Hardware & Firmware for the CMS Global Calorimeter Trigger.

J. Brooke, D.Cussans, R. Frazier, G.Heath, D. Machin, D.Newbold

H H Wills Physics Laboratory, Tyndall Avenue, Bristol, BS1 8TL, United Kingdom
jim.brooke@cern.ch

S. Galagadera, S. Madani, A.Shah

Rutherford Appleton Laboratory, Chilton, Didcot, OX11 0QX, United Kingdom

Abstract

The use of FPGA technology allows the CMS Global Calorimeter Trigger (GCT) to perform a wide variety of data processing tasks with a single configurable trigger processor module. A recent addition to the design of this module is the inclusion of 3.2 Gbit/s serial links to transmit data between modules over a 'cable backplane'. This allows the CMS Level-1 jet-finding algorithm to be implemented in the GCT; a function that would previously have required additional hardware. The cable backplane concept has been verified with custom test hardware. A consequence of the extended GCT design is the need for different firmware in most of the 190 FPGAs in the system. Production of this firmware requires a modular approach and careful code management. A 'build tool' that automates part of the firmware process is described, along with requirements for other scripts and software needed to complete the firmware for the GCT.

I. Introduction

The structure of the CMS Level-1 Trigger is illustrated in Figure 1; greater detail is given in the Technical Design Report [1]. The calorimeter trigger chain is described briefly below.

The digitised calorimeter data are used to generate Trigger Primitives, of coarser granularity than the data read out by the DAQ system. These primitives are then processed by the Regional Calorimeter Trigger (RCT) system [2]. Here, electromagnetic (EM) shower finding algorithms are implemented and local transverse energy sums are calculated. The EM candidates (both isolated and non-isolated) and region energies are passed to the Global Calorimeter trigger, where jet finding is performed and global E_t sums are calculated. The EM candidates and jets (in three categories; forward, central and tau) are then sorted by transverse energy and the highest four in each category are passed to the Global Trigger along with the global E_t sums. The Global Trigger compares all input data (including muons) with the trigger tables and the final Level-1 decision is made.

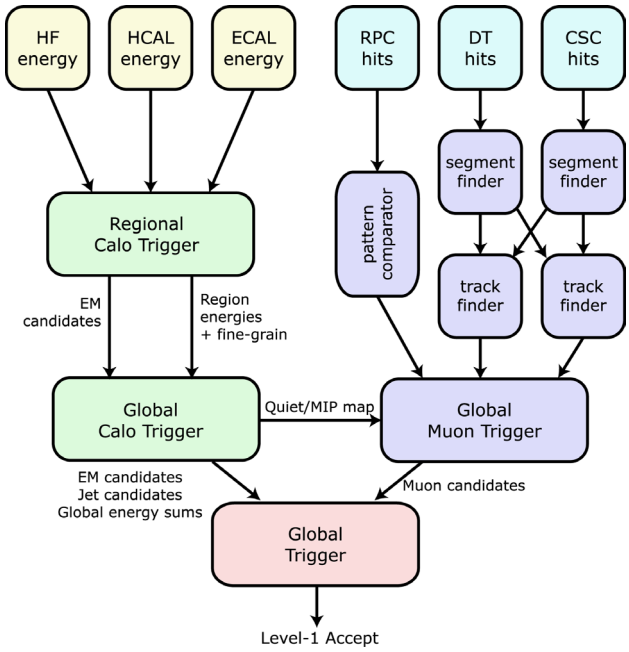

Figure 1: The CMS Level-1 Trigger

II. GCT System Overview

The design concept of the CMS Global Calorimeter Trigger is given in [3]. A single configurable processor module is used to perform all the processing requirements. At the time of publication of [3], these consisted of the sorting of jet and EM candidate streams, calculation of global energy sums and luminosity monitoring. Fast serial links are employed to obtain sufficient bandwidth in and out of the module, and FPGA technology is used for data processing. Different FPGA programs allow a single module to perform the range of processing required.

The current design extends the initial proposal by including point-to-point serial links between modules via a 'cable backplane'. This allows the GCT to perform jet finding (as described in [2]), which would otherwise have required the construction of a special 'cluster crate'.

The GCT is constructed from 18 Input Modules, housed in two crates, and 9 Trigger Processor Modules, housed in a central processing crate containing a point-to-point cable backplane. The processing crate also contains a VME controller and a Communications Module (CM), which

provides interfaces to the TTC and DAQ systems. The design of IM, TPM and the cable backplane are outlined below. At the time of writing, prototypes for each of the modules and backplane have already been produced or will be received soon.

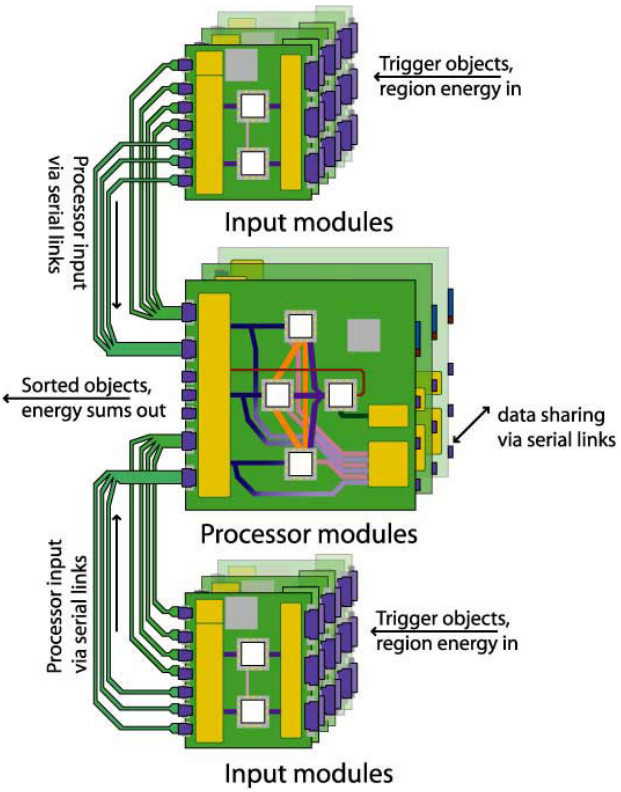

Figure 2: Overview of the GCT hardware

A. Input Module (IM)

The RCT to GCT link consists of 108 34-pair cables, using ECL levels at 80 MHz. Data on six cables is received by each IM, and serialised for transmission to the Trigger Processor Modules (TPMs) or, in the case of the muon feature bits, to the Global Muon Trigger. National Semiconductor DS92LV16 serdes (serialiser/deserialiser devices) operating at 1.44 Gbit/s are used for serialisation. Two FPGAs are used on each IM to perform some routing of the incoming data, and implement spy/playback buffers for debugging and monitoring purposes.

B. Trigger Processor Module (TPM)

The TPM is capable of receiving data on 24 1.44 Gbit/s serial links from the IMs via the front panel. This data can be sent to one of three 1st stage processor FPGAs or the cable backplane. Each 1st stage processor FPGA can also receive data from the backplane. Data coming out of the 1st stage processors is sent to a 2nd stage processor, which also has a connection to the backplane. Finally, the 2nd stage processor sends output data to six 1.44 Gbit/s links on the front panel. The complex data paths around the board are necessary to implement the full jet finding algorithm.

Two TPMs are used to sort the EM candidates (one per stream). Six TPMs are used to implement the jet finding and energy sum calculation. All eight of these TPMs send DAQ data from the spy buffers to a further TPM. This TPM is used to assemble DAQ packets, which are forwarded to the CM, which houses an SLINK 64 interface to the DAQ system. The CM also receives the TTC fibre and distributes the LHC clock and decoded signals to the TPMs.

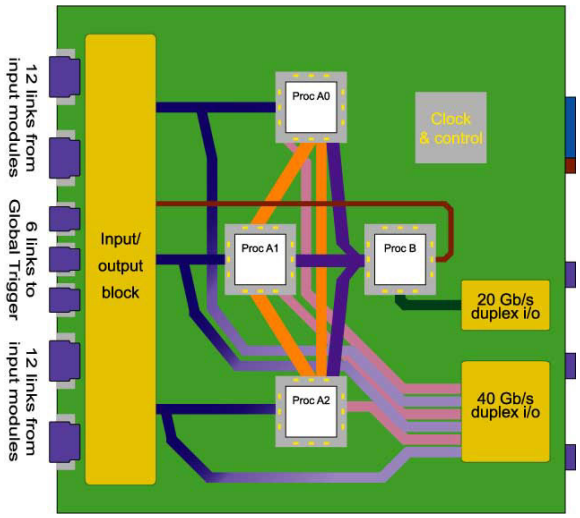

Figure 3: Trigger Processor Module

C. Cable Backplane

Communication between TPMs is implemented via serial links over a cable backplane. Vitesse 7226 serdes are used to provide 3.2 Gbit/s serial data. The backplane houses Teradyne VHDM-HSD connectors on both sides, press-fitted into a single set of holes. The TPM inserts into the front connectors, while Gore Eye-opener 6 cables are used at the rear to complete the desired point-to-point links. Since the only available cables incorporate PTFE dielectric material, for optimum high-speed signal performance, the entire backplane is enclosed in an airtight steel enclosure to prevent escape of halogens in the event of fire.

III. Cable Backplane tests

A small PCB was prepared with two VHDM-HSD connectors on either side, as shown in Fig 4. At the time of the tests, press-fitting the connectors into a single set of holes was not possible, so the high-speed signals were carried on short tracks between the front and rear connectors. This will not be required for the final backplane, as the connectors will be press-fitted into a single set of holes. Two test boards, each containing an FPGA, a Vitesse 7226 and a VHDM-HSD connector were also built to drive the link. SMA connectors were included in the serial data signal path to facilitate measurement of the signal.

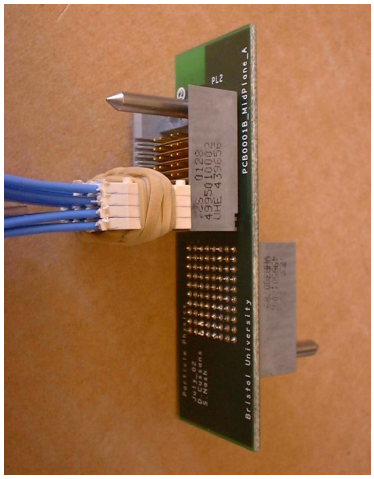

Figure 4: Cable backplane mock-up.

Using one board to transmit pseudo-random data patterns to the other, a data-eye shown in Figure 5 was obtained from the receive end of the link. As can be seen, the eye is well 'open'. The link was left to run for 2 hours, during which time over 10^{13} bits were transmitted, and no errors were observed.

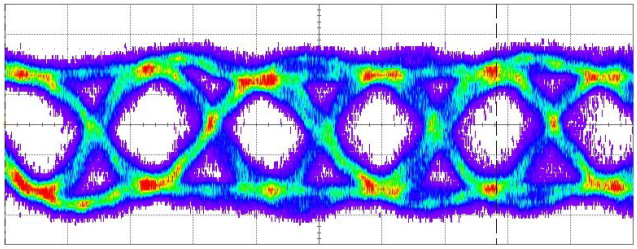

Figure 5: Data-eye at the receive end of a cable backplane link. The time scale is 200ps/div.

IV. FIRMWARE PRODUCTION

Production of the firmware for the GCT is complicated by the number of programs required. The full GCT contains 190 FPGAs, and at least 60 of these will require significantly different programs. In particular, the 32 processor FPGAs generally have entirely different algorithms. The I/O FPGAs have similar programs, but the data routing in each may differ. The firmware is therefore produced using modular VHDL code. Common functionality to be encapsulated in a module which may then be used throughout the system where required. A CVS repository is used to maintain coherent versions of the code.

One particular module is shown in Figure 6. This illustrates the firmware for a Processor FPGA. The spy and playback buffers for debugging, monitoring and DAQ purposes are required in all Processor FPGAs, though the data routing and specific algorithm differ. A simple interface has been defined for the Algorithm modules, which provides input, output and a set of clocks and clock enables. This allows the data processing firmware to be developed entirely separately from the technology dependent memories etc.

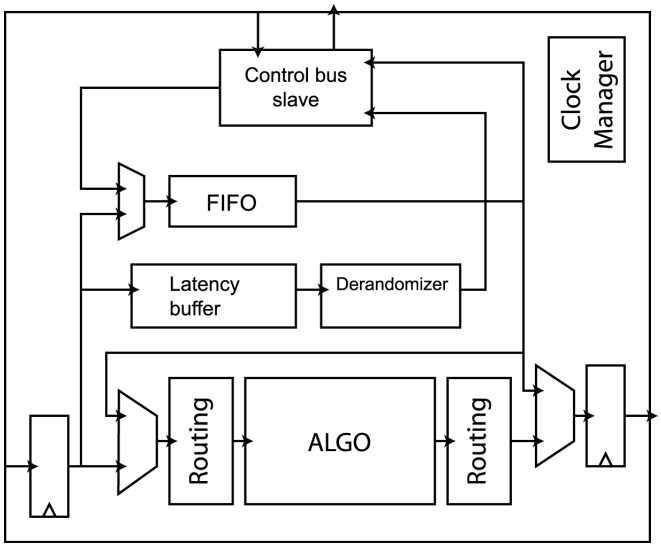

Figure 6: Schematic of the firmware for a Processor FPGA.

In addition to modularisation and version management, the production of a bit stream from the VHDL must be controlled and recorded to ensure the process may be repeated. This is achieved using a set of scripts to control the software tools used; these are currently Mentor Graphics Spectrum for VHDL compilation and Xilinx ISE 5.2i for placement and routing (PAR) of compiled designs.

The firmware process is illustrated in Figure 7. The 'source' files (VHDL and IP core descriptions) are on the left of the diagram, and the 'control' files are at the top. These control files are simple text files, stored in CVS alongside the source code. The 'library files' describe the file hierarchy of the VHDL code. The 'technology' file describes the target device and technology. The most complex file is the 'configuration' file, which describes, in very abstract terms, which modules are to be used. In particular, this file states which algorithm is to be used in a processing FPGA, and how the input and output data is to be routed. One configuration is required for each firmware program, although one file may contain several configurations. The files produced by the software tools are on the right of the diagram; essentially these are either netlists (which may be used as part of a library for other configurations) or the programming files (i.e. bit streams).

The main firmware production script (called 'fbuild') processes the control files to produce a Makefile, which contains relevant commands for the compiler and PAR tools and handles the file dependencies. Running `make` with this Makefile will then produce the required programming file. The script has been generalised wherever possible, to support use of different software tools. Support is also included for use of 'flow files'; scripts specific to a particular software tool that describe the processing required in detail. Hooks are provided in fbuild to allow such a flow file generator to access the information contained in the control files. Flow files are currently used with Leonardo Spectrum for VHDL compilation, but all desired PAR options are passed to the Xilinx tools via the command line.

Future implementations of the GCT firmware production system will include automatic VHDL code generation to simplify the routing of data through the system. In addition, the option to produce fully routed module libraries will be included. This may reduce the CPU time required to place and route 190 different FPGA designs.

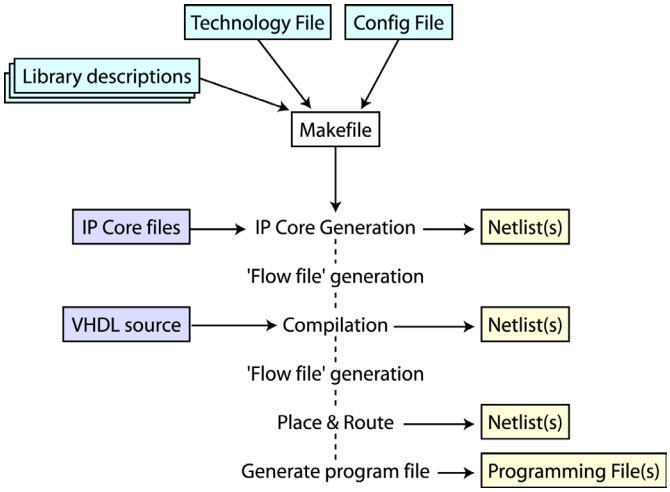

Figure 7: Flow diagram of the firmware production process.

V. CONCLUSIONS

The GCT is capable of performing a wide variety of trigger data processing tasks at relatively low cost by using a single configurable processor module. This module uses FPGA technology and Gbit/s serial links with wide data paths to maximise flexibility. The addition of a high-speed cable backplane allows the GCT to implement the CMS Level-1 jet finding algorithm. The firmware required to perform this processing is complex, with a very large number of individual FPGA programs. In order to produce this firmware in a coherent manner, the designs have been modularised, and special software tools have been written to automate the production process.

VI. REFERENCES

[1] The CMS collaboration, *CMS Trigger TDR*, CERN/LHCC 2000-38, 2000

[2] W. Smith, contribution to these proceedings

[3] D. Newbold et. al., *An FPGA-based Implementation of the CMS Global Calorimeter Trigger*, Proceedings of the 6th Workshop on LHC Electronics, 2000