

The CMS Global Calorimeter Trigger Hardware Design

M. Stettler, G. Iles, M. Hansen ^a,
C. Foudas, J. Jones, A. Rose ^b

^a CERN, 1211 Geneva 23, Switzerland

^b Imperial College of London, UK

Matthew.Stettler@cern.ch

Abstract

An alternative design for the CMS Global Calorimeter Trigger (GCT) is being implemented. The new design adheres to all the CMS specifications regarding interfaces and functional requirements of the trigger systems. The design is modular, compact, and utilizes proven components. Functionality has been partitioned to allow commissioning in stages corresponding to the different capabilities being made operational. The functional breakdown and hardware platform is presented and discussed. A related paper discusses the firmware required to implement the GCT functionality.

I. OVERVIEW

The CMS Global Calorimeter Trigger (GCT) is an integral portion of the CMS trigger system [1]. Its function is to receive and process data from all 18 Regional Calorimeter Trigger (RCT) crates and send the highest ranking electron and jet triggers to the Global Trigger, where they are used for generating the First Level Trigger Accept (L1A) decision. This system has been designed to be modular in design as well as commissioning. Due to a compressed development schedule, it borrows heavily from existing designs.

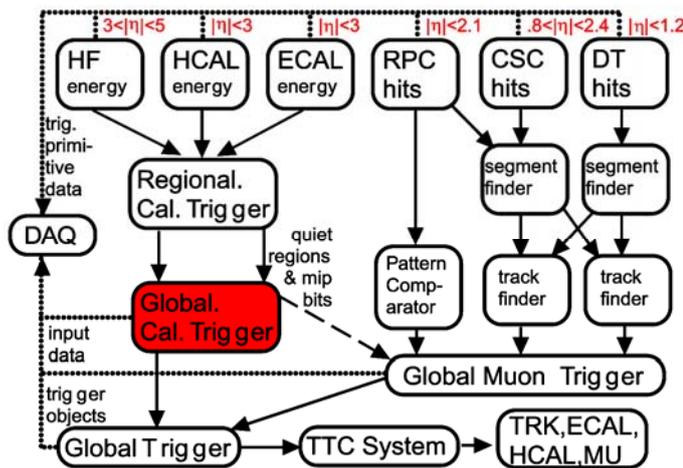


Figure 1: GCT's position in the CMS Trigger system

The GCT is composed of four module types, the Source, Leaf, Wheel, and Concentrator cards.

Source cards receive input directly from the RCT crates, and transmit the data via multi-Gigabit optical links to the GCT crate. They are located in the same racks as the RCT, and provide differential ECL to high speed serial conversion. In addition, they provide a means of data capture and readout for the RCT.

The Leaf cards are configured to receive either electron or jet trigger data on high speed optical fibers. Each electron leaf card processes the electron data from 9 RCT crates, selecting the four highest energy candidates for further processing. Similarly, the jet leaf cards process data from 3 RCT crates, and forward the 4 highest energy jets to the Wheel cards. However, the jet finding algorithm implements a sliding window – which requires data from adjacent RCT crates (corresponding to adjacent physical areas on the detector). Jet leaf cards are linked to their neighbours in a corresponding fashion to facilitate this algorithm.

The Wheel cards are only used for processing jet data, and combine the output of 3 Leaf cards. These 3 Leaf cards process the data from 9 RCT crates, or 1/2 of CMS. The Wheel cards sort the jets generated by the Leafs, and forward the 4 with the highest energy to the Concentrator.

The Concentrator card accepts data from 2 electron Leaf cards and 2 Wheel cards. It performs the final sorting of both electron and jet events, and sends the 4 highest energy candidates of each type to the Global Trigger. In addition, it provides a VME interface (slow control interface), and S-link data acquisition interface for the entire system.

II. REQUIREMENTS

The primary requirements of the GCT are to sort electron triggers, and generate and sort jet triggers. The design has been optimized for these tasks. Secondary requirements include jet trigger counters, total jet transverse energy trigger, total transverse and missing transverse energy trigger, luminosity monitoring, and RCT readout. An in-depth discussion of all the GCT requirements is far beyond the scope of this paper, and are covered in detail in the CMS Trigger TDR [1] and several subsequent CMS internal notes [2,3]. The focus here will be on the primary requirements which have driven the hardware architecture.

A. Jet Trigger

The primary driver of the GCT hardware design is the need to find and sort jet triggers in the RCT data. The GCT jet finder uses a 3x3 sliding window (12x12 trigger towers) to search for patterns of energy deposition consistent with jets. This sliding window requires a contiguous data space, crossing the data boundaries of RCT crates. This implies a data sharing scheme within the GCT that matches the physical configuration of the detector. The detector data is processed by the RCT in 18 longitudinal sections, each 1/2 the length of the barrel and spanning 40degrees in phi. Since the data space spanning the center of the detector (eta 0) must be contiguous

as well, a similar sharing scheme needs to be implemented between the center sections of each $\frac{1}{2}$ barrel. Once identified, jet objects need to be sorted by Et. The 4 highest energy candidates are forwarded to the GT.

B. Electron trigger

In comparison to the jet trigger, the electron trigger is straightforward. The RCT implements a sliding window algorithm similar to the jet algorithm described above and forwards 4 candidates (highest Et) of both isolated and non-isolated electrons. The electron trigger candidates simply need to be sorted and the 4 highest energy candidates sent to the GT.

C. Latency

The GCT has a limited amount of time to accomplish its task of jet identification as well as jet and electron sorting. This is driven by the limited storage in the front end electronics, which can store data from 128 events (which corresponds to 128 40 MHz clocks). The entire trigger process needs to complete in this time period, and the GCTs allocation has been defined as 24 40 MHz clocks.

III. HIGH LEVEL DESIGN

The GCT is a large data processing system, which needs to interface directly to 18 high power ECL crates spanning an entire row of racks located on a different floor or the CMS equipment area. RCT output data needs to be collected and concentrated for efficient transport and processing. The processing of the data needs to be highly parallelized and as efficient as possible. In addition, the design is subject to extreme schedule pressure, and needs to be modularized to both reduce risk and allow efficient parallel efforts in both the hardware and software teams. Lastly, the compressed schedule demanded that the design be based as much as possible on existing modules.

A. Basic design Methodology

Due to the compressed schedule, it is critical to reduce risk by designing conservatively. In the GCT this has taken the form of data transfer rates that are well within the capability of the chosen technologies and minimizing the number of FPGAs in the processing chain. While the reasoning for conservative data transfer rates is obvious, reducing the number of FPGAs may not be. It is well known that the cost of FPGAs is not linearly related to their capacity – so it is almost always less expensive to design with multiple smaller devices than one large one. However, such designs usually have a much higher firmware risk. This is due to difficulty of accurate simulation, poor partitioning efficiency by synthesis tools, and the additional delays introduced when moving data between chips. This additional firmware risk can be substantial, and easily offsets the lower hardware cost in low production or one of a kind systems.

B. Data Transport and Concentration

Interfacing to the RCT system involves several challenges. Electrically, its output is differential ECL, and the data format

is not DC balanced. This implies maintaining an accurate ground reference to 18 ECL crates, which are consuming over 10KW of power. In addition, the data is presented on a total of 108, 68 pin cables. Simply providing panel space for these connectors can be a problem. To provide both electrical isolation and concentration it was decided to build a data transport card, utilizing multi-Gigabit serial optical links to transport the data from the RCT racks to the GCT. Each 68 pin cable is replaced by two fibers, allowing a very high data density input to the GCT. This data transport and concentration card is known as the source card, and was derived from the IDAQ data acquisition module designed by Imperial College (UK).

C. Processing

The RCT produces a massive amount of data corresponding to over 233 Gbps. To process this amount of data, a hardware based signal processing system is required. In addition, much of this data needs to be shared (locally) within the system, so concentrating as much data as possible in individual processing elements is highly desirable.

A processing module was designed to meet this requirement, based on the largest and fastest Xilinx FPGA with embedded Multi-Gigabit serial links available. While these devices are quite expensive, they provide a high level of integration and unmatched data density. This reduces risk in both the hardware design and firmware development. Two of these devices (Xilinx Virtex2Pro70s), which support 16 serial links each, were designed into a module with multi channel fiber receivers to support a total of 32 links on a double PMC. This processing module is known as the Leaf card, and was derived from a satellite processor prototype designed by Los Alamos Laboratory (US).

The most elegant design for the GCT would consist of a chain of these processing modules, each processing and sending serialized reduced data to a following layer. Unfortunately, there is a latency cost involved in serializing and deserializing data, and the integrated deserializers in FPGAs are particularly inefficient. In the case of the GCT, only one level of data concentration by serialization can be tolerated if the latency budget of 24 clocks is to be maintained. The remainder of the processing, which consists almost entirely of sorting, is accomplished by large (9U VME) carriers with wide parallel interface busses.

Two types of carriers have been designed, corresponding to primary jet trigger sorting, and electron and final jet trigger sorting. These are new designs developed by the GCT team at CERN, and have been designed to be compatible with the existing ECAL backplane design.

The primary jet trigger sorting carrier is known as the wheel card, and it supports the unique requirements of the jet finder algorithm. The data from 9 RCT crates that cover $\frac{1}{2}$ the detector barrel is processed by 3 Leaf cards arranged in a ring topology. Each Leaf card shares data with its neighbours, providing the contiguous data space required by the jet finder algorithm. A second set of three Leaf cards is required for the remaining $\frac{1}{2}$ barrel. To provide the required overlap between the two barrel halves, data from the first 2 regions in eta for each $\frac{1}{2}$ barrel (the central region of the detector) is duplicated

and sent to the opposite side for processing. Each wheel carrier supports 3 Leaf cards, and processes jets from $\frac{1}{2}$ of the detector. The wheel receives the jet objects from 3 Leaf cards and sorts them by energy. The resulting data is sent to the concentrator over 3, 80 pair differential cables.

The last stage of the GCT processing chain is the Concentrator card, which directly supports the two Leaf cards required for electron sorting, and receives data from two wheel cards for final jet sorting. In addition, the Concentrator provides interfaces to the GT, the data acquisition system, and slow control (VME). While algorithmically simple, the Concentrator is a complex design involving multiple system interfaces.

D. Overall System

The overall GCT system consists of the four types of modules discussed. There are 54 Source cards, 8 Leaf cards, 2 Wheel cards and one Concentrator card in the final system. A high level block diagram of $\frac{1}{2}$ the GCT (which is symmetric), is shown in figure 1.

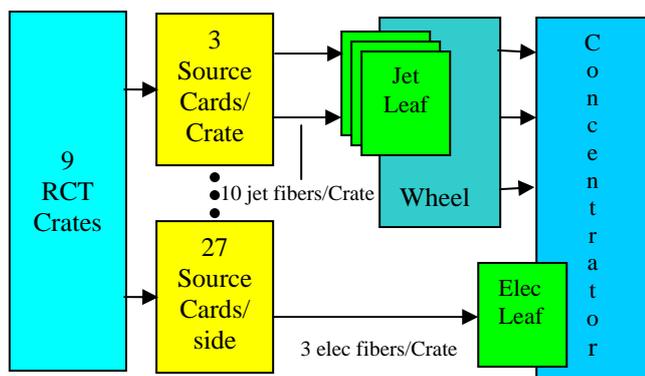


Figure 2: $\frac{1}{2}$ GCT

IV. SOURCE CARD

The Source card is responsible for receiving data from the RCT, and transmitting it to the Leaf cards for processing. In addition to this primary function, it is also used to capture RCT data for diagnostic purposes. Each Source card interfaces to two RCT input cables, and 4 optical fiber outputs. A CMS standard TTCrx is used to relock the RCT data and check synchronization. The Source card is implemented in a 6U VME format, and multiple sets of Source cards are housed in VME chassis installed in the RCT racks.

The IDAQ data acquisition card developed by Imperial College was the original choice for the source card function. However, due to the number of cards required in the final system, a simplified design derived from the original IDAQ with only the required functionality was implemented. This retained most of the support software and related firmware, as well as many components and general design elements.

The Source card logic is implemented in a single Xilinx Spartan 3 FPGA (3S1000), which performs all the data routing and readback functions. Data serialization is accomplished by 4 TI TLK2501 SERDES devices, which also

handle the 8B/10B encoding to insure a DC balanced data stream. Optical outputs are provided by 4 standard SFP modules, in this case Agilent HFBR-5720AL devices. The Source card prototypes have successfully completed testing and are waiting on Leaf card availability for final link tests.

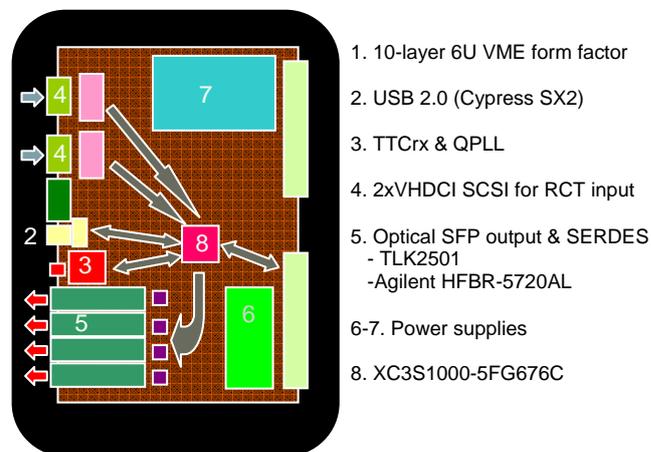


Figure 3: Source Card Block Diagram

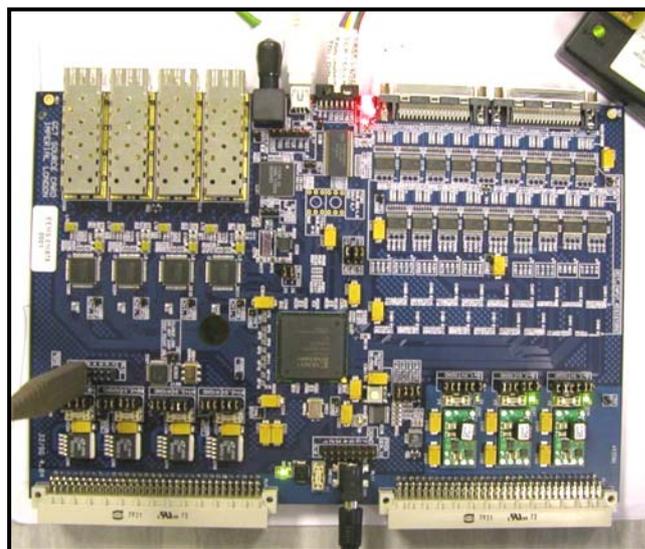


Figure 4: GCT Source card

V. LEAF CARD

The Leaf card is the main processing module in the GCT, and is used for both jet and electron triggers. It can also be used for data snapshots of individual RCT fiber inputs, and to generate test data for Wheel and Concentrator verification.

The processing capabilities of the Leaf card are driven by the requirements of the jet finder algorithm. This algorithm, described in a related paper [4], requires that the data from an entire RCT crate, plus a portion of its neighbours, be available to search for patterns of energy deposition consistent with jets. The logic footprint of the search operation, which is highly parallelized, is more than 3 million gates. In addition, the jet data from a single RCT requires 8 optical links at 1.6 Gbps to transport. The most efficient implementation of this algorithm is to tightly couple adjacent jet finders, preferably in the same chip. The Leaf card design is an attempt to balance the maximum density of multi gigabit optical inputs

and available hardware processing power required to perform this task.

Practical matters of design and schedule risk heavily influenced the choice of processing components and form factor of the module. Early on it was recognized that it would be optimal for the processing FPGAs to have integral high speed serial links, as opposed to external deserializers. The internal devices are far more power efficient, and the intimate connection to the processing fabric of the FPGA gives considerably more design flexibility. Xilinx V2P70 devices were chosen as the best fit for density of high speed serial interfaces and logic capacity. Devices of this size require a reasonable amount of support on the PC board in the form of multiple high current regulators, filtering networks, and precision clock networks. Designing such systems capable of supporting the FPGA near the limit of its performance is a non trivial effort. Fortunately such a design existed in double PMC, a satellite processor prototype (a digital channelizer) designed at Los Alamos Laboratory. This device hosted two of the largest Virtex2 FPGAs and external serializers, and had been verified under high load conditions. A lesson learned from this previous design was that the designs power density (~60W) on a double PMC could not be easily pushed higher. Achieving the required density of fiber optic inputs was not possible with standard SFP modules. Moving to another industry standard, SNAP-12 MFP, provided a good match between serial ports available on the V2P70 and available front panel space.

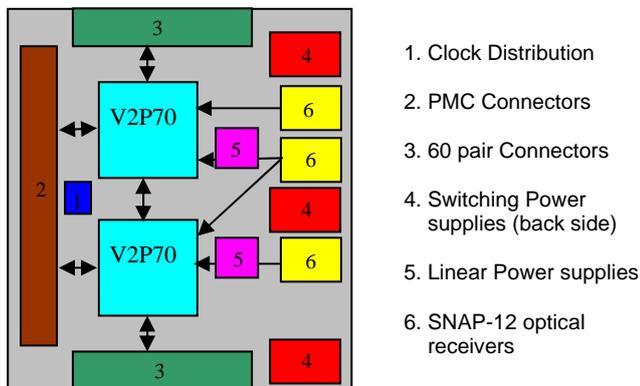


Figure 5: Leaf Card Block diagram

The Leaf card is implemented with 2 Xilinx V2P70 devices, and 3 SNAP-12 MFP optical receiver modules on a modified double PMC form factor. In order to facilitate sharing of data between adjacent jet Leafs, each card supports a 60 pair LVDS connection on the top and bottom edge of the card. The differential clock tree includes a high stability local 80MHz reference as well as multiple inputs from both PMC and micro coaxial connectors. The power supply system includes both phase controlled switchers for FPGA logic and linear regulators for the high speed serial links. The Leaf also includes two flash PROMs to configure the FPGAs, which are configured in a single JTAG chain which is jumper selectable between a header and the PMC connectors.

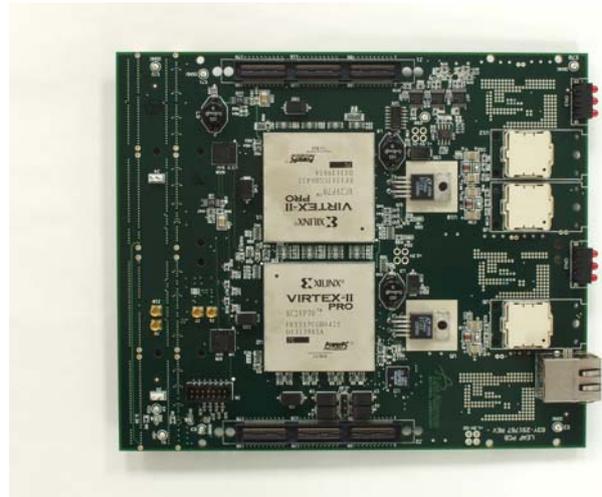


Figure 6: Leaf card

VI. WHEEL CARD

The Wheel card carrier is used to support three Leaf cards for jet processing and each Wheel processes the data from $\frac{1}{2}$ the detector barrel. The Wheel card is functionally straightforward, sorting the jet objects found by the Leafs, and forwarding them to the Concentrator for the final sort. It also supports the slow control, synchronous trigger, and JTAG interfaces to the Leaf cards. The main design challenge on the Wheel was to manage a large number of parallel busses required to transport the jet object data to the sorting FPGAs, where pin limitations drove the selection of large Xilinx V4LX100 devices. In this case, the desire to be conservative in bus speed complicated the hardware design, since the design point of 40 MHz DDR directly relates to the number of I/O pins required. Experience with the Leaf has shown that 100 MHz can be safely transferred through the PMC connectors, so it is likely that the system is considerably over designed.

The Wheel card is implemented as a double wide 9U VME PMC carrier, with a total of 3 double PMC sites (two on the front and one on the back). The Wheel card is compatible with the existing custom 9U backplane used for CMS ECAL. The jet object sorting is handled by one Xilinx V4LX100, and energy summation and control by a second V4LX100. Communication to the Leaf cards is via the PMC connectors, which are fully connected to the FPGAs. A fully differential clock distribution system is implemented, which provides local 40 and 80 MHz clocks (generated by a QPLL), and external clock inputs from the Concentrator. The Wheel communicates with the Concentrator via 3, 80 pair differential cable assemblies. Once again, the design is perhaps over constrained here. The data rate is defined as 40 MHz DDR, but the cables, connectors and electronics are rated for over 4 times this. The JTAG chain, which can include all Leaf cards as well as the local FPGA and associated PROMs, is controlled by a dedicated CPLD, with manual front panel configuration. This allows each Leaf and Local FPGA to be independently added to the chain, and also selects either a front panel header input or a link to the Concentrator.

VII. CONCENTRATOR CARD

The Concentrator Card is the heart of the system, and where all final sorting of jet objects and electrons is done. It carries the electron sorting Leafs directly on two double PMC sites, and interfaces to two Wheel cards via 3, 80 pair differential cable assemblies each. Output to the GT is provided on a double PMC site which houses a multi channel serializer compatible with the receivers on the GT. The Concentrator also provides the system VME, Slink DAQ, and synchronous timing (TTC) system interfaces. As on the Wheel card, the main design challenge is managing all the large parallel busses required to pass the electron and jet object data at 40 MHz DDR. Once again, the desire to be conservative in data rate complicated the hardware design significantly.

The Concentrator is implemented as a double wide 9U VME PMC carrier with three double PMC sites, two on the front and one on the back. It is also compatible with the CMS ECAL backplane. Electron and jet object sorting is handled by two Xilinx V4LX100 FPGAs. The VME and DAQ interface is provided by a Xilinx 2V3000 FPGA. The selection of large vitrex 4 devices was driven by pinout requirements of the data busses, and these devices are only lightly used to implement the sorting function. A fully differential clock tree is provided which includes a QPLL driven by a TTCrx, which is intended to be the timing reference for the entire GCT system. Both 40 and 80 MHz clocks are fanned out differentially to each Wheel and all PMC sites. The JTAG chain is controlled in a similar manner as the Wheel, with a CPLD controlling a segmented chain. The JTAG chain is also accessible by software, allowing the entire GCT to be configured under software control. In order to prevent JTAG errors from disrupting the VME interface, the VME interface FPGA is on a dedicated chain not accessible by software.

VIII. CURRENT SYSTEM STATUS

Development of the GCT hardware has been proceeding in parallel, and first articles have started to arrive over the last few months. The two designs based on existing work, the Source and Leaf cards, are in hand now and under test. Of the two new designs, the Concentrator is in manufacture (due in mid October), and the Wheel is in final layout.

The Source card has been tested extensively, and has met its performance requirements. It is awaiting integration with the Leaf before entering production.

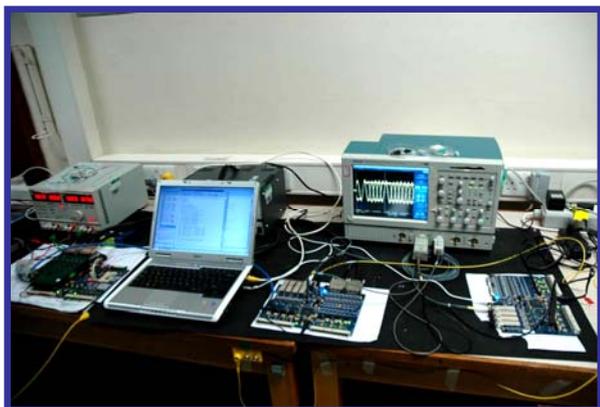


Figure 7: Source card under test at Imperial College

The Leaf card has just started integration tests, and is behaving as expected at this point. The test fixtures and software is in place to start the high speed serial testing, and this critical phase of testing will begin shortly.



Figure 8: Leaf card under test at CERN

The firmware modules for the GCT have been in continuous development and test (both simulation and hardware) for six months. The firmware for several revisions of the Leaf jet finder algorithm has been simulated and successfully tested in hardware on the Los Alamos satellite board. A larger simulation of the Concentrator card and electron Leaf firmware is in progress. The slow control interface has been simulated and proven in initial Leaf card testing.

IX. CONCLUSION

The CMS GCT effort is well under way, with nearly all required hardware in hand or in production. The schedule is optimistic, with 7 Source cards, 2 Leafs and the Concentrator scheduled for installation in the first months of 2007. Integration and testing will be continuing toward this goal.

X. REFERENCES

- [1] The Trigger and Data Acquisition Project, Vol. I, The Level-1 Trigger, CERN/LHCC 2000-038, CMS TDR 6.1, 15 December 2000.
- [2] CMS-IN 04-009, Updated Interface Specification for the CMS Level-1 Regional Calorimeter Trigger to Calorimeter Global Trigger.
- [3] CMS-IN 02-069, CMS Level-1 Global Calorimeter Trigger to Global Trigger and Global Muon Interfaces.
- [4] Revised CMS Global Calorimeter Trigger Functionality & Performance, G. Iles, C. Foudas, G. Hall, J. Jones, A. Rose, Imperial College London, UK, J. Brooke, G. Heath, Bristol University, UK, M. Hansen, J. Nash, M. Stettler, CERN, Switzerland, LECC 2006