in this version. Single ended PECL pads have been used for the I/O. The input data is translated into TTL, which is used for all internal logic functions. All I/O pads are latched in boundary scan registers. The circuit (figure 6) was developed as schematic design based on the AMS 0.8 μm BiCMOS process library and was successfully built in 1997. A Verilog simulation of the Sort ASIC design works above 200 MHz, with the unimportant boundary scan registers failing first.

We were initially only able to test the function of the ASIC up to 50 MHz on a commercial chip tester. The ASIC performed well and sorted the input data in the expected 2 bunch crossing cycles latency.

The next step was the development of a test system for the sort ASIC, which will be able to test the Sort ASIC up to 200 MHz. The system consists mainly of two modules. The Sort ASIC test mod-

ule, which is shown in figure 7, hosts four pseudo-random test pattern generators, a Sort-ASIC and time demultiplexer. The board receives a variable input clock frequency, which is quadrupled and which is used for all fast components on the board. All fast data links have been implemented as impedance matched lines. The demultiplexed output is connected to a compact connector, of the type which we intend to use in the final system. Data is received on a multi purpose memory module, see figure 8. The module is equipped with four dual-ported RAMs, which can be used to record data at variable frequency up to 50 MHz. The data are then read out over a custom field bus system. The memory board can also be used to provide test patterns, a feature which will be used in a later stage of the project.

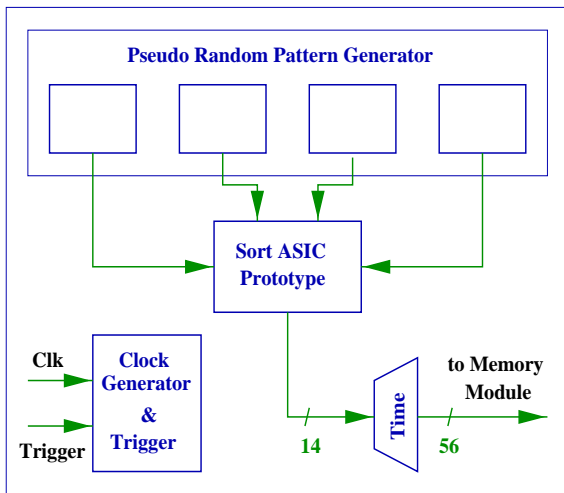


Figure 7. Test Module of the prototype Sort ASIC test system

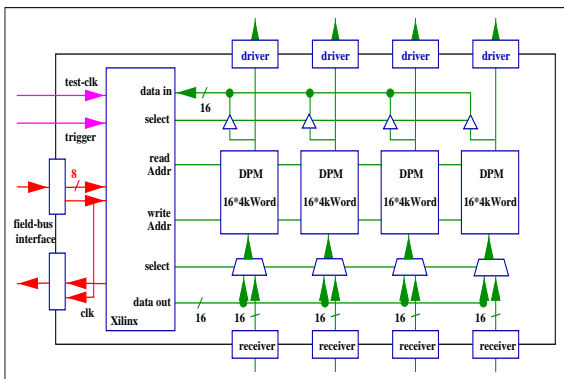


Figure 8. Memory Module of the prototype Sort ASIC test system

The PCBs for the test modules have been manufactured. The components on the boards are currently being assembled. The FPGA-firmware, driver

software and a graphical user interface for the custom field bus system have been developed. We expect to start the high speed tests in October.

7. Summary

The CMS calorimeters provide object information and transverse energy sums for use in the Level 1 trigger decision. The Global Calorimeter Trigger crate is required to collect and sort the objects found during Regional processing, and to complete the summing of energy components from the whole calorimetry. The design of this crate is constrained by the need to collect a large number of input data words. We have devised a fast, low-latency sort algorithm which has been implemented in an ASIC. We have also developed a modular test system to exercise trigger components at LHC clock speeds. This system is currently being used to test prototype ASICs, and first results are expected shortly.

REFERENCES

1. CMS calorimeter trigger group, G.P. Heath *et al.*, Preliminary specifications of the baseline trigger algorithms, CMS-TN/96-10 (1996)
2. G.P. Heath *et al.*, The CMS calorimeter trigger, in Proceedings of the Third LEB Workshop, CERN/LHCC/97-60, p. 393 (1997)
3. W.H. Smith *et al.*, High speed data processing for the CMS calorimeter trigger, in Proceedings of the Third LEB Workshop, CERN/LHCC/97-60, p. 408 (1997)