

8 Calorimeter trigger implementation

8.1 Introduction

In this chapter we explore a number of areas in more detail, which would have obscured the flow of the design description given in Chapter 6. We begin by summarizing the major external interfaces of the calorimeter trigger. We then give a detailed breakdown of the latency of the proposed hardware and algorithms.

A section on technologies follows, in which we discuss some particular implementation issues in more technical detail than in Chapter 6. Wherever possible, we also describe alternative solutions in order to show that the design for the trigger system is not critically dependent on unique technologies, and that we have not arrived at the proposed solutions without investigating other possibilities.

A proposed layout of crates and racks in USA15 is presented. We then discuss the (mainly online) software and computing needs of the trigger system. Some ideas on setting up the timing, a crucial issue at the LHC, follow. We then discuss our proposed procedures for designing, building, and installing the trigger, including review procedures to assure its quality. We end with the time schedule needed to deliver the trigger to ATLAS in time for LHC start-up.

8.2 Interfaces

8.2.1 Central Trigger Processor

The Cluster Processor and Jet/Energy-sum Processor send trigger multiplicity information to the Central Trigger Processor (CTP) in real time, using a differential cable as short as possible carrying ECL or LVDS signals. For each bunch-crossing a total of 84 signal bits are sent. These consist of a three-bit multiplicity for each threshold combination for 'local' trigger objects, and one bit per threshold for global E_T sums. This gives a total of 48 bits from the Cluster Processor, and 36 bits from the Jet/Energy-sum Processor. This is summarized in Table 8-1.

Table 8-1 Summary of trigger results sent to the CTP.

| Trigger | No. thresholds | No. bits to CTP |
|-----------------|----------------|-----------------|
| Electron/photon | 8 | 24 |
| Hadron/tau | 8 | 24 |
| Jet | 8 | 24 |
| Missing E_T | 8 | 8 |
| Total E_T | 4 | 4 |
| Total | | 84 |

8.2.2 Level-2 trigger

The Cluster Processor and the Jet/Energy-sum Processor transfer RoI information to level-2 using readout driver (ROD) modules. In both these subsystems, ROD modules are responsible for the collection of RoIs from individual processor modules, performing zero-suppression and then building data packets with extra information including crate ID, module ID, bunch-crossing number and event number. Once the data are formatted, they are transferred to the level-2 trigger using S-links. The Cluster Processor will require four ROD modules and the Jet/Energy-sum Processor will require one similar ROD module to do this. These are shown in Figure 8-1. As discussed in Section 6.3.10.3, the same ROD modules can be used for both RoI and DAQ data from the Cluster and Jet/Energy-sum Processors.

8.2.3 Data acquisition system

The interface to the DAQ system is via the RODs to the readout buffers (ROBs), connected by S-links. All three subsystems of the calorimeter trigger require this connectivity. Because the data rates are normally low, the Cluster Processor and the Jet/Energy-sum Processor can use RODs which are similar (and if possible identical, see Section 6.3.10) to those used for the collection of RoI information. The higher data rates from the Preprocessor require the use of a different crate readout protocol, and this will use the different ROD design described in Section 6.2.5.3. The arrangement of RODs is shown in Figure 8-1.

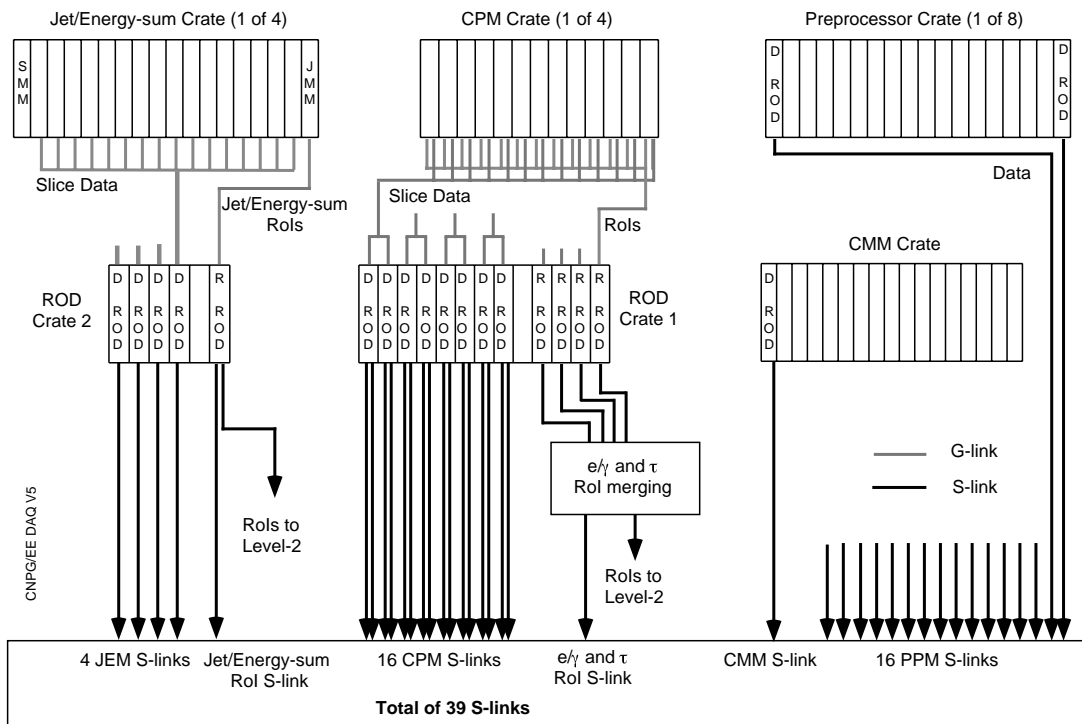


Figure 8-1 Diagram of RODs for data acquisition and Rols.

For all triggered events, the minimum information that must be read out comprises the E_T values from the lookup tables after BCID has been performed, the result bits that are sent to the CTP, and a copy of the Rols — for one bunch-crossing only. In order to monitor, check and

calibrate the calorimeter trigger, considerably more data will be read out for a prescaled or selected subsample of events. Raw trigger data from the calorimeters, extending over several bunch-crossings, is needed for cross-calibration with the calorimeter readout and to verify the BCID and timing. Intermediate results from the Cluster and Jet/Energy-sum Processors, again with the possibility of extending over several bunch-crossings, are needed in order to localize problems and verify the correct operation of the trigger.

A separate network connection is required to carry command and control data, including the threshold settings, lookup table contents, and timing settings required for different modules. It is foreseen that all modules will be provided with a VME interface through which the trigger controls will operate.

8.2.4 Detector Control System

All crates within the trigger system will be connected to a Control Area Network (CAN) bus. This will allow the Detector Control System (DCS) to check for a.c. failure, voltages or currents out of range, power supply overheating, and fan failure. In addition, as some modules will be dissipating 100–150 W, additional temperature monitoring points will be provided on each module. A local CAN-bus will join module nodes and the power supply node for each crate, and a bridge will connect each crate to the master CAN-bus and the DCS. The required logic will reside on the Clock and Control Modules.

8.3 Latency estimate

The latency is calculated in units of 25 ns LHC bunch-crossing cycles (BC), reflecting the pipelined nature of the calorimeter trigger. A full timing calculation starts at the instant the bunches collide, and therefore includes the particle time-of-flight, cables from detector to calorimeter front-end electronics, and the signal propagation time within the front-end electronics. Up to this point the delays, summarized in Table 8-2, are common to both trigger and readout paths.

Table 8-2 Delays in common for trigger and readout data (LAr shown).

| | |
|--|---------------|
| Time-of-flight to endcap at $\eta = 2$ | 0.6 BC |
| Cable to pulse shaper | 1.2 BC |
| Pulse shaper and preamplifier | 0.4 BC |
| Total | 2.2 BC |

The trigger must wait for the latest calorimeter signal before processing can begin. For this reason, the timing refers to the liquid argon endcap at $\eta = 2$. An estimated ≤ 60 m of twisted-pair cable is required to transport signals from the calorimeter at this point to the trigger Preprocessor Modules in the underground trigger counting room USA15. Most of the cables are substantially shorter than 60 m, so finding a way to reduce the longest ones, which come from the endcaps, can help to reduce latency. Including analogue summing on the detector and the required cable lengths, the peak of a calorimeter pulse requires 18.4 BC to reach the PPM, as shown in Table 8-3.

Table 8-3 Latency in analogue signal-handling.

| | |
|----------------------------------|----------------|
| Pulse peaking time (LAr) | 2 BC |
| Cable to tower-summation board | 0.2 BC |
| Analogue summation | 0.4 BC |
| Cable to USA15 (60 m, 5.28 ns/m) | 12.7 BC |
| Receiver station | 1 BC |
| Cable via patch panel and PPM | 2.1 BC |
| Total for analogue signal | 18.4 BC |

The Preprocessor delays all signals except the slowest such that they are all in time, and performs digitization, bunch-crossing identification, and G-link transmission. Signals destined for the Cluster Processor are bunch-crossing multiplexed, while those for the Jet/Energy-sum Processor are summed into 0.2×0.2 jet elements in each calorimeter. The times required in the Preprocessor are given below in Table 8-4. The bunch-crossing multiplex scheme for the Cluster Processor requires a demultiplexing delay of 1 BC at the Cluster Processor. This delay is included in Table 8-5. The most obvious way to improve the latency in the Preprocessor is to use a faster ADC, but smaller savings (e.g. by putting the lookup table in parallel with the BCID) are also possible.

Table 8-4 Latency in the Preprocessor.

| | |
|--|--------------|
| 10-bit ADC (Burr Brown ADS823) | 5 BC |
| Latch data from FADC | 1 BC |
| FIFO at min. latency setting (bypass) | 0 BC |
| Bunch-crossing identification | 7 BC |
| Lookup table | 1 BC |
| MUX odd/even BC's for CP | 0 BC |
| G-link transmission to CP | 1 BC |
| Preprocessor CP total | 15 BC |
| Sum elements for Jet/Energy-sum | 2 BC |
| G-link transmission to Jet/Energy-sum | 1 BC |
| Preprocessor Jet/Energy-sum total | 17 BC |

The electron/photon and hadron/tau algorithms are processed in the same time. This includes reception of G-link signals by the Cluster Processing Modules, bunch-crossing demultiplexing to 160 Mbit/s and multiplexing for the cluster-finding algorithm, and final processing of electron/photon and hadron hits on the Cluster Merger Modules. The time required is shown in Table 8-5.

Table 8-5 Latency in the Cluster Processor.

| | |
|--------------------------------------|--------------|
| Cable from PPM to CPM (≤ 10 m) | 2 BC |
| G-link reception | 2 BC |
| DeMUX even/odd BCs | 1 BC |
| MUX to 160 MHz | 1.5 BC |
| DeMUX; e/γ /hadron processing | 3.5 BC |
| Cluster counting on CPM | 1 BC |
| Transmission to CMM | 1 BC |
| Cluster counting on CMM | 1 BC |
| Transmit to CTP | 1 BC |
| CP total | 14 BC |

Summation in depth to form the 0.2×0.2 elements needed for jet and energy-sum triggers is performed in the Jet and Energy-Sum Modules. Subsequent processing follows separate streams, first in these modules and then in separate merger and summing modules. The timing of the common processing stage is included in both latency calculations. Since the jet and energy-sum triggers determine the overall trigger latency, it is assumed that these systems can use shorter input cabling from the PPMs, as well as a similarly shortened cable to the CTP. The jet-trigger latency is summarized in Table 8-6, and the missing- E_T and total- E_T latency in Table 8-7.

Table 8-6 Latency in the jet trigger.

| | |
|---------------------------------|----------------|
| Cable from PPM to JEM (7.5 m) | 1.5 BC |
| G-link reception | 2 BC |
| Initial summation in JEM | 2.5 BC |
| Transmission to adjacent JEM | 0.5 BC |
| Jet algorithm in FPGA | 4.5 BC |
| Transmission to JMM | 1 BC |
| Jet counting in crate | 2 BC |
| Transmission to system merging | 1.5 BC |
| Jet counting in system | 1 BC |
| Level shift and transmit to CTP | 1.5 BC |
| Jet total | 18.0 BC |

Table 8-7 Latency in the missing- E_T and total- E_T triggers.

| | |
|--|----------------|
| Cable from PPM to JEM (7.5 m) | 1.5 BC |
| G-link reception | 2 BC |
| Initial summing in JEM | 2.5 BC |
| Summation to 0.4×0.2 | 0.5 BC |
| Multiplication to E_x and E_y | 1 BC |
| Board summation of E_x and E_y | 1.5 BC |
| Transmission to SMM in crate | 1 BC |
| Crate summation of E_x and E_y | 2 BC |
| Transmission to SMM 2 | 1.5 BC |
| Final summation and threshold | 3.5 BC |
| Level shift and transmit to CTP | 1.5 BC |
| Missing-E_T and sum-E_T total | 18.5 BC |

In the final stage of processing, the CTP forms the level-1 accept signal, and this is distributed over the TTC system on optical fibres and converted to electrical signals by the TTCrx chip. An overall summary of the level-1 latency can be found in Chapter 18.

8.4 Technologies

The performance requirements of the ATLAS Level-1 Calorimeter Trigger are extremely demanding in terms of advanced technologies. A great deal of work over the last few years has been involved in design studies of various techniques and components, many of which are essential to the operation of the trigger and others which could add significant improvements. In general, these studies have culminated in the design and fabrication of various items of hardware, most of which have been evaluated in a lengthy demonstrator programme. Whenever possible, the demonstrator system has been operated in the demanding environment of the ATLAS test beam at CERN, and fed with signals from prototype calorimeters, although more detailed electronic studies and measurements have taken place in the laboratory.

This section outlines some of the more interesting and challenging areas that have been studied, ranging from ASIC design, through high-density packaging and connector techniques, to exploitation of commercial high-speed link systems. The emphasis is on demonstrating that proven solutions exist in all areas to support the proposed baseline design of the trigger system. In some cases, techniques with potential advantages (performance, cost, user-friendliness) which are still being explored are also discussed, but it is important to note that the baseline design does not rely upon them.

8.4.1 Serial links

The Cluster Processor has to process 6400 trigger towers. Given the algorithmic requirement to process overlapping windows, minimizing fan-out implies maximizing processing per module. If each CPM were to receive 8-bit parallel data from 160 trigger towers it would require 1280 connections (without BC-multiplexing) and massive cable plant. This would clearly be impractical, so it is proposed to transport the data serially. Serialization at 320 Mbit/s would require one link per trigger tower, but by using gigabit chip sets such as HP G-link (HDMP 1012/1014) up to four trigger towers could share one link. The Jet/Energy-sum Processor situation is very similar, though not quite as severe.

Gigabit link operation requires consideration of many technical issues:

- bandwidth and word-length of the protocol;
- attenuation margins (e.g. HP G-link has an attenuation margin of 9.5 dB);
- cable parameters — bandwidth, skin-effect, crosstalk, noise immunity;
- board design — transmission lines, termination techniques, noise emissions, connectors;
- robustness of the link transmission;
- power consumption;
- cost.

8.4.1.1 Hewlett-Packard G-links

As stated previously, our baseline choice is HP G-links. In Section 6.2.1.5 we showed that instead of operating the G-links at 1600 MBd to transport data from four trigger towers by time-multiplexing, it is proposed to use them at a lower speed, still achieving an effective density of four trigger towers per link by means of a bunch-crossing multiplexing (BC-multiplexing)

scheme. By utilizing the fact that no two consecutive bunch-crossings can satisfy the BCID algorithm implemented in the Preprocessor upstream of the links, the effective link bandwidth may be doubled. In addition to being more robust, the lower speed allows the use of longer coaxial-cable links, and thereby imposes looser constraints on the layout of the system in the racks. This scheme is only valid for the Cluster Processor, since the pre-summed information sent to the Jet/Energy-sum Processor does not satisfy the condition on empty bunch-crossings.

One additional bit per trigger tower will be required to identify the time slice, so 18 bits are needed to transmit four trigger towers within 50 ns. As the G-link allows up to 20 bits, the two remaining data bits may be used for additional error checking. (There is also a flag bit, which may either be used as an additional data bit, or as part of the G-link's own internal error-detection mechanism.) The total bit rate on the links is then 960 MBd, with 20-bit data plus four protocol bits every 25 ns. Data integrity via the backplane may also be checked by maintaining multiplexing and error detection over the entire data paths to the Cluster Processor ASICs.

HP G-links have been used in the trigger demonstrator system since 1995, driving coaxial cables, and optical fibres via Finisar devices. As the link lengths for the chosen architecture will be < 10 m, the simple electrical solution is adequate and has been more extensively studied. Although designed to operate as ECL devices the G-links may be operated in PECL mode, allowing interfacing to RAL163 CMOS ASICs (with built-in PECL pads) without additional conversion chips. Provision of clean power supplies, adequately filtered from the TTL supplies, is essential for reliable operation in this mode. The demonstrator system has 18 links (two per module) operating with two channels per link at 800 MBd. Nine of the links may alternatively be run with four channels per link at 1600 MBd. A demonstrator with ~20 links on one module is to be built shortly in order to test the robustness with many links close together.

Operation at 800 MBd

The links have been successfully operated at 800 MBd both in the laboratory and in the CERN test-beam environment. Using a purpose-built real-time hardware tester, high-statistics bit-error rate (BER) measurements were made. A 15-hour test run with one link showed no errors, implying a BER < 10^{-13} . To avoid increasing the ATLAS trigger rate by > 1%, a BER < 10^{-9} per channel will be required. Linkage between transmitter and receiver chips was very robust, with no losses experienced.

Operation at 1600 MBd

In this mode two 16-bit words are time-multiplexed within the 25 ns clock period, thereby achieving a density of four trigger towers per link and requiring only half the number of links and chip sets. Corresponding module real-estate and power-density requirements are also halved. Laboratory tests showed the links to be robust and error-free until VME accesses occurred in the crate housing the G-link transmitters, when link lock was frequently lost. At 1600 MBd these chips appeared very sensitive to +5 V power-supply noise, especially as the links were operated in PECL mode. (HP now offers a version of the chip set to operate with standard TTL logic.) Also, the chip sets are only specified to operate up to 1500 MBd (0°C – +85°C), although this may be extended to 1800 MBd over a reduced temperature range (0°C – +65°C). Further laboratory tests may be carried out to fully understand this problem.

The requirement to operate at 1600 MBd has now diminished since the BC-multiplexing scheme, which also packs four trigger towers into each link, has been adopted.

8.4.1.2 Alternative serial links

Alternative link technologies to G-link are under consideration, ranging from simple 320 Mbit/s links to 1 Gbit/s Ethernet or Fibre-Channel compatible chip sets.

'SimpleLink'

This is a customized serial data link designed to run at 320 Mbit/s using differential ECL [8-1]. At the transmitter end, eight-bit data words (with no error detection, BC-multiplexing, or clock encoding) are simply serialized at 320 MHz. At the receiving end the bit-stream is latched on to two latches with antiphase 160 MHz clocks to generate two 160 Mbit/s data streams. This scheme requires a clock and data alignment strategy as described in Section 8.4.2.2. For testing, 'SimpleLink' daughter cards have been designed to replace the G-link daughter cards in the demonstrator system. This scheme will not require any additional ASIC such as the Serializing ASIC used with the G-links, but additional circuitry would be needed for 'spying' on the data. This scheme uses the Harting 'harlink' connector system, as described in Section 8.4.6.1. If this method were adopted for the final system, 24 of the 'harlink' connectors (144 mm of front panel height) would be needed to input 160 trigger towers to a CPM with BC-multiplexing.

LVDS chip sets

An alternative to the above scheme would be the use of commercial low-voltage differential signalling (LVDS) chip sets, such as 'channel link' from National Semiconductor. As they are designed for the portable PC market to interface colour LCD displays, the modularity is three bytes (three colours) with four control bits, which is inconvenient for the trigger system.

Gigabit alternatives

The use of Gigabit Ethernet and Fibre-Channel chip sets, consuming less power than the current G-Links and available from many vendors, is also under consideration. In these chip sets the clock recovery schemes are different from that of G-link, where a guaranteed transition occurs in the four coding bits added to the data word for correct operation of the PLL. The Gigabit Ethernet and Fibre-Channel chip sets rely on transitions in the data field to correctly recover the clock and hence require data, which is 8B/10B block coded. Schemes are available that guarantee enough transitions for the PLL to operate correctly, but further studies are required to evaluate the operation of these chip sets at 40 MHz (clock recovery, error detection, link robustness, etc.) as was done for the G-links in the demonstrator system.

In this application, the parts would be operated at the slightly slower rate of 120 MHz, and coupled to the 40 MHz data rate using a 3:1 multiplexer/demultiplexer. The raw data capacity of these parts would be 30 bits every 25 ns.

The power supply voltage is 3.3 V, while some parts will also operate from a 5 V supply. In Table 8-8 we compare the power consumption of various receivers that are available.

Some devices are not specified to run at rates other than 125 MHz, therefore operation at the 4% slower rate of 120 MHz (three times the LHC bunch-crossing rate) needs to be evaluated.

The latency of a link using these Gigabit Ethernet (GE) parts will be similar to a link using G-link parts. The GE parts will require a 3-way multiplexer/demultiplexer from/to 40 MHz data.

There have to be sufficient transitions (12%) within the serial data stream for the receiver to extract bit and frame timing information and remain synchronized to the transmitter. A possible implementation is shown in Figure 8-2. The 120 MHz clock generator would be shared between a number of transmitters. GE deserializers have two antiphase receive clocks running at half the word rate of 60 MHz. Recognition of a frame word defines the deserialization of successive

Table 8-8 Comparison of power consumption for various high-speed serial receivers.

| Part | Power consumption | Operation at 120 MHz |
|-------------------------|-------------------|-------------------------|
| G-link: HP HDMP-1024 Rx | 2.5 W | N/A |
| TriQuint TQ9502 | 1.4 W | Yes |
| AMCC S2052 | 0.8 W | Yes |
| Texas TNETE2201 | 0.8 W | No? 125 MHz \pm 0.01% |
| Fujitsu FMM4021 | 0.6 W | No? 125 MHz \pm 2% |
| HP HDMP-1636 | 0.7 W | No? 125 MHz \pm 1% |

words. The serial link data format drawn in Figure 8-3 should allow easy translation to 40 MHz data. The control/padding field is available to implement a data encoding scheme (0/1 balance and transitions) should there be a need to use an a.c.-coupled link.




Figure 8-2 Scheme for using Gigabit Ethernet parts to transmit 40 MHz parallel data serially.




Figure 8-3 Proposed serial link format using Gigabit Ethernet parts.

Custom design

A full-custom 0.6 mm GaAs design for a low-power 1.6 GBd chip set is also being pursued. This is a design carried out by Middlesex University to a specification supplied by RAL. The transmitter chip has been manufactured and is currently under test. Power dissipation of the transmitter and the receiver chips is estimated to be around 600 mW each.

8.4.2 ASICs

8.4.2.1 Use of ASICs in the trigger system

The calorimeter trigger will utilize a number of ASIC types in several different areas. In the Preprocessor there will be two ASIC designs:

- the Preprocessor ASIC (PPrASIC), which interfaces to the FADCs and carries out BCID and LUT functions, providing trigger-tower data to the Cluster Processor, and pre-summed elements to the Jet/Energy-sum Processor;
- the Readout Merger ASIC (RemASIC), which provides the readout merger function on the Preprocessor Modules to merge readout data from the PPrASICs and transfer these data via the custom 'PipelineBus' to the readout driver (ROD) module.

More details on these ASICs can be found in Sections 6.2.3 and 6.2.5.1, respectively.

The Cluster Processor will use two ASIC designs:

- the Cluster Processor ASIC (CPASIC), which implements the electron/photon and hadron/tau trigger algorithms and contains logic to generate RoI information;
- the Serializing ASIC (SASIC), which interfaces to the HP G-link receiver and converts the 40 Mbyte/s parallel data into 160 Mbit/s serial bit-streams for backplane transmission and data input to the Cluster Processor ASIC.

More details on these ASICs can be found in Sections 6.3.7 and 6.3.6, respectively.

8.4.2.2 ASIC design experience

Over the last few years several ASICs have been designed and fabricated by Heidelberg and RAL to evaluate complex digital techniques, and also to gain experience of ASIC design and manufacture.

Using VHDL and synthesis design techniques, the Heidelberg group has successfully designed several ASICs for the demonstrator system using a 0.7 micron CMOS process from Atmel (formerly ES2).

BCID ASIC

This design was done to evaluate an ASIC implementation of one of the BCID algorithms, which had been originally implemented on Xilinx FPGAs. It was tested as part of the demonstrator programme and worked as planned.

FeASIC

This is a prototype of the PPrASIC. It includes the calibration lookup table and the BCID function in the real-time data path, and components of the readout function. The implementation required approximately 80,000 transistors (~20,000 gates) on a 16.4 mm² die.

The two preceding ASICs are described further in Section 7.4.4

RemASIC

This ASIC incorporates the functionality required for the final RemASIC, including a data compression unit, and it interfaces the FeASICs to the custom PipelineBus. The first functional

tests using an HP 82000 chip tester have been performed successfully. This ASIC is fully described in Section 6.2.

Three ASICs have been designed at RAL, two for the demonstrator programme, and one as a test chip to evaluate phase-locked delays and automatic clock-alignment logic. Designs have included full custom, semi-custom cell-based and gate-array techniques, and synthesis from Hardware Description Language (HDL).

RAL114 (0.8 micron CMOS technology from Fujitsu)

This was a 1993 gate-array design, using 20,000 gates, to demonstrate the implementation of the cluster-finding algorithm using pipeline processing elements of adders and comparators.

RAL163 (0.7 micron CMOS technology from ES2)

This ASIC is a semi-custom cell-based design serving two functions:

- serialization of 40 Mbyte/s parallel data to 160 Mbit/s bit-streams;
- conversion of 160 Mbit/s bit-streams to 40 Mbyte/s parallel data (to allow the use of RAL114 cluster-finding ASICs in the demonstrator system).

RAL215 (0.7 micron CMOS technology from ES2)

The Cluster Processor will contain 6400 trigger towers which will be serialized at 160 Mbit/s on the CPMs. As these serial bit-streams have no clock recovery mechanism, an automated calibration scheme is required to align the 160 MHz clock to the incoming data and synchronize them to the 40 MHz system clock. To evaluate one possible method, the RAL215 ASIC employs phase-locked delay techniques to generate controlled delay-elements and to use them to capture correctly the serial data and synchronize them to the 40 MHz system clock. In addition, the design includes elements for testing fast multipliers for BCID FIR-filter coefficients. See Section 6.3.7.3.

8.4.3 FPGAs

The use of ASICs involves large non-recurrent costs, and the chips can only perform the task they were designed for. In an increasing number of cases, the required task can be carried out satisfactorily by field-programmable gate arrays (FPGAs).

The Jet/Energy-sum Processor performs its trigger algorithms using FPGAs, rather than ASICs. In the Cluster Processor, FPGAs will be used for functions such as cluster merging, readout control, readout zero-suppression, etc. More details can be found in Section 6.3. The Preprocessor will use FPGAs for control and configuration purposes, as well as in the ROD where they will allow flexibility and evolution of algorithms and procedures. The Heidelberg group has already used them successfully to implement fast readout in the FEM (see Section 7.6.2).

Recent advances by competing FPGA manufacturers have resulted in programmable devices with high numbers of user I/O pins and large amounts of logic and routing resources which are capable of processing data at 40 and 80 MHz clock speeds. Because they are produced in large quantities, FPGAs can be economical compared with ASICs, especially when limited quantities of components are required, since non-recurrent costs and multiple design cycles usually associated with ASICs are avoided. Features such as JTAG are also standard in most currently-available devices. FPGAs offer maximum flexibility in modifying algorithms and system functionality without the need for hardware redesign or modifications. With SRAM-based

FPGAs, different FPGA configurations can be downloaded remotely without disturbing the system. In addition to changing trigger algorithms, another useful possibility is to download diagnostic configurations designed to test every data path between different FPGAs. On the other hand, FPGAs do not produce the best possible speed and efficiency in carrying out fixed operations, and therefore the latency can be a problem.

8.4.4 Multi-chip modules

8.4.4.1 Use of MCMs in the trigger system

The use of multi-chip module packaging technology brings crucial benefits:

- package efficiency (die size/package size), giving the capability of implementing many channels per module;
- high-speed signals and interconnect routing confined to a small area, minimizing the requirement of transmission lines on the PCBs and providing good EMI performance.

The University of Heidelberg group designing the Preprocessor are investigating the use of MCM-L technology with copper substrates. The design process of the multi-layer structure is based on an industry-standard production technique for high-density PCBs known as DYCOstrate. The Heidelberg group is currently involved in a pilot project to evaluate this technology for the final Preprocessor. Further details are given in Section 6.2.4 and below.

The electron/photon and hadron/tau cluster-finding algorithms operating in the CPMs require a high degree of trigger-tower fan-out to neighbouring modules, which should be kept to a minimum by maximizing the processing in each module. Architectural studies indicate that data from 160 calorimeter trigger towers should be fed directly to each CPM from the PPMs, thus demanding compact MCM packaging solutions. The MCM for the Cluster Processor (CPMCM) is described in Section 6.3.5 and below.

8.4.4.2 MCM design experience

Studies have been carried out at RAL into the use of both MCM-L (laminated) and MCM-C (ceramic) technology for the CPMs. Ball-grid array packaging and lead-less chip-carrier packaging built-in to the substrate have been considered. For reasons of thermal management (~5 W per MCM) and cost, MCM-C technology was chosen for the MCM designed for the trigger demonstrator system.

Cluster Processor demonstrator MCM (RAL2403)

The demonstrator MCM incorporates two HP G-link dies (HDMP-1014D) to receive the trigger towers at 800 MBd and convert them to 16-bit parallel words, and two multiplexing ASICs (RAL163) which convert each 16-bit parallel word into four 160 Mbit/s bit-streams for the cluster-finding ASICs.

The complete specification, design and thermal modelling of the MCM were carried out at RAL, and the track layout was carried out by industry. For the CPMCM in the final system, the complete design including layout could be done at RAL using the Cadence MCM design tools.

It was decided to use a rework strategy in the case of faulty dies. Although the dies are tested by the manufacturers for functionality, they are not tested at full speed. However, dies passing the low-frequency tests and the process parameter tests should be 'known-good' dies, with a high level of confidence. The strategy adopted was to use thermoplastic epoxy for die-to-substrate attachment to permit removal by reheating. In addition, the lid was attached only temporarily until the testing was completed and a permanent hermetic seal made. The MCMs were manufactured and delivered to RAL in September 1997, and were tested (after resolving various manufacturing problems). As the RAL163 ASICs were designed with built-in test facilities they could be rapidly checked, but HP G-link operation could be verified only by observing the lock condition and the resultant data via the RAL163 ASICs. Communication with the RAL163 ASICs worked correctly, but the G-links did not lock reliably. After a thorough investigation the cause of the problem was revealed to be excessive noise on the 5 V supply lines, which the G-links use as the reference when operating with PECL logic levels (to enable direct interfacing to the RAL163 CMOS ASICs). At design time only the ECL version was available from HP but a TTL version now exists. With extra decoupling capacitors added to the substrate the noise levels have been reduced and the two channels lock successfully, but further tests are needed to eliminate some remaining problems, such as loss of lock during DAQ transactions.

Preprocessor demonstrator MCM

The purpose of the demonstrator MCM is to establish MCM design techniques and experience for the final Preprocessor. Crucial design issues such as thermal management, electromagnetic interference, bonding techniques, mixing of high-speed digital and analogue signals, testing and production techniques will be addressed. The demonstrator MCM has similar partitioning into dies as the final PPrMCM. DYCOstrate four-layer technology will be used as the MCM design technique, as already described for the final PPrMCM in Section 6.2.4.

The MCM processes four trigger channels using a prototype Preprocessor ASIC (FeASIC) developed at the ASIC laboratory of the University of Heidelberg. Digitization is done by a dual-channel 8-bit FADC (AD9058) from Analog Devices, and two HP G-links operating at 800 MBd are used to serialize the output data. A Flip-chip Interconnection ASIC (FINCO) interfaces the processing FeASIC to the G-link dies. The FINCO was submitted for manufacture in April 1998 using a 0.8 μm BiCMOS process from AMS (Austria Micro Systems). It provides level conversion from TTL to PECL, since the available G-link device has no TTL inputs.

The FINCO die also contains a multiplexer, to double the operating rate of the G-links to 1.6 GBd. This baud rate is inside the typical operating range (≤ 1.8 GBd), but a bit-error rate less than 10^{-14} is only guaranteed up to 1.5 GBd. An increase in the MCM's temperature has a dramatic influence on the bit-error rate. This makes temperature measurement an important issue in the MCM design. The FINCO includes a thermal sensor to measure the temperature inside the MCM, on its own silicon substrate. The high pin-count of the FINCO results in a very 'pad-limited' design. This, and the possibility of using vias inside a pad in the MCM design, suggest the use of the flip-chip soldering technique for the FINCO. Otherwise, wire bonds would increase the total MCM area needed for this chip. The final PPrMCM would benefit from this bonding technique if a further reduction of the MCM size is needed. In that case, this demonstrator chip can show the reliability level of flip-chip mounting. In-circuit (JTAG) testing is implemented for the FINCO. A JTAG interface provides boundary-scan I/O, to preload defined pad stages and to scan the values received from the FeASIC. This is a very useful feature for testing wire connections between MCM dies.

8.4.5 Backplanes

In order to minimize the high degree of fan-out required by the Cluster Processor algorithms, each CPM should process as many trigger towers as possible. Crate-to-crate fan-out is eliminated and timing problems are simplified by duplicating trigger towers shared between CPM crates upstream, in the Preprocessor. Trigger towers shared between modules within a crate are transmitted via a backplane.

When designing the backplane, the following technical issues need to be addressed:

- maximum ASIC I/O speed — currently 160 Mbit/s with available technology;
- connector pin-count limitations for a 9U module;
- module insertion and removal forces;
- signal density on the backplane — number of layers required;
- maximum signal propagation distance — need for transmission lines;
- track impedance — board thickness, aspect ratio (manufacturing limitations);
- driver technology (e.g. ECL, LVTTTL, GTL, etc.);
- crosstalk — permissible bit-error rates;
- timing margins for data capture (the clock is not encoded with the data).

The demonstrator system (see Section 7.5.2) was designed to evaluate all critical technologies before embarking upon the final backplane design. It included data transfer from the Preprocessor at 800 MBd using HP G-links, conversion to 160 Mbit/s data using RAL163 ASICs, single-ended data transport at 160 Mbit/s via a backplane, and data reconversion by RAL163 ASICs to 40 Mbyte/s parallel data for input to the RAL114 cluster-finding ASICs.

The backplane was designed to demonstrate that low crosstalk is achievable with single-ended 160 Mbit/s data transport between modules up to ten slots (five double-width modules in the demonstrator) apart. The specifications were drawn up at RAL, with the layout, manufacture, and assembly carried out by industry.

The design features of this prototype were:

- height 3U, impedance 33 Ω , stripline, four signal layers + eight power/ground layers, grounded guard tracks between signal lines to minimize crosstalk;
- eight-way fan-out of 160 Mbit/s data to the backplane via ECL line drivers;
- received data registered to provide adequate timing margins for differential delays.

Measurements of the performance showed timing margins of ~ 3.5 ns. Signal crosstalk was below noise margins (200 mV maximum) with four neighbouring signals switching. Hardware real-time bit-error tests give a BER $< 6 \times 10^{-14}$, statistics-limited. Note that we estimate an overall BER of 10^{-9} per channel might increase the trigger rate by $\sim 1\%$.

For the final Cluster Processor, the proposed ϕ -quadrant architecture backplane will be much simpler than that described in the *ATLAS Technical Proposal* [8-2] and above, as the maximum signal propagation distance will be only one slot. With communication only between adjacent modules (no module-crossing), significantly fewer backplane layers will be required, thereby possibly avoiding the need for low-impedance transmission lines to keep the backplane thickness within manufacturing capabilities (aspect ratio). I/O buffers, such as 3.3V TTL,

CMOS or GTL (if necessary for driving transmission lines), could be implemented on the ASICs to drive and receive the backplane signals. The baseline solution will be to use single-ended ECL signalling, as already demonstrated. Table 8-9 compares our previous and present proposals.

Table 8-9 Comparison of backplane features.

| | ATLAS Technical Proposal | Present ϕ -quadrant design |
|---------------------------------------|--------------------------|---------------------------------|
| Number of signals | 456 | 320 |
| Signal propagation distance | 5 slots | 1 slot |
| Transmission line | Yes | Yes |
| Low-impedance (25Ω) drivers | Yes | No |
| Number of layers | 12-16 | 4 + power/ground |

The Jet/Energy-sum Processor will use a similar architecture and backplane. However, the smaller number of signals being shared between modules makes a slower, 80 MHz data rate possible.

8.4.6 Connectors and cables

The complete trigger system has to handle several different types of signals:

- 960 MBd signals on 50Ω coaxial cables from the Preprocessor — 40 inputs per CPM, 22 inputs per JEM;
- 40 Mbyte/s signals on twisted-pair cables to convey results to the CTP;
- 160 Mbit/s single-ended signals via the backplane — ~ 300 per CPM;
- 80 Mbit/s single-ended signals via the backplane — ~ 330 per JEM.

8.4.6.1 Cable connectors

To ensure easy module removal, the ~ 40 input cables, for example, should use a compact modular connector. Several potential candidates have been identified.

Coaxial cable connectors

A high-density coaxial-cable connector is available from Harting — the 'harpak' eight-contact, 50Ω mini-coaxial module, 30 mm high. Each CPM would require five such modules, using ~ 150 mm of vertical panel space. Used with CERN-approved coaxial cable type C-50-2-1 the system is capable of transmitting 2.2 GHz signals over about 6 m [8-3].

Twin-ax cable connectors

This connector system from GORE is based on shielded twin-ax cables fitted to five-pin, 2-mm-pitch wafers. Each wafer carries two 1.2 GBd signals, and is stackable on a standard 2 mm five-row connector. Compliance with CERN fire-safety standards has still to be determined.

Mini-coaxial backplane connectors

Using the BC-multiplexing scheme means that fewer pins will be required on the CPM backplane connector for the fanned-out signals. It may therefore be possible to input the signals from the Preprocessor via backplane mini-coaxial connectors, such as the METRAL™ backplane mini-coaxial connector system which is compatible and stackable with other 2 mm METRAL™ connectors. These may also be used for the 160 Mbit/s backplane signals (see below) as well as for module power. Nine mini-coaxial positions are possible in a 12 mm connector module requiring only 60 mm height of the backplane edge space.

High-density connectors for twisted-pair cables

The Harting 'harlink' 10-contact (five pairs) modular metric I/O connector system is designed for high-speed data transport in a compact package. The connectors can be PCB-mounted on a 6 mm pitch, thereby allowing 160 contacts in 100 mm of vertical panel space. They are currently being tested carrying 320 Mbit/s signals as part of the present demonstrator system.

8.4.6.2 Backplane connectors

Without the BC-multiplexing scheme, high-density backplane connectors would be needed in the Cluster Processor to transmit and receive 456 single-ended 160 Mbit/s signals. However, by using this scheme the number drops to around 320. In addition, a number of signal ground-pins and pins for other functions (e.g. slow controls) will be required. Three possible connector systems are under consideration for the trigger system.

Futurebus+ type, 2 mm METRAL™ connectors

Four-row, 2 mm METRAL™ connectors were used for the trigger demonstrator backplane. For the final Cluster Processor, a five-row connector, with the middle row used for signal grounds, could be used. Assuming about 400 signal pins (320 for 160 Mbit/s signals and 100 for other signals), 9 SU (1 SU = 25 mm) connectors would be required. The insertion forces involved will be approximately 240 N per connector, placing significant mechanical constraints on the backplane design and requiring a suitable method of module insertion and extraction (e.g. via screws or levers).

Siemens DensiPac

DensiPac is a high-density surface-mount backplane connector with a grid of 1.25 mm, giving 576 pins per 100 mm card edge with insertion forces of approximately 180 N. Designed for automatic surface-mount process lines, it can be assembled with standard pick-and-place machines, needing no special tools to press fit and offering easy in-service replacement.

AMP Z-PACK 2 mm connectors

Similar to the METRAL™ connectors mentioned above, these are 2 mm pitch, five-row connectors meeting the DIN 43356 and IEEE 1301 hard-metric equipment practice. The insertion force is about 65 N per 5 cm of card edge. These connectors are used in 6U CompactPCI backplanes, for which industrial insertion/extraction hardware has been developed. A seven-row option, with the two outer rows used as a ground screen, is also available.

Siemens SpeedPac

This is a high-speed backplane connector supporting data-rates up to 2.5 Gbit/s and rise-times down to 50 ps, with crosstalk < 0.4%. It is designed as zero-insertion-force (ZIF) for differential or single-ended lines. Instead of using the conventional male and female contacts, with

through-hole connections to the PCBs, SpeedPac is designed as a Beam-On-Pad connector. It addresses the two major issues concerned with backplane connectors:

- insertion force — a front panel lever opens the connector for insertion and extraction;
- signal/ground-pin ratio — all contacts can be used for signals, with grounding via the surrounding connector frame.

The differential-pair connectors are now in production, and the single-ended version will be available in early 1999. Each CPM would need two 4 SU double-sided connectors for signals, and a separate power connector. The 25 mm width of the double-sided connector would unfortunately allow only 15 crate slots, compared to 21 in a conventional crate, but with only 13 CPMs per crate in the Cluster Processor this may be acceptable.

8.4.7 Programmable delay elements

Unlike the 800 MBd serial links, where the clock is encoded with the data and recovered at the receiving end, the 160 Mbit/s data transmission used in the Cluster Processor does not have a clock recovery scheme. Therefore the function of the clock alignment logic is to select the appropriate 160 MHz clock phase to capture the incoming serial data at 160 Mbit/s using 1.25 ns delay elements, and then to synchronize to the 25 ns clock period using 6.25 ns delay elements. For example (see Figure 8-4) if A5 (hex) is sent as a calibration data pattern the receiver could see four possible patterns: A5, 4B, 96, and 2D. The first phase of the calibration logic is to use a histogramming process to determine which phase setting (1.25 ns taps) produces the highest capture frequency for one particular data pattern. The second phase synchronizes the pattern correctly by using the 6.25 ns delays. For example, if A5 were received no delay would be required, if 4B were received one 6.25 ns delay would be required, and so on up to three 6.25 ns delays.




Figure 8-4 Calibration sequence pattern.

Performing this procedure by software in the demonstrator system took several minutes to calibrate 36 channels, which would clearly not be feasible for the full trigger system. A hardware implementation in the Cluster Processor ASIC would take only a few microseconds to carry out the calibration process for all channels in parallel by activating the calibration sequence. An ASIC (RAL215) to evaluate this technique has been designed and successfully tested.

The Preprocessor will use phase-locked delay chips developed at CERN. If the RAL215 ASIC tests show that there is any risk involved in implementing the phase-locked delays inside the CPASICs, then these phase-locked delay chips could be used externally. It would also be necessary to use external delay elements in the case of the CPASIC being designed as a gate array (e.g. for commercial reasons), as the phase-locked technique is a full-custom process.

8.5 Rack and crate layout of the calorimeter trigger

The rack and crate layout has not yet been finalized. In this section we present the considerations governing the layout, and show one possible arrangement which meets these criteria.

The main constraints on the layout are the total length of all cables including the analogue inputs and results to the CTP, in order to minimize latency, the maximum length of high-speed digital cable runs, and also the power consumption and cooling capacity in each rack. Within such limits, we would clearly prefer layouts which are conceptually simple, ergonomically favourable, and easy to maintain. The preferred layout should also be able to cope with architectural changes such as a move to lower density PPM or CPM crates.

The signal-cable runs within the trigger comprise those from the detectors to the Receiver Stations; from the Receiver Stations to the Preprocessor; from the Preprocessor to the Cluster and Jet/Energy-sum Processors; and finally from the two trigger processors to the CTP, level-2 and the DAQ. The layout of the receiver station racks are primarily a concern of the LAr groups, but clearly they must be close to the trigger racks. The cables from the receiver stations will be passing analogue signals, whereas those from the PPMs to CPMs and JEMs will carry digital signals. The latter should not be longer than 10 m (Section 7.5.1), whereas the analogue cables have no such restrictions apart from considerations of latency.

All three main processor modules are expected to consume considerable power, with each CPM for example, dissipating about 150 W. Typically each crate will draw about 2 kW. We therefore propose to have no more than two such 9U crates per rack. Even so, careful consideration is required when specifying the fans and designing the air flow through the crates. Spare space in the racks may be used for other crates such as for DAQ, DCS connection, etc.


This limit of two large crates per rack is also favoured on ergonomic grounds when working on the system during the installation and commissioning phase, and subsequently for the ease and speed of maintenance should it be necessary to change faulty modules or crates.

One of the design goals of the architecture and layout is to have as conceptually simple a mapping as possible from trigger granularity to that of modules and crates. It is hoped that this will minimize errors and ease the learning curve for new people. Given the overall ϕ -quadrant architecture and a limit of two crates per rack, a layout with the four quadrants of each processor arranged in two adjacent racks is attractive.


Two alternatives have been considered for the Preprocessor. It could be located in the same row of racks as the receiver stations, as illustrated in Figure 6-9. In this case the analogue cable runs are short and the complexity of the PPM to CPM/JEM cabling, including the intercrate fan-outs, could be taken care of in the cable runs under the floor and/or above the ceiling. However this solution requires fairly long digital cable runs. The currently preferred option is to place the Preprocessor crates between the two trigger processors, bringing the analogue cables from the adjacent row. In this case the digital cables can all be less than 5 m long, but must all be accommodated within the same row of racks. This scheme is illustrated in Figure 8-5.

If it proves necessary to reduce the density of channels per PPM or CPM and consequently divide the system into octants, the cluster processor, for example, would double to fill four racks. One possible advantage of such an arrangement is that it might allow all connections to be made at the rear of trigger modules, allowing their removal from crates without the need to disconnect any cables.

(a) Adjacent rows of racks (viewed from above)



(b) Trigger processor racks (viewed from front)



(c) Cabling from Preprocessor to Cluster Processor

Cabling from Preprocessor to the Jet/Energy-sum Processor is similar, but has only one quarter of the channels and has been omitted for clarity.




Figure 8-5 Possible layout of calorimeter trigger racks.

8.6 Software

8.6.1 Future development of test and diagnostic software

When prototype modules of the final trigger system appear, comprehensive diagnostic software will be needed in order to test them. This software will be the first version of the final diagnostic software, i.e. that used to test the production modules. Trigger experts will use the diagnostic software to trace faults during the lifetime of the ATLAS detector, but the bulk of the functionality must be implemented within the next two years. Experience has also shown that software models of the hardware may highlight problems in the hardware design. The scale of the software effort required is comparable to that needed for the design of the hardware.

The main requirements are similar to the functionality of the existing diagnostics software, described in Section 7.8.2. The engineers and physicists using the software should be presented with a graphical view of the system configuration. It should be possible to select which module(s) to investigate in detail. For each module it must be possible:

- to display each register and memory cell;
- to set and display every bit;
- for the significance of each register bit to be separately explained and reflected in graphical user-interface 'widgets';
- to exhaustively exercise each register and memory cell, highlighting errors;
- for the interaction between various registers to be modelled;
- for incorrect behaviour to be flagged.

Online help in using the software must be provided as well as paper documentation.

In addition to the detailed diagnostics of single modules, it must be possible to test the communication between pairs and chains of modules, and eventually across the whole trigger system. For higher-level diagnostics it will probably be too slow and cumbersome to use the detailed software model of each module that is necessary for the single-board tests. We will require a simpler functional model of each module, one that just maps its input data to its outputs. Since it will be desirable to use the same software framework for both individual module and larger scale diagnostics, the ability to enable or disable the detailed software model dynamically on a module-by-module basis will be required. Flexibility is also essential, since experience has shown that it is not always possible to predict precisely what tests will be needed, or how the software will be used in practice to solve problems with the hardware.

In the existing diagnostics software, the system configuration is kept as a text file which can be updated graphically by the user interface. In future we will want to use the OO database of the ATLAS DAQ system when this becomes available. This should include both the parameters of every module and the interconnections between them.

The diagnostics software is required to be user-friendly and interactive. We will also need some 'batch' versions. These should be packages of tests giving a simple pass or fail result. Such packages can be run by the 'Test Manager' component of the proposed ATLAS back-end DAQ framework [8-4]. In order to allow both humans and a possible (computer-based) expert system to identify faults, test packages should be available on scales ranging from the single module to crates, chains of connected modules and the whole trigger system. Such test packages will

clearly share most of their code with the interactive diagnostic software. There may be also some overlap with the online calibration software described below.

8.6.2 Future development of calibration software

Some calibrations, such as timing, can best be performed online. Others, such as the calorimeter energy calibration, will probably be done online at first and then refined by offline analysis.

The PPMs contain lookup tables which will provide final E_T scaling. We will want to compare, offline, the trigger-tower energies we receive from the tower builders with the full calorimeter readout of the individual cells. Monitoring of noise will be required to set suitable thresholds in the lookup tables. Apart from the energy calibration, there will be other parameters to be calibrated for many modules, such as on-board clock delays, FADC pedestals, etc.

The trigger is absolutely dependent on correct bunch-crossing identification. The proposed BCID algorithm uses a filter matched to the pulse shape. We will have to collect pulse-shape data and tune the BCID parameters accordingly. This may be done online or offline. We will also need to use the calorimeter calibration system to tune the BCID behaviour for saturated pulses.

8.6.3 Online control and monitoring of the calorimeter trigger

The ATLAS DAQ system will allow us to spy on the data read out from the calorimeter trigger. We will look at real data and monitor triggers to verify its correct operation. This will require an online simulation of the processor using a functional model of each of our modules and of the tower builders. This could be the same as that used in test and calibration packages. The online monitor should also present summary information such as overall trigger rates and maps of activity across the detectors. It should raise alarms if required.

In addition to high-level monitoring via the DAQ system, we will also want to perform local monitoring within the trigger system. This may include spying on the DAQ data flow and also using information which is not normally read out.

We will require software to download the whole (or partial) trigger configuration. This includes thresholds, delays, FPGA programs and so on. We will also want to check periodically that the correct configuration is still downloaded, i.e. check for corruption of module registers. It must be possible to load the trigger locally or under ATLAS central control. Reloading of the trigger may have to be done in conjunction with the CTP and Muon Trigger.

We will also want software to allow us to create the configuration files. For example, to change thresholds globally or across a large fraction of the detector it should be easy and intuitive for a physicist to set up new trigger conditions. Details of the configurations will be kept in the DAQ database, which should provide the necessary controls on who can change it and the logging of changes made.

We will use an offline trigger processor simulation (from DAQ data) to make detailed checks of correct operation of the processor. This might be the same as (or a more detailed version of) that used for online monitoring of the system. This would complement the existing simulation that treats the processor at an abstract level.

Many of the calibration tasks described above will be run offline using the reconstructed data from the trigger, in particular the energy calibration and probably the determination of BCID parameters for every channel. We will also want to check offline, using longer runs, those calibrations which would be initially performed online.

8.7 Strategy for setting up the timing

In this section we discuss the procedures foreseen for setting up the timing of the calorimeter trigger. This will be a complex process, and cannot be done completely in isolation from the rest of ATLAS. It is therefore likely to be a long and iterative procedure. We can compare with experience from HERA, where the detectors are smaller and the bunch crossing interval is 96 ns. In H1, for example, only near the end of the first year of data-taking were all triggers and all subdetector readout branches synchronized for all events. In ATLAS we will have several bunch-crossings in the detector at the same time, although we will have better BCID than H1.

The timing problem and its solution can be broken down into different areas. Firstly we need to ensure that the trigger is correctly timed; secondly we must read out the right set of bunch-crossings to level-2 and the DAQ when the level-1 accept signal is received.

The procedure for timing in the trigger consists of several logical steps: synchronizing all the trigger tower signals from the calorimeters on the PPMs; synchronizing FADC strobes and BCID with the analogue pulses (see Section 5.3.3); internal timing of trigger modules and the connections between them; adjustments for the different latencies of the various calorimeter triggers; synchronization of calorimeter, muon and other triggers at the CTP.

The facilities we can use, in order of their likely availability, will be: internal test signals within the trigger hardware; calorimeter test pulses; cosmics; particles from one beam; and finally colliding beams.

Since the internal timing can be done stand-alone, this is likely to be the first part of the procedure to be completed. It involves both the timing of signals and clocks on individual boards and of the signals transmitted between boards, either via cables or the backplane. We need to set both the phase of the signals with respect to the LHC clock, and also their relative latency in numbers of whole clock cycles.

In the context of our demonstrator programme (see Section 7.8.3) we designed two different ASICs, RAL163 [8-5] and RAL215 (see Section 8.4.2.2), for evaluating data transport and re-synchronization at 160 Mbit/s. We also developed software for calibrating the timing of the backplane connections between CPMs, by loading known data into memories on each board and capturing the data on the receiving module. A similar procedure can be used to set the latencies between the chains of modules in the trigger: PPMs to CPMs and JEMs, and CPMs and JEMs to their respective merger modules.

8.7.1 Synchronization of input signals

Both the LAr and tile calorimeters have pulser systems for calibrating their electronics. They can thus generate trigger-tower signals for the PPMs. These can be used to set the FADC strobe phase with respect to the LHC clock and the pipeline delay for each channel. According to their respective TDRs [8-6][8-7], the pulser system for each calorimeter should provide signals with

the same relative timing as those expected from beam particles, at least to within 1–2 ns. This will of course have to be checked when real particles are available.

Software to perform the timing calibration of PPM input signals will have to access both the trigger hardware and calorimeter controls. It may be written as part of the DAQ or as a stand-alone calibration program.

With correctly timed inputs, the BCID parameters can be tuned. Then, with suitable choices of trigger thresholds adapted to calorimeter calibration pulses, the timing of triggers from a range of pulses from small to saturated can be checked through the trigger system.

8.7.2 Use of real particles

Before the LHC start-up, it is proposed to run with cosmic rays. However the timing of cosmics is very different from that of beam particles, so some special timing regime may be required for this.

During the initial LHC machine studies, there will probably be tests with a single bunch (or a few bunches) in one beam, and the beam backgrounds will be quite high. In this situation the calorimeter may be well illuminated with beam halo or beam–gas interactions. However, like cosmics, these will not have the same timing as particles from pp interactions. However it may be possible to use them for some gross timing checks, such as looking for any evidence of reflections on cables. Since the total cable lengths from detector to PPMs are of the order of 60 m, a gap of 24 bunch-crossings between filled LHC bunches would be required to avoid any possibility of confusion.

The next stage is likely to be a small number of colliding bunches. This will finally allow the true timing to be checked against that from the pulser systems. Beam backgrounds will probably still be high, but the luminosity will be low so that it may take some time to accumulate adequate statistics from all parts of the calorimeters.

8.7.3 Timing of the readout

It should be reasonably easy to read out the correct bunch-crossing from the calorimeter trigger in response to a level-1 accept originating in the calorimeter trigger itself. However, ensuring that the readout is correct for all triggers (e.g. muons) requires correlating the latencies of calorimeter, muon and other triggers. This task, which is the responsibility of the CTP, is likely to take some time with colliding beams and offline analysis before we can be confident that all subdetectors are reading out the same event on all types of trigger.

8.8 Design, construction and assembly procedure

Producing the Level-1 Calorimeter Trigger is the joint responsibility of six institutes in three countries. In broad terms the trigger is partitioned into three subsystems, each of which is the responsibility of one or more institutes:

- Preprocessor;
- Electron/photon and hadron/tau Cluster Processor;
- Jet/Energy-sum Processor.

There are a number of hardware items which are (or could be) common to all subsystems, and for which all subsystems should share common specifications and/or designs. These could include items such as crate assemblies, TTC receiver modules, ROD modules, as well as control and computing infrastructure, etc.

8.8.1 Design procedure

The design and manufacture of a trigger system module will in general proceed through several phases. Where there are areas of the design which rely critically on techniques or technologies not yet fully tested, it may be decided to demonstrate their use in a small-scale pre-prototype module. The purpose of this exercise would be to answer some very specific questions to inform and guide the design of the final system module.

Given the overall architecture and partitioning of the entire system, the first stage will be to prepare a detailed functional specification of the module, including a programming model and a full definition of the I/O requirements relating to all interconnecting subsystems. Detailed design of a prototype module will then commence. This will in general have full functionality, possibly with enhanced diagnostic facilities, but with a reduced channel count. A module test plan will be drawn up during this period to enable stand-alone module testing. Customized test equipment will be designed and manufactured, and software prepared. A subset of the standard TTC system will also be required as part of each overall test system.

Only a small number of these prototype modules will be manufactured, but sufficient to test all forms of intermodule signalling. At this stage, it will be particularly important to rigorously check all interface connections to the module, which in general will require the availability of other prototype modules. Module design and production will be carried out in several different institutes, and must therefore be carefully integrated into an overall schedule. However, in some cases, such as interfacing the PPMs to the calorimeter trigger-tower signals, or the RODs to level-2, it will be necessary for the test system to emulate these links.

Another very important aspect of the prototyping phase will be the evaluation and verification of the ASIC and MCM designs during testing. Provision has been made in the cost estimates and in the workplans for a second iteration for each design (if required).

The results of this extensive test programme will be fed back as possible module design modifications before the detailed design of the final system module can start. This will now be to the full specifications, but with any changes which may have emerged from the prototyping phase. A small number of these pre-production modules will then be produced, which after undergoing similar testing to the prototype modules should require very few modifications to perform at full specification. Any such small changes needed will be made to the module design

and/or layout, following which production of the full number of final system modules can start.

To summarize, the terminology used for the different stages of module design is as follows:

- **Pre-prototype** (Module '-2') refers to a module (possibly stand-alone) that is designed to test only a small subset of specific features which are crucial to the final design. Some parts of the phase-2 demonstrator system can already be considered to be in this category.
- **Prototype** (Module '-1') refers to a module that is designed to test all features and functions of the final design, but which will not necessarily have a full complement of channels. It must be fully compatible at all its interfaces with other prototype modules.
- **Pre-production** (Module '0') refers to a module that is of the final design, with the full number of channels. After test it would be expected that only a small number of minor design/layout changes would be needed before full production could start. After minor modifications, the pre-production modules themselves would probably be fully compatible with the final production modules (and could even be used as 'emergency' spares).
- **Production** refers to a module that is essentially identical to the pre-production module, but with the few minor modifications identified during pre-production tests incorporated. If the number of final modules needed were small enough, the production stage may not be required and the pre-production modules (with appropriate hardware modifications) would be used.

8.8.2 Assembly and pre-testing

It is envisaged that each of the three trigger subsystems (Preprocessor, Cluster Processor and Jet/Energy-sum Processor) will initially be assembled independently at the institutes responsible for their designs, and tested as stand-alone systems. As an example, the full Cluster Processor will be assembled as four crates of 13 CPMs feeding a crate of CMMs and a crate of RODs. A small number of PPMs (in playback mode) will be used to supply signals, and the overall test system will monitor the data intended for the CTP, level-2 and the DAQ system. The main purpose of this stage of testing will be to gain experience of system aspects (timing, control, monitoring, etc.) and to exercise and refine the online software.

8.8.3 Installation and integration

After the full stand-alone testing phase, the three calorimeter trigger subsystems will be shipped to CERN. The first part of the electronics to be installed in USA15 will be the Preprocessor, which will be connected to the calorimeter cabling and to the TTC system. A thorough verification of all calorimeter trigger-tower signals will then be carried out, and the timing will be set up (using the calorimeter test-pulse system). Links to the DCS and between the RODs and the DAQ ROBs will also be established as these become available.

Installation of the Cluster and Jet/Energy-sum Processors can then proceed. This will first involve cabling each of them to the Preprocessor via the serial links, connecting the TTC system optical fibres, and linking the RODs to the DAQ ROBs. The DCS connection will also be made. At this stage, the full trigger system will be timed-in using the established timing strategy and a systematic channel-by-channel functional verification will be performed.

The final links will be made from the CMMs, JMMs and SMMs to the CTP, and from the RODs to the level-2 trigger via the RoI Builder (which is the responsibility of the level-2 trigger). This will complete the system integration phase.

8.8.4 Quality assurance and review procedures

There will be a semi-formal technical review procedure set up for all subsystem elements, including software, covering the period from mid-1998 until final production. It will be operated jointly by the six collaborating institutes. The guiding principle is to identify a small pool of specialists from the six institutes, from which a group of four will be drawn to review the design of each subsystem module. Each review group will contain at least one person with expertise in one of four main areas — DAQ, electronics, software, and system engineering. In general, all the specialists will be actively involved in design or specification of at least one other subsystem module themselves, so will approach the review procedure from different perspectives. The composition of the group will change as different subsystem elements are reviewed to ensure objectivity, but a group reviewing a given subsystem element should remain together throughout the different design phases to ensure continuity. An added benefit of this procedure will be that in-depth knowledge of all subsystems will be spread amongst all the institutes.

The design and specification of each subsystem element will be subject to at least two review exercises — a Preliminary Design Review (PDR) and a Final Design Review (FDR) — with the possibility of an Interim Design Review (IDR) between them.

The PDR will be used to examine and assess the full requirements and specifications for the subsystem element, to identify any missing functionality, to ensure full compatibility with all connecting subsystems and to determine overall feasibility. This is perhaps the most important part of the review procedure, as it will determine the direction of subsequent engineering effort. Detailed written specifications will be supplied to the review group two weeks in advance of the review itself, and following the review the final agreed conclusions will be distributed to the level-1 calorimeter trigger community.

IDRs may be held at any time during the design phase, but most usefully at the completion of schematic capture when many engineering issues (timing margins, interfaces to other subsystems, detailed latency calculations, etc.) can be explored. By monitoring progress at this stage some potential problems may be detected and resolved early, thereby minimizing wasted effort.

The FDR will be held before the module design is sent for manufacture, and is intended to catch any design or engineering errors before budgetary commitment. This review will necessarily be of a more technical nature than the initial review, but there should be few problems to detect by this stage.

8.9 Timetable

The proposed schedule for design and construction of the Level-1 Calorimeter Trigger is given in Figure 8-6. Milestones are given in Section 23.4.

8.10 References

- 8-1 A.J. Maddox, *Simple-links*. RAL hardware write-up, January 1998.
<http://hepwww.ph.qmw.ac.uk/~efe/Simple-Links.html>
- 8-2 *ATLAS Technical Proposal*, CERN/LHCC/94-43, 1994
- 8-3 A.J. Maddox, *Evaluation summary of the Harting 'harpak' mini-coaxial module with IS23 approved cable*. RAL hardware write-up, February 1998.
<http://hepwww.ph.qmw.ac.uk/~efe/coax&harpak.html>
- 8-4 *Back-End DAQ User Requirements Document*. Draft Revision 3, 1996.
<http://atddoc.cern.ch/Atlas/DaqSoft/document/draft.ps>
- 8-5 J. Edwards and V. Perera, *User guide for the first level calorimeter trigger processor demonstrator dual function ASIC - RAL163*, RAL hardware write-up, 1994.
- 8-6 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN/LHCC/96-41, 1996.
- 8-7 *ATLAS Tile Calorimeter Technical Design Report*, CERN/LHCC/96-42, 1996.

