CERN/DRDC 94–52
RD27 Status Report
9 January 1994

FIRST-LEVEL TRIGGER SYSTEMS FOR LHC EXPERIMENTS

I. Brawn, R.E. Carney, Y. Ermoline[1], J. Garvey, D. Grant, R. Harris, P. Jovanovic, I. McGill, D. Rees, R. Staley, A. Watson, P.Watkins
School of Physics and Space Research, University of Birmingham, UK

O. Buyanov, N. Ellis *), C. Jacobs, B.G. Taylor, J.–P. Vanuxem, H. Wendler
CERN, Geneva, Switzerland

P. Hanke, M. Keller, E.-E. Kluge, M. Wunsch
Institut für Hochenergiephysik der Universität Heidelberg, Germany

P. Bödo, H. Hentzell, C. Svensson, J. Yuan
Linköping University, Sweden

J. Fent, W. Froechtenicht, C. Kiesling, H. Oberlack, P. Schacht
Max-Planck-Institut für Physik, Munich, Germany

E. Eisenhandler, M. Landon, J.M. Pentney, G. Thompson
Queen Mary and Westfield College, University of London, UK

J. Dowdell, J. Edwards, C.N.P. Gee, A. Gillman, R. Hatley, J. Leake, V. Perera
Rutherford Appleton Laboratory, UK

B. Green, J. Strong
Royal Holloway and Bedford New College, University of London, UK

E. Gennari, A. Nisati, E. Petrolo, G. Piccinelli, S. Veneziano, L. Zanello
Intituto Nazionale di Fisica Nucleare and Universita La Sapienza, Rome, Italy

R. Cardarelli, A. Di Ciaccio
Intituto Nazionale di Fisica Nucleare and Universita Tor Vergata, Rome, Italy

G. Appelquist, C. Bohm, M. Engström, S. Hellman, S-O. Holmgren, E. Johansson, N. Yamdagni, X. Zhao
Stockholm University, Sweden

W.H. Smith, S. Dasu, C. Foudas, T. Gorski, M. Jaworski, J. Lackey, D. Panescu, W. Temple
University of Wisconsin, Madison, Wisconsin, USA

(*) Spokesman

---

[1] Now at CERN, Geneva, Switzerland.

1. Introduction

The RD27 project [1] was approved in June 1992 [2] to perform a broad-based study of level-1 trigger systems. The progress of the project was reported [3] in Summer 1993 and continuation of the project for a further year was approved [4] with the following milestones:

- Make a detailed description and simulation of the central trigger processor in a high-level design language.
- Continue study to optimise input to the bit-parallel calorimeter processor.
- Make detailed design studies of a bit-serial calorimeter trigger processor and compare its expected performance to the performance of the bit-parallel processor implementation.
- Evaluate the performance of the muon trigger prototype in the RD5 test beam.

Since the last status report, two new groups have joined the collaboration: the group from University of Wisconsin led by Wesley Smith and the group from INFN/Rome-2 led by R. Cardarelli. The Wisconsin group is studying calorimeter trigger processors, while the Rome-2 group is studying muon triggers.

Members of RD27 have been very active during the last year in the preparation of the technical proposals for both ATLAS [5] and CMS [6]. In particular, the design work for the ATLAS calorimeter, muon and central trigger processors, and for the CMS calorimeter trigger processor, was performed by RD27 groups. Members of RD27 (N. Ellis and W. Smith) are the trigger coordinators for the two proposed LHC experiments.

Members of RD27 have also continued to contribute to the development of the optical Timing, Trigger and Control (TTC) distribution system [7] which, following a decision in the last DRDC review of RD27, is now funded only through RD12. Latency stability has been studied, encoder jitter and laser transmitter performance have been improved, the functionality of the timing receiver ASIC (application-specific integrated circuit) has been enhanced, and its development by CERN Microelectronics Group is now well advanced. Both ATLAS and CMS Technical Proposals have specified implementation of the RD12/RD27 TTC system, and close contacts will be maintained to ensure efficient integration of the system with the overall trigger and DAQ systems.

RD27 has been found to be a very useful forum for exchanging ideas between ATLAS and CMS, including detailed reviews of the different designs, comparison of algorithms and cross-checking of trigger-rate calculations. In some areas, for example fast, synchronous optical data-transmission to the calorimeter trigger processors, there is close collaboration within RD27 between members of the two LHC experiments. Other areas where members of ATLAS and CMS will collaborate closely are in evaluating high-speed transmission-line backplanes and high-density connectors, and comparing different ASIC technologies (full custom and gate arrays) and manufacturers. We also consider the possibility of using some ASICs (or custom cells) in the trigger processors for both the experiments.

In this document we report on the work that has been performed and propose a programme of R&D for the next 12 months. We discuss in turn the calorimeter, muon and central trigger processors, and suggest milestones for the coming year.

2. Calorimeter Trigger

Introduction

The design work for the calorimeter trigger processors of ATLAS and CMS has been performed by members of RD27. These designs are documented in detailed technical notes [8, 9, 10] and are described briefly in appendices A–C of this document. Although they represent a very substantial amount of effort by members of RD27, here we concentrate only on the technical R&D that has been done in the last year and that is proposed for the next 12 months. Another aspect of the work that is not discussed here is the physics simulation studies [8, 11] performed to evaluate different algorithms by comparing rates and efficiencies. Simulation studies have been performed with full detector simulation (GEANT); many details are considered in the simulations, such as the effects of pile-up and electronic noise after pulse shaping, and bremsstrahlung of electrons and conversions of photons.

All the design studies are for purpose-built systolic processors, programmable at the level of parameters, and implemented making extensive use of ASICs. They assume that digitisation is performed on the detector, either using a dedicated digitisation system for the trigger or using the FERMI system [12]. A large number of high-speed, synchronous optical links is used to bring the data to the trigger processors — with typically ~7000–8000 trigger ADC channels, 8-bit precision and 40 MHz bunch-crossing rate, the aggregate data rate is ~2000 Gbit/s. This bandwidth is very high compared to typical applications in the computer and telecommunications industries, and the combined requirements of low latency and synchronous data transfer are specific to the level-1 trigger application. An important part of our work therefore relates to issues of data movement.

All of our processor designs are based on data links with a usable data rate of at least 800 Mbit/s, allowing the transmission of two 10-bit values every 25 ns (e.g. E_{T} data for two trigger cells). We are evaluating both commercial products (so-called ‘chip-sets’) and custom solutions for these links; the latter may be cheaper and better suited to the special requirements of the level-1 trigger. The possibility of receiving optical fibres on Multi-Chip Modules (MCMs) is being investigated as a means to achieve high density, and is assumed in the designs for ATLAS. The issue of error detection is also being studied — a single-bit error in the most-significant bit of any of the ~7000–8000 trigger channels would lead to a fake trigger.

At the time of our last status report, we were considering using asynchronous data transmission on lower-bandwidth links. This relied on zero suppression to reduce the average data rate. This option was evaluated in detail [13] and would have been feasible. However, the need to buffer the data before transmission increased the latency and led to the possibility of data loss (in case of buffer overflow). The design was therefore abandoned in favour of one using high-speed synchronous links (the design described in the ATLAS technical proposal); the new design incorporates much of the old one for the part of the system that performs the actual trigger processing.

Another aspect of data movement that we are investigating is high-speed communication between modules within a crate over a backplane. For both ATLAS and CMS, the designs use custom backplanes with 160 Mbit/s point-to-point links. Different possibilities will be investigated for the signal levels (e.g. differential ECL versus single-ended BTL), and for hardware choices such as high-density connectors.

For ATLAS, it is not yet decided whether or not the FERMI system will be used for the calorimeter readout. We are therefore studying dedicated trigger digitisation systems that may

be needed in the experiment. In the more immediate future, these are in any case required for performing beam-test demonstrations of the trigger systems. A new 36-channel trigger ADC system is being constructed (see below).

The signals from the calorimeters, in particular the liquid-argon calorimeters of ATLAS, extend over many bunch crossings. Bunch-crossing identification (BCID) logic has to be included in the system so that energy deposits are only considered by the trigger for a single bunch crossing. Such logic is included in the FERMI system, but has to be considered also in the case of independent digitisation for the trigger. As discussed in more detail below, studies were performed in the last year using a single-channel BCID module using data in real time from an ATLAS liquid-argon prototype calorimeter. Next year, a 36-channel system based on field-programmable gate arrays (FPGAs) will be used for more extensive studies.

The ATLAS and CMS triggers are both based on 'trigger cells' with a granularity $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ in pseudorapidity–azimuth space. The electron/photon and jet algorithms are similar (e.g. the electron/photon energy is measured in an area of 2×1 **or** 1×2 trigger cells in both cases); both ATLAS and CMS include a missing transverse energy trigger. All the designs make extensive use of ASICs to achieve the required speed (low latency) and density of processing power. We have already reported [3] on a first ASIC (gate array) implementation of a cluster-finding algorithm, which was operated at full LHC speed in tests with calorimeter prototypes [14]. During the last year, we have designed and submitted to the manufacturer a very high speed full-custom GaAs circuit, described below, that contains adder trees — addition of many numbers is required in all the designs, for example for calculating missing transverse energy. The development of such demonstrator ASICs provides experience of different technologies from various manufacturers, lets us evaluate different processing architectures, and supplies prototypes of key parts of the trigger processors. It also gives us the practical experience of designing ASICs that will be important for building the final systems in which larger (and more expensive) ASICs will be used.

The first RD27 trigger demonstrator system showed the feasibility of a pipelined ASIC for finding e.m. clusters. A phase-2 demonstrator programme will concentrate on getting the data from the detector into the trigger. Apart from the need for fast, compact links, the trigger algorithms demand massive fan-out of signals in order to provide environment information for the e.m. and jet cluster-finding logic. The large number of trigger cells to be examined demands that the data be in serial format in order to keep the pin-count feasible on the modules and ASICs. The following are the key items to study:

- Optical transmission at 800 Mb/s, using laser diodes and/or LEDs.
- Optical receivers and optical-to-electrical converters. Both commercial receivers and custom-built MCMs will be evaluated.
- Feasibility of operating the optical links at 1.6 Gbit/s in order to reduce the cost.
- Transmission of serial data at 800 Mb/s from crate to crate on coaxial cable.
- Fan-out of serial data on backplanes at 160 Mb/s using single-ended signals.

In the following, we describe hardware developments performed over the last year and planned for the next phase of the project.

Bunch-crossing identification

The analogue pulses from the calorimeters typically extend over many bunch crossings, while the level-1 trigger must identify without ambiguity a single bunch crossing that contained an

interaction of interest. Signal processing is therefore included in the digitisation systems so that, as far as possible, transverse-energy values are passed to the trigger processors only for the bunch crossing currently under consideration. Such signal processing is included in the FERMI system [12]. However, the calorimeter readout system for ATLAS has not yet been chosen. In case FERMI is not selected, a separate trigger ADC system with associated bunch-crossing identification logic will be required. An initial study of this aspect of the problem of triggering at LHC has already been made, and more extensive studies are planned for the coming year [15].

Various approaches to the problem of BCID can be envisaged, using analogue or digital signal processing, or a mixture of the two. We are concentrating on digital solutions with programmable parameters. This allows the system to be adapted to calorimeters with various pulse shapes, and for changes be made 'at run time' if necessary. The algorithms under study use digital filtering (Finite Impulse Response filters) to modify the pulse shape, followed by peak finding. This combination (filter + peak finder) allows one to identify the bunch crossing by constant-fraction discrimination or by the zero-crossing time, as well as by the peak time of the original pulse, depending on the choice of parameters in the digital filter. In our studies, we have used the ADC value at the pulse peak to derive the transverse energy value for the trigger cell. This gives adequate performance provided the phase of the clock to each ADC is correctly adjusted to sample at the pulse maximum, although the resolution could in principle be improved by digital filtering (e.g. by integrating the pulse). Hence, the BCID logic is used to flag valid transverse-energy data, out-of-time data being set to zero before transmission to the trigger processor. For low-energy signals, it may not be possible to uniquely identify the bunch crossing. In this case, the data can be passed to the trigger processor without BCID validation, at the cost of slightly increased pile-up noise.

The latency is a critical parameter for all parts of the level-1 trigger system. This has to be taken into account when designing the BCID logic. For example, algorithms such as constant-fraction discrimination that use the leading edge of the pulse give an earlier signal than that from zero-crossing logic. Studies have been carried out [16] to determine which algorithms are best to suited to the calorimeter signals that will be encountered in ATLAS. They are based on data collected in beam tests of a trigger-processing module connected to a prototype liquid-argon electromagnetic calorimeter. Work on algorithms for scintillator-tile hadronic calorimeter data, collected recently using relatively slow analogue pulse shaping, is in progress.

Single-channel BCID demonstrator

A demonstrator BCID module was built to evaluate filtering and peak-finding algorithms, using signals from prototype calorimeters in beam tests in real time. The module was designed to act on a single channel of trigger-cell data and was programmable, allowing a wide range of algorithms to be evaluated at the full LHC clock speed of 40 MHz.

A schematic overview of the module is shown in Fig. 1a. Successive samples of digitised eight-bit data from an ADC (E_{Tin}) enter the module at 25-ns intervals and are stored in a pipeline memory until the BCID decision is available. The FIR digital filter is implemented as pipelined logic (Fig. 1b) — the five coefficients of the filter can be programmed to achieve the desired pulse shaping (in fact the multiplication was performed using look-up tables). The results from the digital filter (E_{Tconv}) are passed to the peak finder, which is also implemented using programmable pipelined logic (Fig. 1c). The peak finder flags valid data (corresponding to pulse maxima). The multiplexer (MUX) component outputs the appropriate datum from the pipeline if the BCID flag is set, or zero if the flag is not set. For diagnostic

purposes, three data values can be read out for each bunch crossing — the delayed input data (E_{Tout}), the output of the filter (E_{Tconv}), and the output of the multiplexer (E_{Tzsup}). Figure 2 shows a block diagram of the complete module with the pipeline, filter logic, and peak-finding logic.

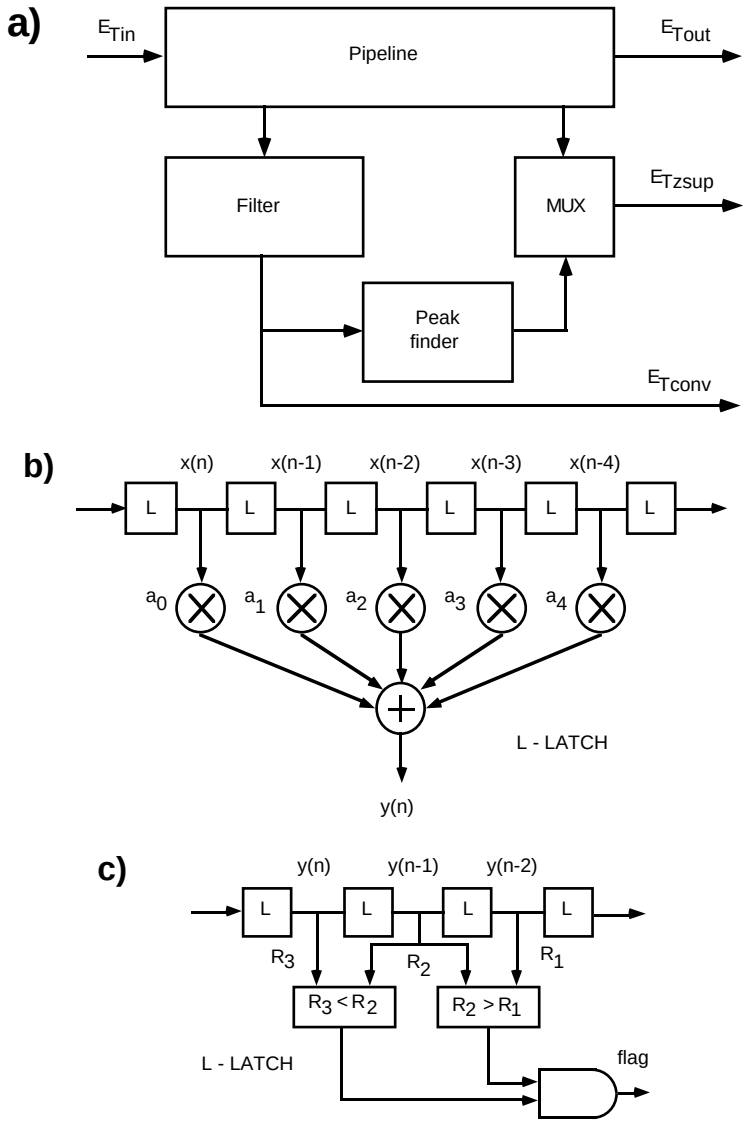

Figure 1: (a) Block diagram of the BCID processor module, (b) Finite Impulse Response (FIR) filter, (c) peak finder.

Beam tests

The experimental set-up used to test the single-channel BCID module is shown in Fig. 3. The demonstrator test rig, already used in previous beam tests [14], was designed to allow for the inclusion of BCID modules, inserted in the data-flow path between the trigger ADCs and the cluster-finding processor module. One channel of digitised data entered the BCID module, and the three results channels (E_{Tout}, E_{Tconv}, E_{Tzsup}) were output to the processor which provided the pipelined readout of the results.

Data were collected with a high-energy (200 GeV) electron beam incident on the RD3 prototype electromagnetic liquid-argon calorimeter. The trigger processor system was operated with a free-running 40 MHz clock, and a beam trigger was provided by scintillator hodoscopes. Following a beam trigger, a time frame of up to 248 25-ns samplings that

contained the full pulse from the calorimeter was read out . A standard CAMAC TDC was used to record the phase of the free-running clock relative to the arrival time of the beam particle. Three channels were used to read out E_{Tout}, E_{Tconv}, and E_{Tzsup}; an example of the recorded data is shown in Fig. 4. Note that in this test no low-level threshold was applied to the input data to the filter. Such thresholding, which is foreseen for the final trigger processor, would eliminate most of the spurious low-E_{T} values seen in the bottom-most plot of Fig. 4.

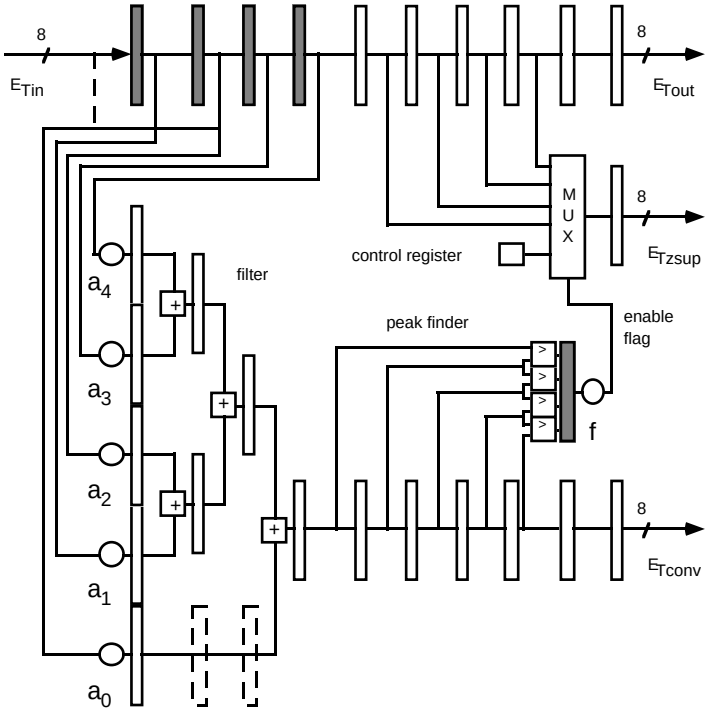

Figure 2: The single-channel BCID demonstrator module.

A total of about 13.5k events, corresponding to more than a million data samples, were analysed. The recorded input data were used to compute the expected results, which were then compared with those from the real-time processing. The value of E_{Tzsup} that would be used in the subsequent processing showed no error, although there were a few (less than a hundred out of more than a million samples) cases in which E_{Tconv} differed from the computed value.

An example of the results on the reliability of BCID algorithms is shown in Fig. 5, which shows the position within the time frame (sampling number) of the pulse peak, as determined in real time by the BCID module. The TDC system was used to select the subset of events in which the phase of the free-running 40 MHz clock sampled the analogue pulse near its maximum, which corresponds to the situation in LHC. For the high-energy data obtained in this beam test, the BCID module reliably identified the correct sampling. We have also studied the performance of BCID algorithms for low-energy data [16], where reliable identification was achieved for measured energies down to 3 GeV.

During the beam tests, the programmable nature of the BCID module allowed the study of several filter and peak-finding algorithms. Further details can be found in Ref. [17].

The results described above were obtained in November 1993 using signals from the RD3 liquid-argon electromagnetic calorimeter, with a pulse peaking time at the input to the ADC system of about 50 ns. Recently, during the ATLAS beam test of September/October 1994, data were collected with signals from the RD34 hadronic scintillator-tile calorimeter for which relatively slow pulse shaping was used — these data are currently being analysed.

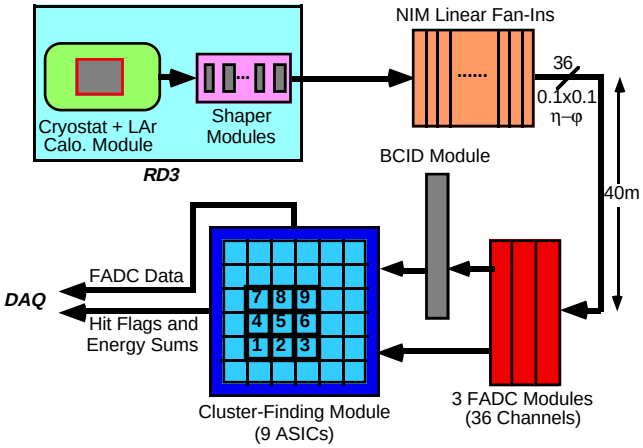

Figure 3: Block diagram of the calorimeter and trigger system used in the test beam.

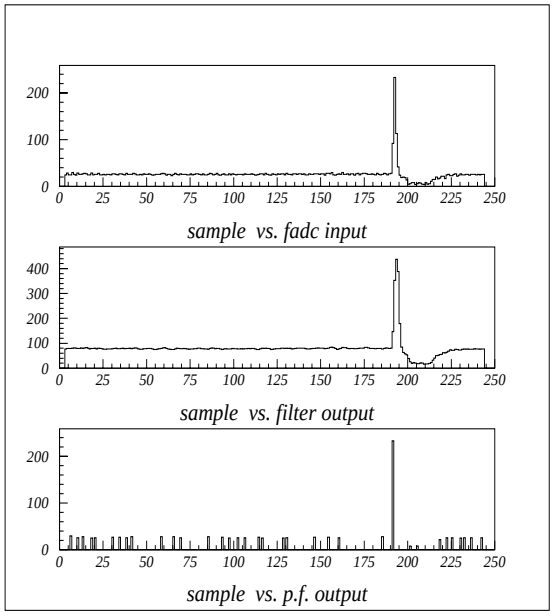

Figure 4: Results recorded in real-time operation of the BCID demonstrator module: the delayed input data (upper plot), output of the FIR filter (middle plot), final BCID output after a peak-finding algorithm (bottom).

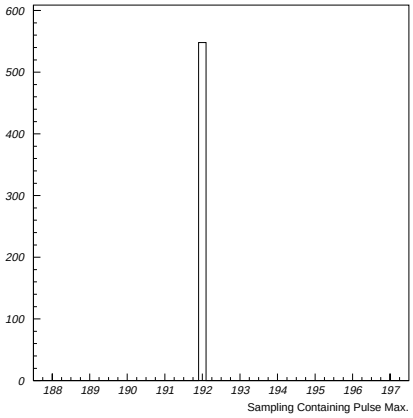

Figure 5: Sampling number containing the pulse maximum.

Future work on BCID

A new BCID demonstrator system is currently being built which will consist of three 12-channel modules. This has been designed to process all 36 trigger channels available in the present trigger demonstrator system. The BCID logic will be implemented using Xilinx field-programmable gate arrays (FPGAs), which leads to a compact design and can be reconfigured to perform alternative algorithms as required. The initial configuration will include filter and peak-finding algorithms optimised for pulses from the ATLAS liquid-argon electromagnetic calorimeter. However new algorithms, including those needed for slower pulse processing, can be tested using the same system after a simple software reconfiguration.

Trigger ADC system for test-beam studies

A new trigger ADC system is currently being constructed as part of the phase-2 calorimeter-trigger demonstrator programme. It will be used in 1995 to provide input data to the BCID modules described above. The ADC system used up to now in tests with calorimeter prototypes is not compatible with the new demonstrator system (signals and connectors), and also lacks some useful diagnostic facilities. The new ADC system improves on it in a number of respects. In particular, the monitoring and test facilities will be much more extensive.

The new ADC system will provide digital data for various tests of calorimeter trigger systems. It will consist of nine four-channel VME modules covering 36 trigger towers. The system digitises the analogue signals linearly into eight-bit values at the full LHC bunch-crossing frequency (40 MHz). The data are transmitted directly to the trigger logic under test.

Data are written simultaneously in a circular manner to 256-sampling deep fast memories until digitisation is stopped following an external signal (i.e. beam trigger). All digitisation results produced during the preceding time interval of 6.4 μs (25 ns × 256) can be read out via VME, e.g. for diagnostic purposes.

The ADC modules also allow one to ‘play back’ data into the trigger processor under test. Previously-recorded data (e.g. from a test beam) can be reloaded through VME into the memories. The clock-driven output allows testing of new trigger components with ‘real’ data in the laboratory at the full 40 MHz rate.

An additional feature is the ≤ 100 MHz sampling capability of the new ADC system. This permits recording of two digitisation samples per 25 ns interval (i.e. 80 MHz). Shaped signals from the liquid-argon calorimeters span as little as 3–5 bunch crossings. Double sampling allows the study of energy resolution and/or efficiency of BCID as a function of the sampling rate.

Adder ASIC design and prototype

We have designed and contracted with Vitesse to build an adder ASIC, which is conceived as a four-stage pipeline with eight input operands and one output operand. Each operand has 10 bits of value, one bit of sign, one bit of input value overflow (tower overflow) and one bit of arithmetic overflow for a total of 13 bits. There are only three stages of adder tree, but an extra level of storage has been added to ensure that chip I/O times are isolated from the adder tree itself. The ASIC has a full implementation of J-Tag/Boundary Scan. We have determined that the ASIC must work reliably at a clock frequency of 200 MHz in order to ensure safe operation at an in-circuit frequency of 160 MHz. Evaluation by Vitesse of the design confirms operation at frequencies of 240 MHz. Delivery of first prototypes should take place in the first half of 1995.

The present adder ASIC forms part of the CMS trigger design, but similar processing is required in ATLAS. Both experiments will benefit from the experience of building high-speed, full-custom ASICs in GaAs.

This ASIC uses four-bit adder macro cells to implement twelve-bit-wide adders. Eleven bits are wired, left justified, to the inputs of the adders. The LSB of each adder will be internally set to ZERO. So, although the adder tree may be constructed from three four-bit adders, the width of the operand data paths is limited to eleven bits.

A chip is designated as 'master' if it is in the first rank of the adder tree and as 'slave' elsewhere. Masters generate 'Tower overflow' (T_{OV}), but slaves only propagate T_{OV}. Both masters and slaves can generate and propagate arithmetic overflow/underflow (A_{OV}). These bits are appended to each input and output operand producing eight 13-bit inputs and one 13-bit output.

These bits are ORed with all the other T_{OV} bits on the first rank and propagated forward through the tree. T_{OV} becomes, effectively, the twelfth bit of the output result. When the ASIC is used as a slave, T_{OV} is the twelfth bit of each input operand. This effectively propagates T_{OV} through to the bottom of the tree regardless of which channel(s) generate the tower overflow. The twelfth bit is T_{OV} and the thirteenth is A_{OV}. The data outputs of the chip are forced to $3FF_H$ when there is either and overflow or underflow.

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

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

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

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

The chip also contains boundary-scan support. The ASIC boundary-scan implementation, along with a proper board-level implementation, should provide full testing capability of the ASIC while it is in circuit. The boundary cells can also be used to verify circuit integrity

(shorts, opens, and stuck at one/zero) at the board level. IEEE standard 1149.1 has been strictly adhered to in order to ensure compatibility with other ASICs or board-level boundary-scan controllers. The full J-TAG controller and a major subset of the commands has been implemented. All inputs and outputs, with the exception of the five boundary-scan control signals, have scan I/O cells.

Initial investigations of integrated optics

A novel feature of the trigger-processor designs for ATLAS is the use of integrated optics to achieve a high-density system. In order to investigate the feasibility of mounting multiple high-speed optical fibres onto MCM substrates, a first demonstrator MCM was built by IMC corporation in Linköping. In this demonstrator, seven banks of eight V-groove fibre mounting pads were placed along each of three sides of an MCM (Figs. 6 and 7). Although the design was intended for pin-diode arrays, discrete diodes were used to validate the construction. So far only two adjacent pin-diodes have been mounted and tested.

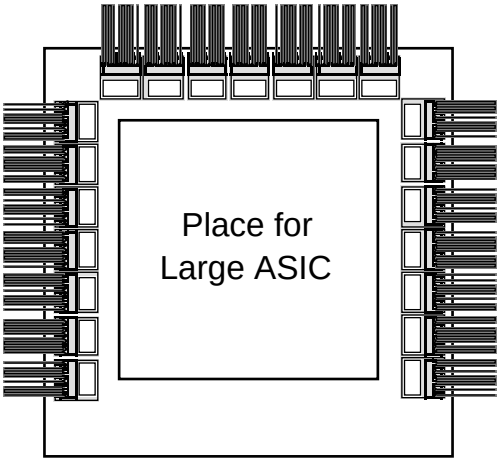

Figure 6: High-density integrated-optics MCM.

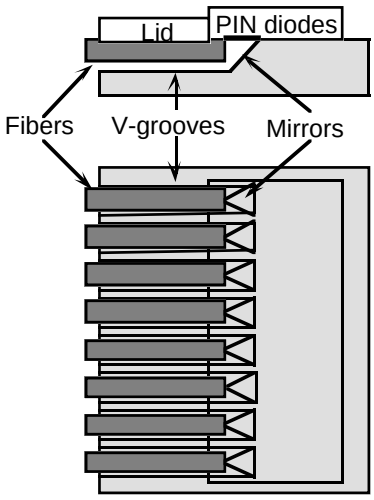

Figure 7: Expanded views of high-density 8-channel integrated-optics demonstrator.

The tests proved that it is indeed possible to mount fibres densely, with a pitch of 0.25 mm which would extrapolate to a $30 \times 30\,\mathrm{mm}^2$ MCM substrate with 168 fibres. Electrical tests showed that 400 Mbit/s input rates would be possible when fed via multi-mode fibres. This choice was a consequence of using LEDs as a light source. (LEDs were chosen since they have better long-term reliability than lasers, especially when exposed to radiation.) The measured bandwidth limit was 250 MHz, mainly due to contact problems. All components

were designed for higher rates and we are confident that the bandwidth limit can be increased after further work.

The V-grooves were fabricated with a pitch of 0.25 mm to fit standard fibre ribbons. The mirror was etched using the 1-1-1 crystal plane with an angle of 54.7°. In the tests, 71% of the light was recorded in the PIN diodes.

Phase-2 demonstrator systems

A phase-2 demonstrator programme [15], which will extend into 1996 including beam tests, has been funded and will investigate key elements of the design described in the ATLAS technical proposal. Digitised data from the 36-channel ADC system will be transferred to three 12-channel transmitter modules (Fig. 8). Each module will process the incoming data through look-up tables and BCID logic (using FPGAs as described above). Pairs of channels will be combined and converted into serial bit-streams each running at 800 Mbit/s. This conversion will initially be performed by a commercial chip-set, but R&D work is underway which may provide a low-cost custom alternative. The serial bit-streams are converted to optical format, by means of laser diodes or LEDs, and injected into optical fibres. Each fibre will be at least 65 m in length to represent the ATLAS data-transmission environment as realistically as possible, and will carry data from two calorimeter trigger cells.

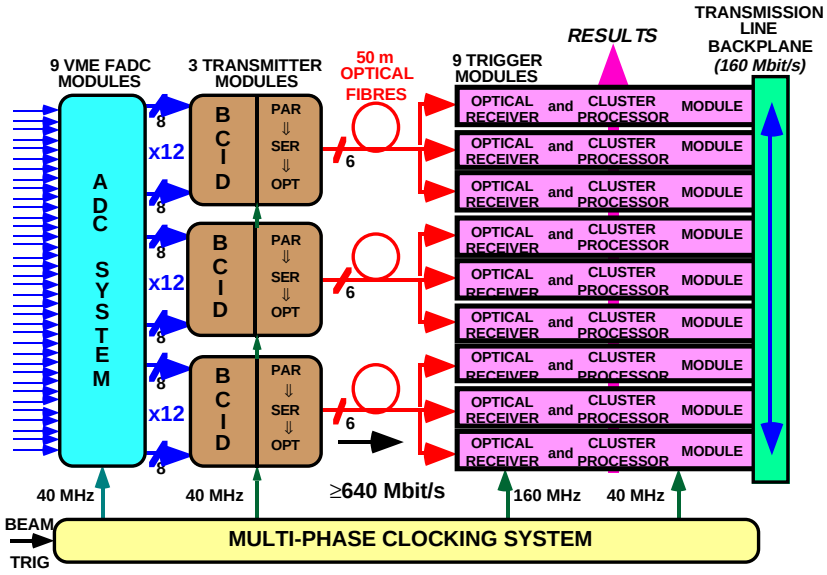

Figure 8: One configuration of the phase-2 demonstrator system.

The trigger modules (Fig. 9) will combine the functions of optical receiver and cluster processor. A total of nine modules, each receiving and fanning-out data from a 2×2 array of trigger cells, will be required to process a 3×3 array in a 6×6 trigger cell window. Input data from the transmitter modules will arrive at each trigger module on two fibres. The fibres will connect to fibre pig-tails on a single 4-channel MCM per module where much of the signal processing will take place. (Only two of the MCM channels will be used, but we wish to demonstrate the feasibility of 4-channel devices.)

Serial data bit-streams will be transmitted from the MCM to the cluster-processing ASICs (via serial/parallel ASICs) and also to a transmission-line backplane for fan-out to neighbouring trigger modules. To execute the cluster-processing algorithm for a maximum of four trigger cells, each trigger module will require data from up to 25 trigger cells, four of which are fed directly from the MCM and the remaining 21 coming from the backplane.

The front-end MCM, shown in Fig. 10, is a crucial component in the system. It contains four identical processing channels, each consisting of an electrical-to-optical converter and two digital processor dies, and performs three main functions:

- Optical-to-electrical conversion of the four incoming serial bit-streams.
- Serial-to-parallel conversion to 16-bit words *(two trigger cells).*
- Time-multiplexing each trigger channel byte into two serial lines, each running at 160 Mbit/s.

The first function — optical to electrical conversion — will be achieved by a novel type of embedded-fibre MCM, which will consist of an optical-fibre pig-tail, accurately located in a chemically-etched V-groove in the silicon substrate, and viewed from above by a PIN diode. The resulting signal will feed a trans-impedance amplifier/discriminator die to produce a logic-level signal. Four of these daughter MCMs will be mounted on the main MCM, with the fibre pig-tails emerging via sealing collars through holes in the package.

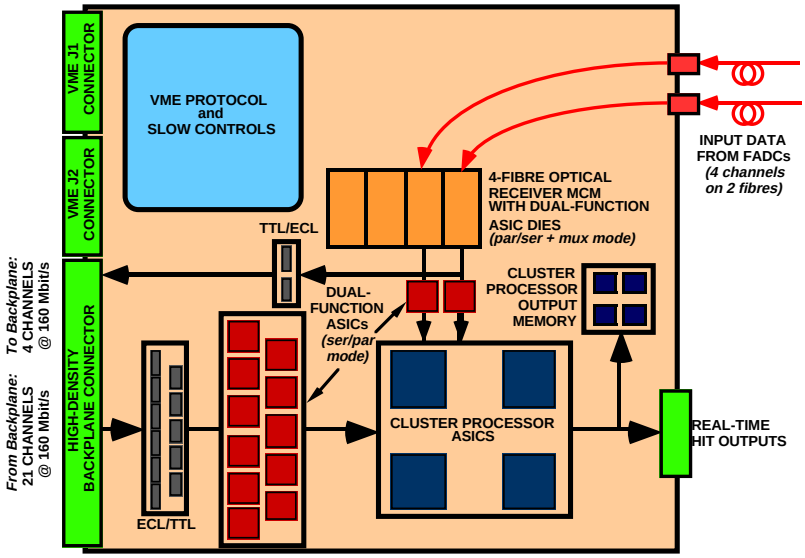

Figure 9: Layout of trigger module.

The second function will utilise commercial serial-to-parallel dies (part of the chip-set to be used in the transmitter modules), one of which will be mounted on each channel of the main MCM to re-generate 16-bit words at 40 MHz. It should be noted that the only place in the system where high-speed (≥ 640 Mbit/s) signals will be found is in the extremely short connections between the embedded-fibre MCMs and these serial-to-parallel dies, which are all internal to the main MCM. No fast clocks will therefore appear at module level.

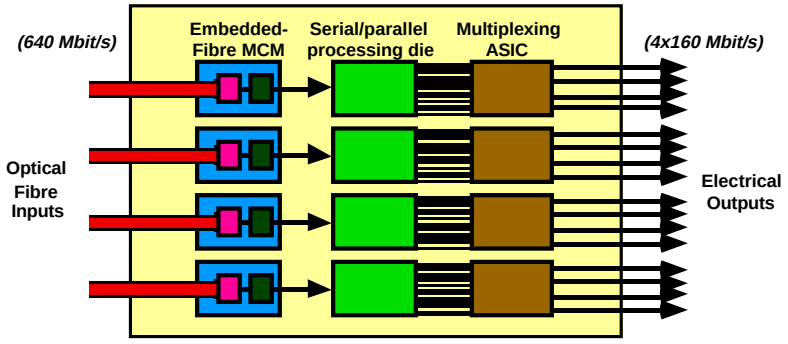

Figure 10: Four-channel integrated-optics MCM.

An important advantage of using such a commercial chip-set is that the fast clock ($\geq$ 640 MHz) is recovered from the data stream via a phase-locked loop, independently on every channel, and therefore does not need to be distributed. This neatly eliminates numerous complications inherent in signal-handling at these very high speeds.

The final processing function — to time-multiplex the eight data bits corresponding to each trigger cell onto two 160 Mbit/s serial lines — will be performed by part of a dual-function semi-custom ASIC mounted in die form on the main MCM.

The cluster-processing ASICs in the ATLAS technical proposal design will process a 4×4 trigger cell array, taking in both electromagnetic and hadronic information. Most of the complexity of such a large new ASIC will be in the processing arithmetic for a 16-trigger-cell algorithm, an operation which has been thoroughly studied in the Phase-1 demonstrator programme. To save costs, therefore, the Phase-1 ASICs will again be used for the cluster-processing function in the phase-2 demonstrator. These chips accept only parallel data, so it will be necessary to provide serial-to-parallel conversion for the 160 Mbit/s serial bit-streams coming from the MCMs. By combining this with the time-multiplexing function on a single chip, a dual-function ASIC can be designed, almost halving the non-recurring engineering (NRE) cost. Four of these ASICs will be mounted on the integrated-optics MCM in unpackaged die form. The serial-to-parallel conversion function, enabling the Phase-1 ASICs to be re-used in the new demonstrator system, will be performed by four more of the conversion ASICs, this time used in a packaged format.

To test the multi-layer high-density transmission-line backplane which will provide electrical data fan-out in the full system, a scaled-down version with restricted capacity will be designed for the demonstrator system crate, fanning-out 36 channels of data between nine trigger modules, but offering a similar high track density. This will enable a study of potential problems of cross-talk, signal reflections and timing synchronisation to be made in as realistic way as possible.

A number of variations in configuration will also be built and evaluated. These include the use of separate optical receiver modules. This has a number of advantages, including the possible use of commercial optical receivers should custom MCMs prove to be too expensive. However, this option demands serial transmission of high-speed data to the cluster-processing modules. Therefore, transmission at 800 Mb/s over short lengths of coaxial cable will be studied. Finally, the feasibility of operating the optical links at 1.6 Gbit/s in order to reduce the cost will be tested.

Alternative technologies

Other technologies are also being considered for ATLAS, as discussed in Appendix C. The feasibility of the suggested design depends on some key technical issues, such as the possibility of bringing a large number of 800 Mbit/s fibres to an MCM, the reliability of bit-serial operations at high clock rates and the performance of the opto/electric components.

At present, a functional demonstrator of these technologies is being designed [9]. A detailed VHDL description down to a Register Transfer Level will serve as a system definition. Simulation results are used for verification. The purpose of this demonstrator project is to prove the feasibility of the system and provide an opportunity to study the system properties before finalising the system itself. The system components used will be identical to the ones envisaged for the final system, if the specifications are not changed. The first step is to build the processing ASIC and the opto/electric MCM and to test them separately using simulated

and test beam data. The second and third steps are to test the ASIC and MCM together and to test a board with several MCMs and processing ASICs operating together.

The demonstrator will be built and tested in two steps over a period of two years. The first step will include the design of a bit-serial processing ASIC and its testing with simulated data, recorded data and then finally on-line data. The prototype ASIC will be designed in BiCMOS. Board design and specification of the interfaces to the front end are still in an early stage. Initially the demonstrator will receive electrical signals from the ADC system being developed within the RD27 collaboration and described above. The demonstrator board will be designed to be read out via VME in a way compatible with the test-beam environment.

In parallel there will be an independent development and test of the opto/electric conversion system presented above.

The following step will be a combined test of a partial system including several processing ASICs and their associated opto/electric circuitry. Whether this stage will also include manufacturing and tests of the merger ASICs is largely a question of funding.

Calorimeter trigger R&D programme for 1995

A number of activities in calorimeter trigger R&D are planned for 1995. The design work for ATLAS and CMS will continue as will the physics simulation studies.

As discussed above, the new trigger ADC system will provide data to FPGA-based BCID modules in beam tests with ATLAS prototype calorimeters. The existing first-demonstrator processor system will receive the data from the FPGA-based BCID system, identifying electrons in real time at the full LHC speed.

For both ATLAS and CMS, an important element in the trigger-processor design is electrical data transmission within a crate using point-to-point links over a 160 MHz high-speed backplane. The connector density between the modules and the backplane has to be high (~300–500 per module). We will perform a number of backplane design studies of high-volume signal transmission, high-density connectors (possibly on both sides of the backplane), clock distribution, power distribution, cooling, location of cabling and accessibility for debugging and maintenance.

We will continue to study synchronous trigger data transmission to the calorimeter trigger processors using fast optical links. Alternatives, such as sending analogue-summed electrical signals from the detector to a remote ADC system, will also be investigated.

We will continue our work on the ASICs that form the heart of the trigger processors. The prototype GaAs adder ASICs that will be manufactured by Vitesse will be evaluated in the laboratory. We will continue design work on the various ASICs that are needed for the ATLAS and CMS trigger processors, sharing our experience of different logic architectures (bit-parallel versus bit-serial) and technologies (both gate arrays and full-custom circuits, in CMOS, GaAs, BiCMOS, etc.).

We will also continue our board and circuit design work for the ATLAS and CMS processors. Design requirements will be developed in order to set up data-flow diagrams and VHDL descriptions. We will study and develop models for the overall calorimeter trigger latency, which is a critical parameter for the design of the front-end electronics. Finally we will address the issue of diagnostics in order to develop a global philosophy that can be applied across trigger systems, involving both ASICs and boards. For example, a study of the J-Tag/Boundary Scan scheme will be made.

3. Muon Trigger

Members of RD27 prepared the level-1 muon trigger design that is described in the ATLAS technical proposal [5]. Only a very brief description is given here, but details can be found in Ref. [18]. The trigger is based on dedicated trigger detectors that have a sufficiently fast response time to uniquely determine the bunch crossing that contained the muon. Resistive plate chambers (RPCs) are used in the barrel, while thin-gap chambers (TGCs) are used in the end-cap. Here the system for the barrel is described as an example, but the same electronics can be used in the end-cap.

The barrel muon spectrometer of ATLAS uses a large air-core toroid magnet. RPC layers are located near the middle and at the end of the toroid, as shown in Fig. 11. High-p_T muons are identified by the fact that they penetrate the absorber and are only deflected by small angles in the toroidal field, and hence 'point back' to the interaction region.

A low-p_T muon trigger is provided, requiring hits in three out of four layers of trigger chambers located near the middle of the toroid. The four chambers are arranged in two groups of two, separated by about 40 cm, and the hits must lie within a road. The high-p_T trigger requires, in addition, hits in two out of three additional chamber layers located after the toroid, at a distance of about 3 m from the other trigger chambers. The chambers give two-coordinate readout and the majority conditions are required independently in each of the two projections.

The trigger has considerable flexibility because the width of the roads can be programmed. By requiring hits in a large number of layers ($3/4 \oplus 2/3$ in each projection for high-p_T muons), backgrounds from random hits are suppressed. The efficiency of the trigger is not degraded by the presence of noise hits in the vicinity of muon tracks. The simplicity of the algorithm allows a low-latency implementation as described below.

Our design study [19] is for an implementation using a programmable coincidence-matrix ASIC. This will provide three programmable road definitions, corresponding to three different p_T thresholds. The design uses two identical ASICs in each projection for each group of 'reference' strips, as illustrated in Fig. 12, giving three thresholds for each of the low- and high-p_T triggers. The final ASICs for ATLAS will contain 32×48 coincidence matrices, corresponding to 32 reference strips and allowing a road width up to 16 strips. About 6000 such circuits will be required to instrument the whole of the barrel and end-cap muon trigger in ATLAS.

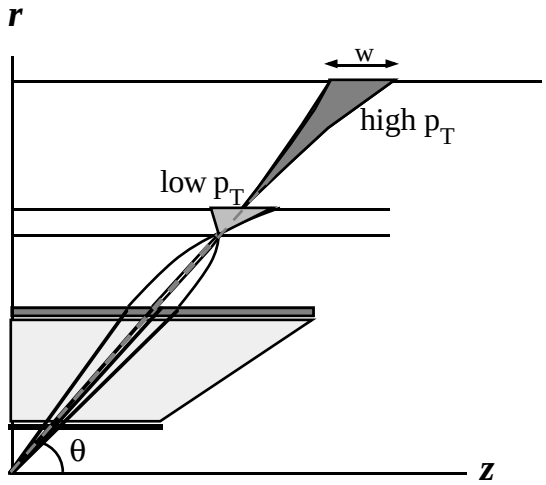

Figure 11: Conceptual design of the level-1 muon trigger for ATLAS.

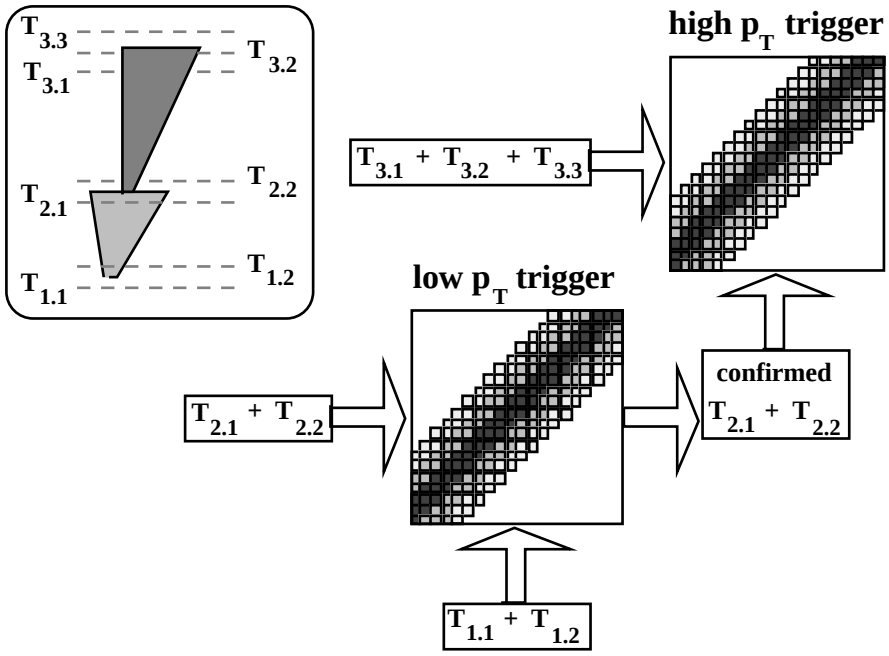

Figure 12. Muon trigger local logic for one projection. The symbols $T_{i,j}$ refer to trigger chamber layers located at different points in the spectrometer. The inputs to the two axes of the coincidence-matrix circuits are the patterns of hits in the different chamber layers. The matrix is programmed to identify combinations of hits that correspond to valid tracks for the different p_T thresholds.

The coincidence matrices are mounted on or near the trigger detectors. This 'local logic' determines if there is a muon candidate in a small area of the detector ('trigger tower', typically $\Delta\eta \times \Delta\phi \sim 0.1 \times 0.25$). The trigger system is divided into sectors in ϕ, and four regions in η (two end-caps plus two half-barrels). So-called 'sector logic' combines the information from all the local logic within a sector, summarising the information for up to two muon candidates per sector. The following information is sent to the muon central trigger processor (located in the underground counting room close to the central trigger processor):

- Number of muon candidates in the sector.
- Trigger tower number within the sector.
- Highest p_T threshold fired for each trigger tower.

The general design concept of this muon trigger has already been validated in a demonstration by RD27 using signals from RPC detectors in RD5, implementing the coincidence matrix using GaAs cross-bar switches [20]. This showed that the time resolution was sufficient to reliably assign the trigger to the correct bunch crossing. In fact, the r.m.s. width of the distribution of the trigger decision time, relative to a timing reference given by scintillation counters, was under 2 ns.

A demonstrator coincidence-matrix ASIC [19], similar to that required for ATLAS, has recently been submitted to a manufacturer. It provides an 8×24 coincidence matrix with two thresholds and is field programmable (see Fig. 13). Each cell provides logic to require hits in three out of four inputs, which allows the full algorithm described above to be implemented using two ASICs per projection — details of how this is done can be found in Ref. [19]. Micro-pipelines are included on the inputs to compensate for different propagation delays from different trigger-chamber planes. The ASICs will be included in a demonstrator trigger system to be evaluated in test beams in 1995, using signals from RPC and/or TGC detectors.

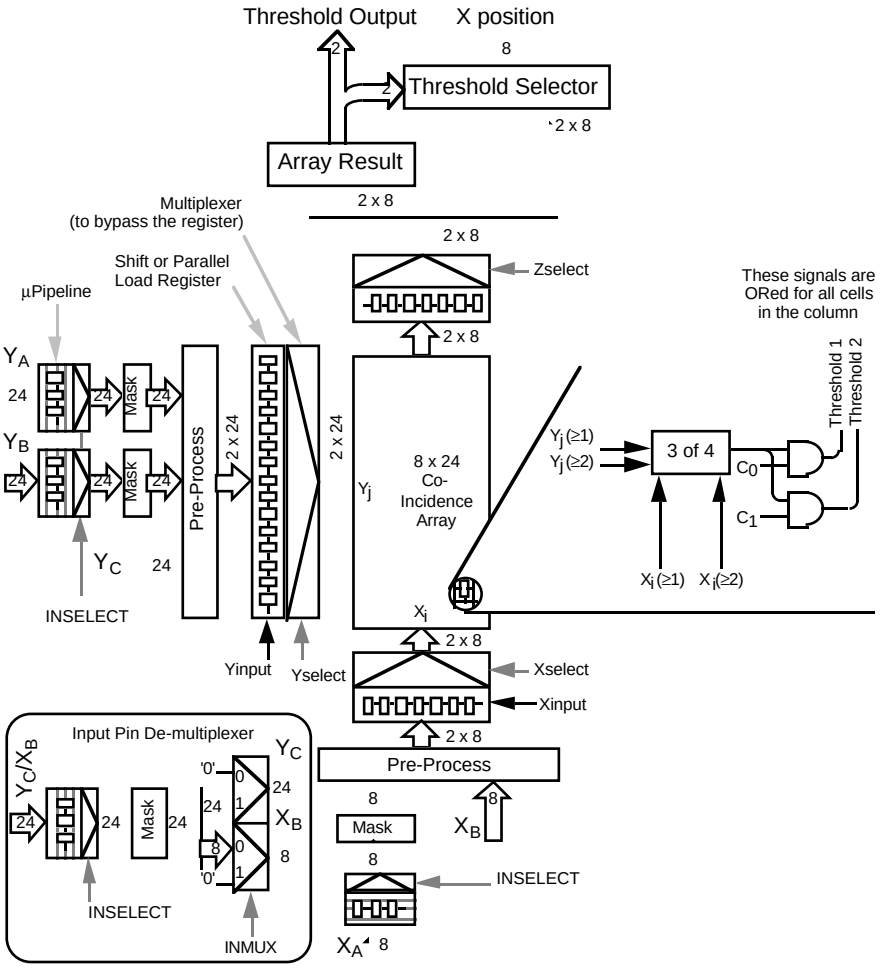

Figure 13: Block diagram of muon-trigger demonstrator ASIC.

4. Central Trigger Processor

The level-1 central trigger processor (CTP) [21] combines information from the different subtrigger processors and makes the overall level-1 yes/no decision. A block diagram of the CTP system that is being studied in RD27 is shown in Fig. 14. The design described here has been adopted by ATLAS, and some of the components are also of interest to CMS.

The logic functions that have to be implemented to combine the different trigger inputs are relatively simple — trigger signatures can be required in coincidence, veto or ‘don’t care’. For example, one might require at least one high-p_T electron in coincidence with large missing transverse energy as a trigger for $W \rightarrow e\nu$. The present design uses a mixture of look-up tables and combinatorial logic to implement the actual CTP algorithm. However, we are also considering the use of FPGAs.

An analysis of the system requirements reveals that the CTP is a rather complex system with many challenging aspects:

- Latency (processing time) is critical since data from millions of detector channels have to be stored in pipeline memories until the level-1 trigger decision is delivered. In the case of ATLAS, only 125 ns is allowed for the CTP processing. This is most easily achieved in a system that is compact, so that time is not used sending signals between different modules.

- The CTP must accommodate a large number of inputs from different subtrigger processors, and allow for test and calibration triggers; it must also allow a large number of trigger combinations to be required in disjunction (OR). Following experience from hadron-collider experiments, and on the basis of work performed in ATLAS and CMS, it will probably be necessary to allow for ~100 trigger combinations of ~100 trigger inputs. Our design study suggests that, using gate arrays, this can be achieved in a single 9U crate. The system we have designed is scalable so that, within the limits of a single crate, the available number of trigger inputs or combinations can be increased by adding extra modules.
- The inputs from the different subtrigger processors will not initially be aligned in time. Phase-adjust circuits (PA) and variable-length pipelines (VPL) are therefore required.
- The ability to prescale individual high-rate triggers has to be built into the system.
- The CTP has to include so-called 'dead-time logic'. This must limit the number of triggers arriving in a short period of time to avoid the overflow of buffers in the front-end systems. Given the short (25 ns) bunch-crossing interval, it is not feasible to implement this using feedback from the front-end systems. A system with programmable parameters is envisaged that will generate a veto signal before problems can occur in the front end.
- Information about which triggers fired must be read out with all events selected by the level-1 trigger. This will be needed in the subsequent trigger levels and in the off-line analysis, e.g. for trigger acceptance studies. The readout has to be pipelined and must introduce only negligible deadtime. It will be necessary to be able to read out the data from a few bunch crossings before and after the triggered one. Experience from experiments at HERA has shown that this is an essential feature when 'timing-in' the trigger system.
- All the input signals and the trigger combinations have to be scaled to measure the rates of the different triggers. This information is essential to monitor the correct functioning of the trigger and detectors, as well as beam conditions.
- The need to monitor the luminosity also has to be considered — the trigger rates provide information on the (relative) luminosity of the LHC. The luminosity may vary from bunch-to-bunch. Precision analyses will have to correct for pile-up effects that depend on the number of events in the bunch crossing that triggered (possibly a few bunch crossings before and after the one that triggered). We are therefore going to investigate the possibility of monitoring the rates of triggers for each of the ~3000 pairs of bunches in the LHC.
- In addition to measuring rates, extensive test and diagnostic facilities must be included in the system. This has already been considered in the design, which, for example, includes circuits to generate test data.
- Of course, a control system has to be provided for setting programmable parameters.

We consider it important for design of the CTP to progress in parallel with that of the subtrigger processors to arrive at a coherent overall level-1 trigger system. This will make it easier to arrive at common standards, for example for control, readout and crate mechanics. Furthermore, a reliable estimate of the overall level-1 trigger latency (a critical parameter for the design of the detector front-end electronics) can only be arrived at once all parts of the system have been designed in detail.

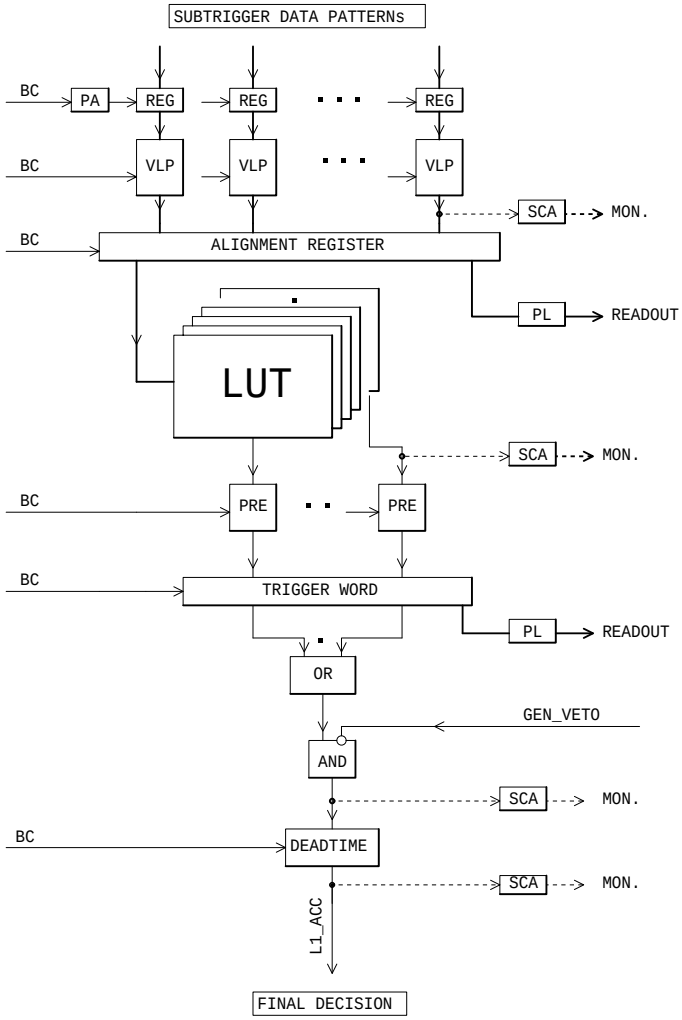

Figure 14: Block diagram of the central trigger processor.

During the last year, much progress has been made in simulating the central trigger processor in a high-level design language. A complete model of the central trigger processor data path has been written in VHDL. There are three levels in the simulation — system level, interface level and functional model. An important feature of the model is that component-level VHDL descriptions can be imported; from the same component-level descriptions, FPGA designs can be synthesised. This feature of the VHDL model has already been used to import a component-level description of an array of scalers into the system-level model of the CTP, and it is planned to extend this to all critical components.

Work has started to implement the critical components in FPGAs. This has already been done for the prescaler circuit and results are expected soon for the variable-length pipeline and scaler array. It is planned to extend this work during the coming year to cover all critical components in the CTP system. Different solutions for some of the components will be investigated and compared. Although the initial hardware work is being done with FPGAs, there are advantages in terms of speed and density of logic in using gate-array ASICs. FPGA implementations provide important information for future ASIC developments. The higher density of gate-array ASICs could reduce the number of physical components and interconnections on modules, thus leading to better system reliability.

While a complete design has been made of the data path of the CTP, much more work is required on the control, monitoring and readout parts of the system. The interface to the

subtrigger processors also has to be defined in more detail. Finally, the physical implementation of the system has to be explored in more detail since signal-path lengths cannot be ignored in the latency calculations.

5. Suggested milestones

We propose the following milestones for the next year of the project:

- Beam tests of a second-generation muon trigger processor incorporating a coincidence-matrix ASIC.
- Beam tests of bunch-crossing identification logic for the calorimeter trigger based on FPGAs, using a new FADC system.
- Development of demonstrators for key components of calorimeter trigger processors, including evaluation of the following: fast optical links and integrated optics; fast electrical data transmission between boards and crates; fast, high-density processing ASICs.
- Evaluation of key components of the central trigger processor in gate arrays.

6. Request to CERN

We request funding from CERN at the level of 100 kCHF in 1995. We also request continued electronic-engineering support from ECP division at the same level as in 1994; this is required for the work on the central trigger processor.

7. Division of work

The areas of responsibility of the participating groups are summarised in the following table:

	Physics Simulation	Timing & Control	Central Trigger	Calorimeter Trigger	Muon Trigger	Level-1 / Level-2
Birmingham	•			•		
CERN	•	•	•	•		•
Heidelberg	•			•		•
Linköping				•		
Munich-MPI	•			•		
QMW, London	•			•		
RAL				•	•	•
RHBNC, London	•					•
Rome 1 and 2	•				•	
Stockholm	•			•		•
Wisconsin	•			•		

Acknowledgements

We express our thanks to ATLAS and CMS, and to other R&D groups with whom we have held useful discussions. We are particularly grateful to RD3, RD5, RD33, RD34 and to the ATLAS calorimeter groups for their help and cooperation during joint beam tests.

References

[1] RD27 proposal, CERN/DRDC/92-17.

[2] Minutes of Research Board 30 June 1992. Minutes of DRDC meeting 4 June 1992.

[3] RD27 status report, CERN/DRDC/93-32.

[4] Minutes of Research Board 30 September 1993. Minutes of DRDC meeting 2 September 1993.

[5] ATLAS Technical Proposal, CERN/LHCC/94–43.

[6] CMS Technical Proposal, CERN/LHCC/94–38.

[7] B.G. Taylor, "Timing, trigger and control distribution for LHC detectors", working reference document, CERN/ECP-RA.

[8] I.P. Brawn et al., "The level-1 calorimeter trigger system for ATLAS", RD27 note 34 and ATLAS-DAQ-NO-30.

[9] G. Appelquist et al., "An R&D programme for alternative technologies for the ATLAS level-1 calorimeter trigger", RD27 note 36 and ATLAS-DAQ-NO-32.

[10] J. Lackey et al., "CMS calorimeter level-1 trigger conceptual design", RD27 note 37 and CMS TN/94-284.

[11] S. Dasu et al., "CMS level-1 calorimeter trigger performance studies", RD27 note 38 and CMS TN/94-285.

[12] FERMI Status Report, CERN/DRDC/94–17.

[13] V. Perera, "A first-level calorimeter trigger processor for the LHC", RD27 note 8.

[14] I.P. Brawn et al., "Beam Tests of a Prototype Level-1 Calorimeter Trigger Processor for LHC experiments", Nuclear Instruments and Methods **A349** (1994) 356–366.

[15] I.P. Brawn, "A demonstrator programme for the ATLAS level-1 calorimeter trigger", RD27 note 35 and ATLAS-DAQ-NO-31.

[16] I.P. Brawn, "Preliminary Investigations into Bunch-Crossing Identification for the Level-1 Trigger", RD27 note 11.

[17] I.P. Brawn et al., "Beam Tests of a Single-Channel Bunch-Crossing Identification Module for the Level-1 Trigger", RD27 note 31.

[18] E. Barberio et al., "Implementation of the first-level muon trigger", ATLAS-DAQ-NO-39.

[19] J. Dowdell et al., "A coincidence array demonstrator ASIC for the RD27 muon trigger", RD27 note 30.

[20] E. Petrolo and S. Veneziano, "Use of GaAs circuits for first-level muon triggering at LHC", Nuclear Instruments and Methods **A344** (1994) 194–198. RD5 Status Report, CERN/DRDC/93–49.

[21] C. Jacobs et al., "First-level central trigger processor for LHC experiments", RD27 note 39.

Publications from RD27

- N. Ellis et al, “A calorimeter-based level-one electromagnetic cluster trigger for LHC”. Proc. CHEP92, CERN 92-07, pp. 210–213.
- B.G. Taylor, “Multichannel Optical Fiber Distribution System for LHC Detector Timing and Control Signals”. Proc. CHEP92, CERN 92-07, pp. 689–692.
- B.G. Taylor, “Multichannel Optical Fiber Distribution System for LHC Detector Timing and Control Signals”. Conf. Record IEEE Nuclear Science Symposium, Orlando, USA, 25–31 October 1992, pp. 492–494.
- E. Eisenhandler et al, “First level calorimeter trigger system for the LHC”. Presented by V. Perera at IEEE Nuclear Science Symposium, Orlando, Florida, USA, 25–31 October 1992.
- E. Petrolo and S. Veneziano, Nucl. Instr. and Meth. A344 (1994) 194.
- I.P. Brawn et al., Nucl. Instr. and Meth. A349 (1994) 356.
- C. Bohm et al, “A bit-serial first-level calorimeter trigger for detectors at LHC”. Presented by C. Bohm at 4th Int. Conf. on Calorimetry in HEP, Elba, 19-25 September 1993.
- I.P. Brawn et al., “First results from a prototype level-1 calorimeter trigger system for LHC”. Presented by A. Gillman at 4th Int. Conf. on Calorimetry in HEP, Elba, 19–25 September 1993.
- B.G. Taylor, “Optical timing, trigger and control distribution for LHC”, presented at IEEE Nucl. Science Symposium, San Francisco, 31 Oct – 6 Nov 1993.
- I.P. Brawn et al., “Implementation of a first-level calorimeter trigger for use at the Large Hadron Collider at CERN”. Presented by R. Staley at IEEE Nucl. Science Symposium, San Francisco, 31 Oct – 6 Nov 1993.
- C. Bohm et al, “A bit-serial first-level calorimeter trigger for detectors at LHC”. Presented by X. Zhao at CHEP94, San Francisco, 21–27 April 1994.
- S. Dasu et al., “The level-1 calorimeter trigger for the CMS detector”. Presented by S. Dasu at CHEP94, San Francisco, 21–27 April 1994.
- I.P. Brawn et al., “A first-level calorimeter trigger for LHC experiments – design studies and beam tests of a first prototype trigger system”. Presented by A.T. Watson at CHEP94, San Francisco, 21–27 April 1994.
- W. H. Smith et al., “Level-1 Calorimeter Trigger for the CMS Detector”, Proceedings of DPF '94, Albuquerque, NM, August, 1994, U. Wisconsin Preprint WISC-EX-94-340.
- S. Dasu et al., “Level-1 Calorimeter Trigger for the CMS Detector at the LHC”, Proceedings of the International Conference on Calorimetry in High Energy Physics, Brookhaven, NY, September, 1994, U. Wisconsin Preprint WISC-EX-94-342.
- C. Bohm et al., Presented by C. Bohm at RTD’94, Dubna, Russia, 27 June – 1 July 1994.
- I.P.Brawn et al., “Bunch Crossing Identification in a Level-one Calorimeter Trigger for LHC”. Presented by Y.Ermoline at RTD’94, Dubna, Russia, 27 June – 1 July 1994.
- V. Perera et al., “A first-level calorimeter trigger for the ATLAS experiment”. Presented by V. Perera at IEEE 1994 Nuclear Science Symposium, Norfolk, Virginia, USA, Oct 30 – Nov 5, 1994.

RD27 notes

1. N. Ellis, Level-1 and level-2 triggering at LHC.
2. J. Garvey et al., A calorimeter-based level-1 em cluster trigger for LHC.
3. BG.Taylor, Multichannel optical fibre distribution system for LHC detector timing and control signals.
4. N. Ellis, Level-1 and level-2 triggering at LHC.
5. V. Perera et al., First level calorimeter trigger system for the LHC.
6. C. Bohm et al., A bit-serial first-level calorimeter trigger for an LHC detector.
7. N. Ellis, Ideas for a local/global level-2 trigger system.
8. V. Perera, A first-level calorimeter trigger processor for the LHC.
9. J.P. Vanuxem and H. Wendler, Central trigger processor for the first-level trigger for LHC experiments.
10. V. Perera, Crate readout controller (ROC).
11. I. Brawn, Preliminary investigations into bunch-crossing identification for the level-1 trigger.
12. R.E. Carney, The level-1 bit-parallel em trigger processor - CERN beam tests November 1992.
13. A. Watson, Physics simulation studies of the first-level electron/photon trigger.
14. E. Petrolo, S. Veneziano, Use of GaAs circuits for first-level muon triggering at LHC.
15. A. Nisati, Accidental trigger rates in the ATLAS muon system.
16. S. Navas Concha, Data analysis for first-level trigger in ATLAS calorimeters.
17. R. Staley et al., Implementation of a first-level calorimeter trigger for use at the LHC at CERN.
18. B.G. Taylor, Optical timing, trigger and control distribution for LHC detectors.
19. A. Gillman, First results from a prototype level-1 calorimeter trigger system for LHC.
20. N. Ellis, Forward calorimetry and the level-1 trigger.
21. E. Eisenhandler, Level-1 triggers for LHC experiments.
22. I.P. Brawn et al., Beam tests of a prototype level-1 calorimeter trigger for LHC experiments.
23. U. Dickow, An analysis of the RD27 level-1 prototype calorimeter trigger processor test beam data.
24. N. Ellis, Distributed processing for fast triggering of DAQ.
25. P. Hanke, M. Keller, ROI Building, Results from a first demonstrator.
26. C. Bohm et al., A bit-serial first-level calorimeter trigger for detectors at LHC.
27. C. Bohm et al., A bit-serial first-level calorimeter trigger for detectors at LHC.
28. S. Dasu et al., The level-1 calorimeter trigger for the CMS detector.
29. A. Watson et al., A first-level calorimeter trigger for LHC experiments – Design studies and beam tests of a first prototype trigger system.
30. J. Dowdell et al., A coincidence array demonstrator ASIC for the RD27 muon trigger.
31. I.P. Brawn et al., Beam tests of a single-channel bunch-crossing identification module for the level-1 trigger – November 1993.
32. S. Dasu et al, The level-1 calorimeter trigger for the CMS detector at LHC.
33. V. Perera et al., A first-level calorimeter trigger for the ATLAS experiment.
34. I.P. Brawn et al., The level-1 calorimeter trigger system for ATLAS.
35. I.P. Brawn, A demonstrator programme for the ATLAS level-1 calorimeter trigger.
36. G. Appelquist et al., An R&D programme for alternative technologies for the ATLAS level-1 calorimeter trigger.
37. J. Lackey et al., CMS calorimeter level-1 trigger conceptual design.
38. S. Dasu et al., CMS level-1 calorimeter trigger performance studies.
39. C. Jacobs et al., First-level central trigger processor for LHC experiments.

Appendix A: The calorimeter trigger design for the ATLAS technical proposal

We briefly describe here the design of a bit-parallel first-level calorimeter trigger processor which has been adopted as the ATLAS baseline system. Full details of the design are given in Ref. [8]. This design is still evolving as our studies and prototype work progress and as new technology becomes available. However, it should be noted that the design has been studied in considerable detail, is based on technology available now or in the near future, and appears feasible both on cost and performance grounds.

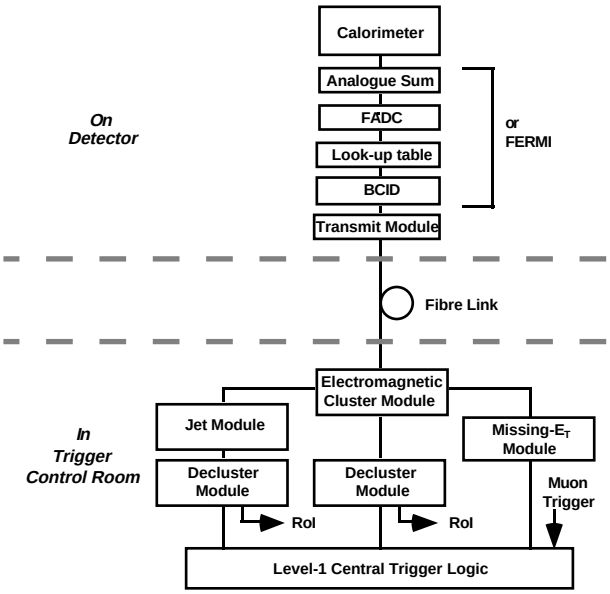

Figure A1: Block diagram of the level-1 trigger system for ATLAS.

The trigger architecture is shown in Fig. A1. The main components are the front-end digitisation and bunch-crossing identification (BCID) logic, a high speed transmission system, and processors for e.m. clusters, jets and missing-E_T. The data reception and e.m. cluster processing is the most demanding part of the system, and this will be housed in four crates. The subsequent processing will require two further crates. The trigger will use six different ASIC designs and seven different circuit-board designs. Figure A2 shows a block diagram of the trigger processor system. The latency of the entire system has been evaluated in some detail, and it can be made to be well under the ATLAS requirement of 2 μs.

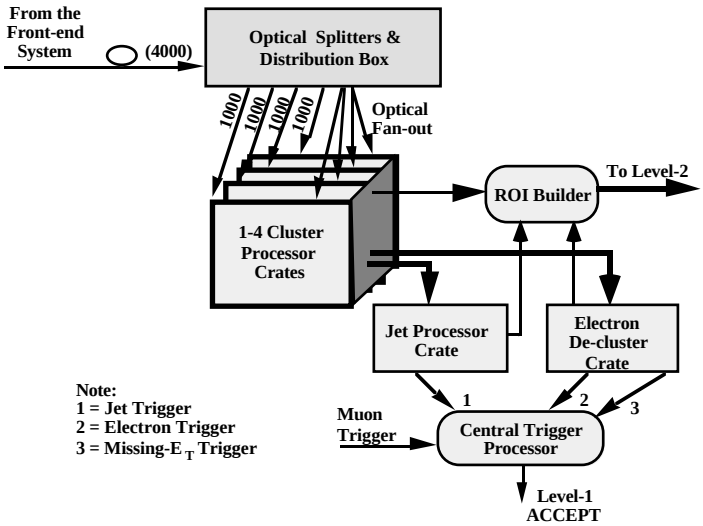

Figure A2: First-level calorimeter trigger system.

Front-end processing

Analogue signals from the calorimeters are summed on the detector to form trigger cells of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ and digitised by an 8-bit FADC with a nominal resolution of $E_T = 1$ GeV per count. The options under consideration include both the FERMI system and an independent trigger ADC system. In the non-FERMI option, the digital data pass through a look-up table where pedestal subtraction and final calibration are applied, yielding an 8-bit E_T measurement for the trigger cell. The data then pass through a BCID stage where the digitisation corresponding to the pulse peak is selected, and other samples are set to zero. Following BCID, two trigger cells are combined for transmission as a 16-bit word from the detector on a fibre-optic link carrying data at 640 Mbit/s. Approximately 65 m of fibre is needed to carry the signals to the underground counting room where the trigger processor is situated.

Cluster-processing ASIC

The design calls for a cluster-finding ASIC to fully process 16 trigger channels. This operation requires E_T from a total of 98 channels (7×7 trigger cells for each of the e.m. and hadronic calorimeters). Each cluster-processing module houses four ASICs and uses E_T from 242 channels. A total of 64 modules are required. Transfer of data into the modules and ASICs at the required rate has proved to be the single most demanding aspect of the trigger design.

At the time of the previous RD27 status report, a preliminary design had been developed in which information was to be sent asynchronously as a single, zero-suppressed 160 Mbit/s serial stream for each channel, requiring derandomizing buffers at each end of the link. Further study showed that the latency associated with the derandomizing buffers was unacceptably high.

In the revised design, data enter the cluster-processing ASIC synchronously as two serial 160 Mbit/s streams per channel, using two I/O pins. Optical signals from 64 fibres are brought directly onto each module

The ASICs produce results for eight sets of threshold values (cluster threshold, e.m. and hadronic isolation thresholds) as hits and as a ‘region-of-interest’ array for the level-2 trigger. A 160 MHz clock is needed for the serial input and output links, but the internal logic will probably run at 40 MHz. We anticipate using a 0.5 μm CMOS gate array with up to 820k gates. ECL pads required to receive the 160 Mbit signals are under development. Approximately 256 cluster-processing ASICs will be needed to process the entire calorimeter.

As the first prototype ASIC runs at more than 67 MHz, the trigger latency will be minimised by reducing the number of pipeline stages and doing more processing per stage.

Integrated-optics multi-chip module

The optical fibres will be received in four-channel integrated-optics multi-chip modules (MCMs) which will generate the 160 Mbit/s electrical signals used in the trigger crates. Each MCM will include four types of unpackaged dies:

- An integrated optical-to-electrical converter with a fibre pig-tail. This device will operate at up to 1 Gbit/s.
- A serial-to-parallel converter to translate the incoming serial bit-streams to 16-bit words (data from two trigger cells) every 25 ns. The parallel-to-serial and serial-to-parallel

converters at the two ends of the transmission chain form a chip-set using the same serial protocol.

- A 4-bit $\times$ 4 serialising ASIC, which will transmit the 16-bit word on four serial links operating at 160 Mbit/s, as required by the cluster ASIC. This ASIC will also provide programmable delays to compensate for variations in fibre delays and different detector response times, and will record the incoming data for readout following a level-1 trigger-accept decision.
- Diagnostic and readout memory.

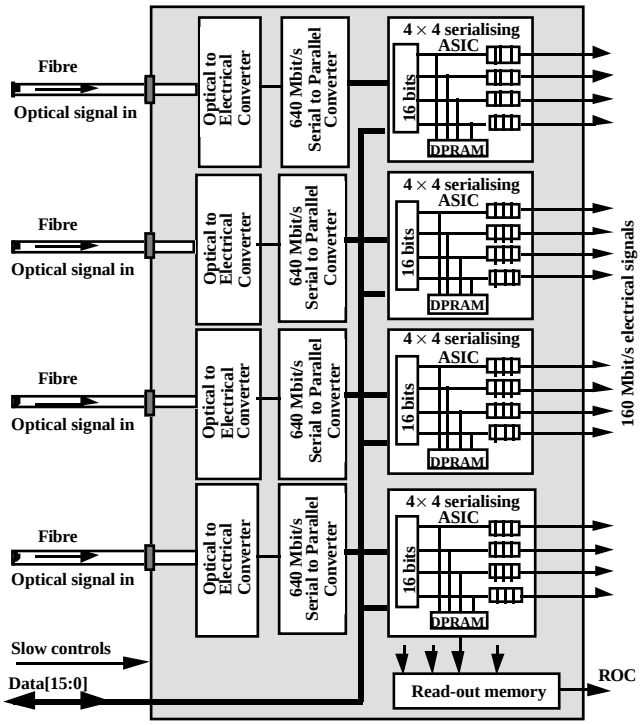

Figure A3: Integrated-optics MCM block diagram.

Sufficient dies for four input fibres will be bonded and packaged as a single MCM, with the fibres entering through the side of the package as shown in Fig. A3. The MCM will thus handle data from eight trigger cells.

E.M. cluster-processor module

The cluster-processor module will contain 16 four-fibre MCMs to receive 64 optic fibres, four cluster-finding ASICs, and one adder ASIC to combine the E_{T} sums, as well as look-up tables to convert E_{T} to its components, E_{x} and E_{y}. A total of 64 such modules will be required for the full ATLAS calorimeter trigger processor.

Results module

The results module will receive E_{x} or E_{y} values from 16 cluster-processor modules and carry out addition using adder ASICs. The twos-complement results of this adding process will be sent to the missing-E_{T} module. A total of eight such results modules would be required.

Missing-E_T module

The function of the missing-E_{T} module is to receive the partial E_{x} and E_{y} sums from the eight results modules and carry out further addition using the adder ASICs before finally testing the

missing transverse energy $E_T = (E_x^2 + E_y^2)^{1/2}$ against four thresholds using look-up tables. Only one such missing-E_T module will be required.

Jet processor module

The jet algorithm will be performed using the 13-bit E_T sums calculated over 4×4 trigger-channel areas of the calorimeter that are available from the cluster processor modules (i.e. $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$ supertowers). These supertowers will be further summed into 2×2 sliding windows, each of which will be compared to eight threshold values. A jet ASIC will be designed to perform this algorithm using data from 4×4 supertowers, i.e. a total of nine jet windows. Approximately 30 jet ASICs are needed in the system, mounted on about eight jet-processor modules.

Cluster-counting module

Simply counting e.m. or jet hits would result in an overestimate due to double-counting of contiguous hits. Therefore a "veto" procedure to look for "corners" will be used. The cluster-counting electronics for jets and e.m. clusters will be identical. The module will examine a 256-pixel array of hits from either the e.m. cluster-processor modules[1] or the jet-processor modules and count non-vetoed pixels. It will then compare the multiplicity with eight multiplicity-threshold values. Part of the vetoing and counting logic will be implemented on a veto ASIC, and to complete the counting an adder ASIC will be used. Each veto ASIC will process 16 pixels, so each module will have 16 ASICs. A total of eight e.m. and eight jet cluster-counting modules would be required.

Readout and crate-controller module

Each crate will be organised by a readout and crate-controller module, to allow communication with the trigger modules and to provide an interface to the level-2 trigger. A built-in CPU might be used to control and format the data and to provide test facilities.

System crates

The four e.m. cluster-processor crates will each process 1024 e.m. and 1024 hadronic trigger cells. In each crate there will be 16 cluster processor modules with 64 fibre optic connectors. There will also be two results modules.

The jet processor crate will include eight jet-processor modules and the eight cluster-counting modules needed for jets.

The e.m. cluster-counting crate will include eight cluster-counting modules and the missing-E_T module.

The crates will be 450 mm high (18 SU) with 20 slots.

[1] For the e.m. cluster counting, the declustering logic acts on a granularity of 4×4 trigger cells, corresponding to the OR of hits from an ASIC for a given threshold. Simulation has shown that this gives satisfactory performance, at least for high-luminosity physics.

Appendix B: The calorimeter trigger design for the CMS technical proposal.

Electron/photon trigger

We have designed an electron/photon trigger based on the recognition of a large and isolated energy deposit in the electromagnetic calorimeter. The algorithm implemented in the hardware design involves two separate cuts on the longitudinal and transverse isolation of the ECAL energy deposit. The first cut involves the hit tower HCAL to ECAL energy ratio, $H/E < 0.05$. A second cut requires transverse isolation, i.e. a cut on a sum of HCAL transverse energies in the nearest eight towers surrounding the hit tower, H1 < 2.0 GeV. In order to reduce the number of bits of information exchanged between electronics cards we limit the dynamic range of neighbouring tower HCAL information to 2 bits.

Jet and isolated-hadron triggers

We have designed a jet trigger based on sums of 0.35×0.35 η–ϕ regions using a dynamic range that covers energies up to about 200 GeV. A second category of jets that exhibit the presence of isolated hadrons can be found by testing the high-E_T 0.35×0.35 regions to determine if a high fraction of the total observed energy is contained in a single trigger tower. Such regions are candidates for isolated hadrons. These isolated hadron jet triggers are used in detecting taus.

Missing transverse energy trigger

We have designed a neutrino trigger that consists of calculating the event missing-E_T vector and testing it against a threshold. The calorimeter trigger calculates both sums of E_T and missing E_T. The transverse energy vector components are calculated from the 8-bit compressed-scale digitised HCAL and ECAL pulse heights converted to a linear scale with a 10-bit dynamic range, and multiplied by entries in look-up tables containing the tower angular coordinates. The HCAL and ECAL sums are then combined into single-tower sums. The tower sums over threshold are routed through the digital summing networks.

Conceptual design for the CMS level-1 calorimeter trigger

We have produced a conceptual design for the CMS level-1 calorimeter trigger. A block diagram of this design with the detail of one calorimeter trigger processor crate is shown in Fig. B1. This crate consists of receiver cards, electron isolation cards, and jet/summary cards, along with support/service cards. Each crate is designed to fully process 256 trigger towers. This density of packaging is achieved by eight pairs of receiver and electron isolation cards. The jet/summary card summarises the data processed in the crate and drives the output signals to the global level-1 trigger processor.

The majority of cards in the trigger processor crates are dedicated to receiving and processing data from the calorimeter, however they require only three different designs. There are eight receiver cards, eight electron isolation cards, and one jet summary card for a total of 17 principal cards per crate, along with DAQ processor, trigger timing and control interface, and a crate monitor card. The backplane is a monolithic, custom, 9U-high printed-circuit board.

Receiver card & synchronisation ASIC conceptual design

The receiver card is 9U by 400 mm. Its basic function is to receive the calorimeter data from fibre, convert to electrical, synchronise with the local clock, and check for data transmission errors. There are also look-up tables on the card to translate the incoming information to transverse energy on several scales. The energy summation tree is also started on these cards.

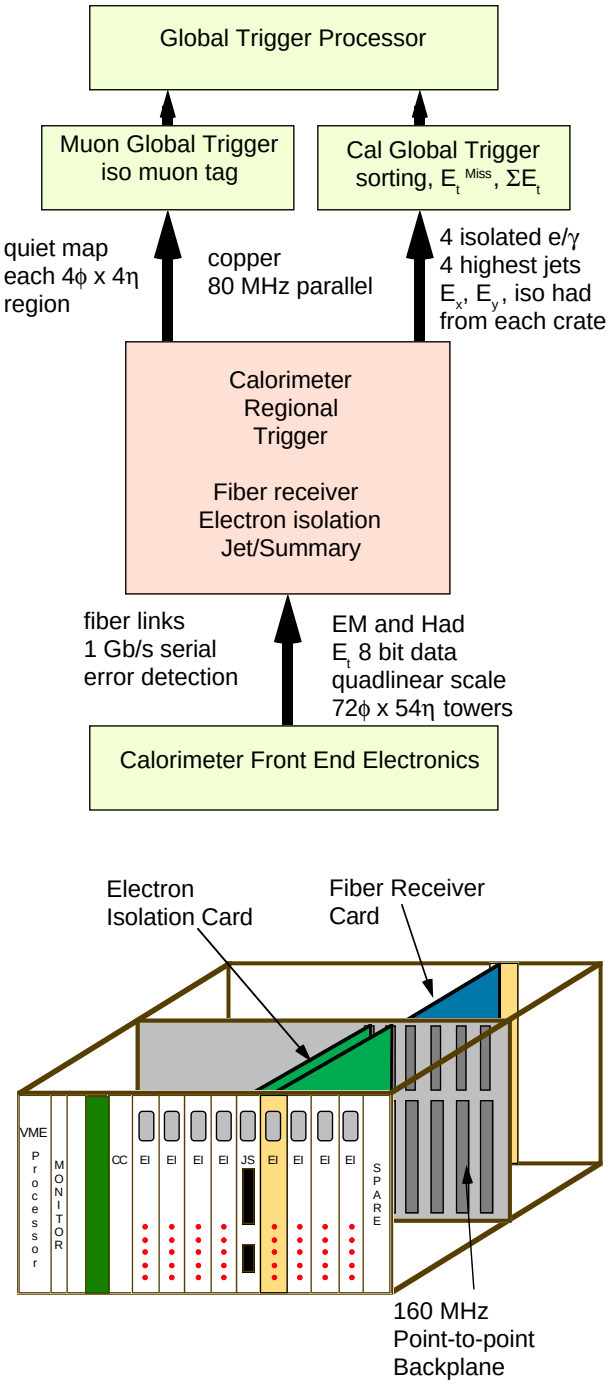

Figure B1. Level-1 calorimeter trigger overview and details of one crate.

Each receiver card is designed to handle up to 32 fibres. Each fibre will have either two towers of hadronic or of electromagnetic information, transmitted in a compressed eight-bit format with 5 bits of error code. The present design uses transmitter and receiver links capable of handling 21 bits of information in 25 ns at a rate of 960 Mbaud.

The outputs of the serial-to-parallel receivers are not only unsynchronised with the local clock but are also not aligned to the same bunch crossing. We have produced a conceptual design for a synchronisation ASIC to receive and synchronise/align four channels of parallel data (84 bits) from the serial/parallel converters.

In order to achieve maximum utilisation of board space, all the logic following the synchronisation ASIC is run at 160 MHz. A significant saving is realised by placing the

multiplexing circuitry, necessary to convert the 40 MHz data flow into 160 MHz, at the output stage of the ASIC. Four input channels, containing eight pieces of 8-bit information, are placed on two output channels of 8 bits. This 4:1 compression requires a corresponding increase in frequency to keep up with the incoming data flow.

Adder ASIC design and prototype

We have designed and contracted with Vitesse to build an adder ASIC, which is conceived as a 4-stage pipeline with eight input operands and 1 output operand. Each operand has 10 bits of value, one bit of sign, one bit of input value overflow (tower overflow) and one bit of arithmetic overflow for a total of 13 bits. There are only three stages of adder tree, but an extra level of storage has been added to ensure chip I/O times are isolated from the adder tree itself. The ASIC has a full implementation of J-Tag/Boundary Scan. We have determined that the ASIC must work reliably at a clock period of 5.0 ns in order to ensure safe operation at an in circuit period of 6.25 ns. Evaluation by Vitesse of the design confirms operation at frequencies of 240 MHz. Delivery of first prototypes should take place in the first half of 1995.

Electron isolation board conceptual design

We have produced a conceptual design for a board that implements the electron isolation trigger described above. Thirty-two towers are processed on each electron isolation board. Data from twenty-eight neighbouring towers is required to determine isolation for towers on the edge of the 4×8 region. All data is transferred between the receiver card and the electron isolation card at 160 MHz. The electron isolation card receives data at 160 MHz in a staged fashion from at most five neighbouring receiver cards and performs the isolated electron algorithm described in the introduction. Some of the data originates in neighbouring crates, but is transmitted through the local receiver cards. The electron isolation card is $9\mathrm{U} \times 280$ mm and resides in the front of the crate. The electron isolation algorithm is performed on this card and the final results sorted to identify the four largest isolation candidates.

The electron isolation algorithm will be implemented in a custom ASIC. The results from the electron identification ASIC are sorted in a second ASIC (sort ASIC) and the top four candidates region are transferred to the jet/summary card. The jet/summary card does a further sort using another sort ASIC to output the top four electron, jet, and isolated hadron candidates in addition to the total E_{T}, E_{x} and E_{y} information in the crate region.

The algorithm used to determine isolation compares the two tower sums of any given tower with its four nearest neighbours. The maximum sum is chosen, and two cuts are applied to the longitudinal and transverse isolation of the ECAL energy deposit. The first cut requires the central tower HCAL to ECAL energy ratio to be < 0.05. The cut on transverse isolation requires the sum of HCAL transverse energy in the eight nearest neighbours to be < 2.0 GeV.

Isolation & sort ASIC conceptual designs

We have produced conceptual designs for the two ASICs to be used on the electron isolation board. The isolation ASIC is designed to shift in the data for 16 towers, 4 towers at a time, over a single bunch-crossing time. The data for 20 neighbouring towers must also be entered in the same time period. The entire 4×8 region can be processed by two ASICs in four 160-MHz cycles. The output of the ASIC is four two-tower sums each 160 MHz cycle and four 1-bit results indicating whether the eight nearest hadronic sums are less than 2.0 GeV.

The four two-tower sums from each of the isolation ASICs are presented in parallel to a single sort ASIC.

The sort ASIC receives all 32 results from the isolation ASIC in one crossing and appends 5 bits of location information to each input. The 5-bit location follows each datum through the sort ASIC and uniquely identifies the four largest. The result from the ASIC is the four largest two-tower sums. There is a four-crossing latency for the result, but the pipeline architecture ensures that once filled with data, a new result will appear every crossing.

Appendix C: An R&D programme for alternative technologies for the ATLAS level-1 calorimeter trigger

Introduction

We briefly describe the design of an alternative first-level calorimeter trigger processor which takes advantage of new possibilities that arise as a consequence of modern design techniques and components, such as optical interconnections, application-specific integrated circuits (ASICs) and multi-chip modules (MCMs). Full details are given in Ref. [9].

This design is homogeneous down to the trigger-cell level. This means that no boundary effects occur due to the system partitioning. The construction presented relies mainly on two different types of highly complex ASIC for processing, and an MCM for opto-electrical conversion of input data.

The trigger processor performs electron/photon identification, jet detection and missing-E_{T} calculations for the central first-level trigger and region of interest (RoI) selection for the second-level trigger. Exploring the possibilities given by advanced technologies leads to a first-level trigger architecture with advantages over more traditional designs, allowing, for example, higher precision calculations. They may also allow a future implementation of more advanced algorithms. The compact design will contain few connectors and a comparatively small number of parts (ASICs and MCMs), a fact which promotes reliability.

A demonstrator programme intended to verify the system performance has been funded and is under development with the aim to manufacture modules and start tests of these during 1995. Depending on the outcome of these tests, the trigger design presented in the ATLAS technical proposal may be modified, wherever it can be shown to bring advantages in terms of cost or performance.

Trigger implementation

The compact calorimeter trigger operates on merged electromagnetic and hadronic calorimeter data from a mesh of 64×64 trigger towers each represented by 8-bit data words.

The electron/photon identification is based on a fully-symmetric algorithm with a local-maximum requirement to obtain declustering. The magnitude of the cluster-pair sum, the isolation environment and the leakage environment is divided into 8, 4 and 4 programmable ranges, respectively, allowing a compressed representation of 3, 2 and 2 bits. For each trigger cell the value of these codes, together with a flag signalling whether the trigger cell E_{T} is larger than its immediate neighbours, is fed into a programmable look-up table to derive a feature code containing a physics classification of the state of that cell. Eight independent classes are foreseen. The global occupancy of these classes is counted and reported to the central trigger processor. Certain of these features will also generate RoIs that will be reported to the second-level trigger. The report will contain the centre position and the feature code. The total energy is in principle available but storing it would greatly increase the memory requirements.

Two types of jet identification algorithms are used, one for the central trigger processor and one for RoIs. The first step in both processes is a smoothing of the combined calorimeter data with a 4×4 kernel to reduce statistical effects, followed by conversion to a 3-bit code using programmable levels. The RoI algorithm reports 3-bit code maxima as RoI centres. The central trigger processor algorithm, on the other hand, uses cluster counting algorithms on binary images obtained by applying seven different thresholds to the data.

The missing-energy calculations are made with 10-bit precision in sine and cosine but without further approximations. The result is translated into one of eight programmable ranges, i.e. compressed into a 3-bit code for transfer to the central trigger processor. The second-level trigger will obtain uncompressed data representing the square of the total transverse energy as well as its x and y components.

The trigger processor design is equipped with programmable classification levels and programmable look-up tables to specify feature definitions and RoI selections.

The first-level calorimeter trigger operates on digital data where each trigger tower is represented by two 8-bit words, one for each calorimeter type. This implementation assumes that data is received from the FERMI digital read-out system, or a system with similar functionality.

The fact that the first-level trigger operations can be expressed as mostly local followed by global merging of results suggests partitioning the system as in Fig. C1 below. The main part of the processing is here performed in weakly interacting local units, which preferably should be entirely located inside ASICs or MCMs.

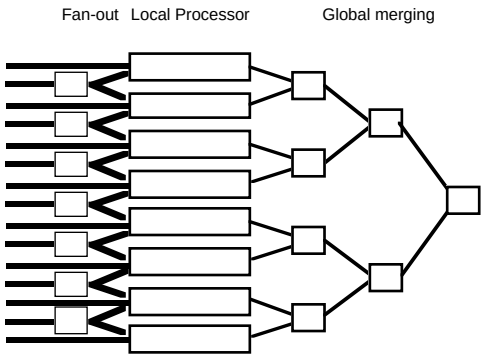

Figure C1: A possible partitioning of the trigger system.

The different 4×4 environment kernels used in the electron/photon identification require access to two extra rows and columns of data above and to the left of the region to be processed, and two rows and two columns below and to the right. Such information sharing may be implemented using intermediate or source fan-out of input data. Fast bit-serial operations will lead to a compact design since the majority of the operations performed are additions.

System layout

The trigger system proposed is based on 128 large processing ASICs, each performing trigger calculations on 4×8 blocks of trigger cells out of the total 64×64 matrix. To allow for 5×5 environment operations around each trigger cell within the block, each ASIC will need information from $(4+4) \times (8+4)$ (i.e. 8×12) cells. This means that $8 \times 12 \times 2 = 192$ input channels are needed, where the factor two has been included to account for the two calorimeters. Each such channel would need to carry 400 Mbit/s to allow for eight serial data bits/trigger cell, plus two flag bits each 25 ns (one flag to signal pulse detection and one for parity).

The transmission-channel implementation is assumed to be based on optical fibres, connected to multi-chip modules (MCMs) which provide the opto/electrical conversions necessary to serve the processor ASICs.

If each processor board can support eight ASICs, these boards may be organised to serve an entire ring around the calorimeter with a width of four trigger cells. In order to give each ASIC access to the extended 8×12 environment, this ring must be able to access two rows of trigger cells to the left and two to the right. This means that each ring will be divided into two regions of width two, one of which will be shared with the ring (board) to the left and one with the ring to the right. A convenient way to solve the information sharing problem is to use passive optical fan-out in the shape of fibre splitters. The additional fan-out that will be required on the boards will occur after the opto/electric conversion.

Data sent via fibres are thus initially grouped into mutually-exclusive 4×8 blocks. These are then expanded to overlapping 8×12 blocks by information sharing and converted into electrical signals.

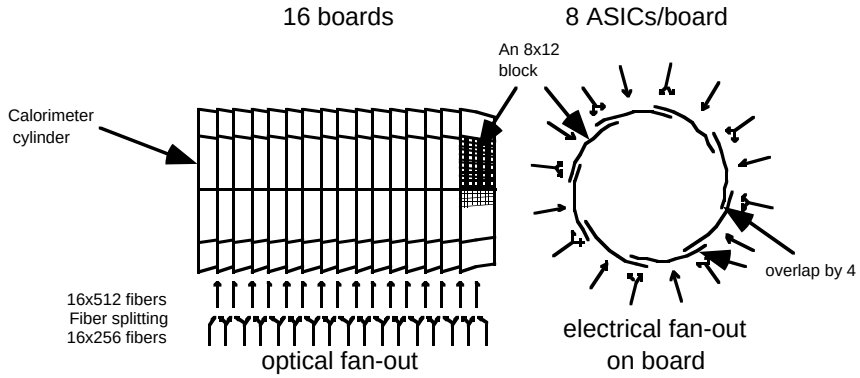

Figure C2. Physical partitioning of the trigger and the corresponding input data fan-out

The major part of the compact first-level trigger is implemented on 16 processor boards. These may be located in one or several crates. All inputs are entered via fibre-optic connectors on the front panel and most outputs are point-to-point links from the processor boards to the supervisor board. This means that the system can be spread out over a number of crates without seriously endangering the signal quality if this is preferred for practical reasons.

The processor board

The processor board will consist of three parts: an opto/electric part, a processing part and a result merger part (Fig. C3). The opto/electric part will be responsible for converting 800 Mbit/s optical information on 512 input fibres to differential electrical signals. The optical signals will be derived from an external fan-out box which provides a duplication of fibres as required to supply the boards with sufficient environment information via passive fibre splitting. The electrical signals are then fanned out to eight large processing ASICs on each board. Results from the processor ASICs are merged in specially-designed merger ASICs to be transferred to the supervisor board via point-to-point links.

The opto/electric conversion is designed on a silicon MCM substrate (Fig. C4) which contains V-grooves for retaining 8 fibres. The light from each fibre is reflected in a 54.7° mirror and projected onto a PIN diode. The signal from the diode is fed to an amplifier bonded to the same substrate and then propagated to the output. The current design assumes that each 8-fibre substrate is contained in a thin SIL package (mounted on the edge) and that 64 of these MCMs will be located immediately behind the 64 8-fibre MT connectors mounted on the front edge of the board. The large processor ASICs will be located behind the MCMs (Fig. C3) and the eight merger ASICs behind these. With a 17-layer 9U standard size PCB (8 signal and 9 power + ground planes) it will be possible to allocate one signal layer to each

differential signal pair from the opto/electric MCM. Signals without fan-out will pass directly from the MCM output to terminated inputs on the processor ASIC. The fanned-out signals will first connect to unterminated inputs on one ASIC and then to terminated inputs on another. The processing ASICs will probably be mounted in 500–600 pin BGA (ball grid array) packages.

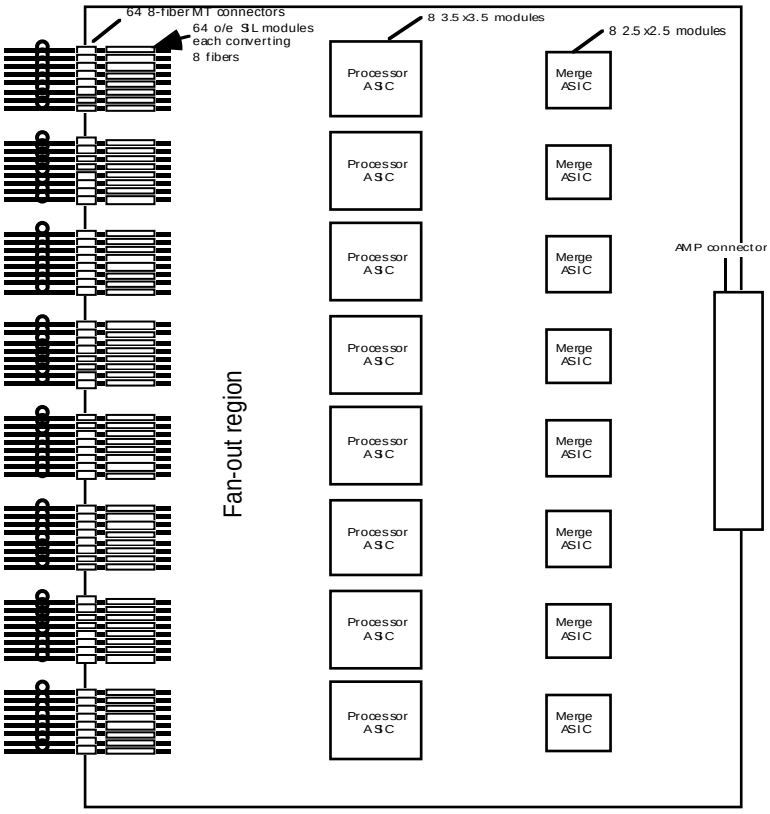

Figure C3. Processing board.

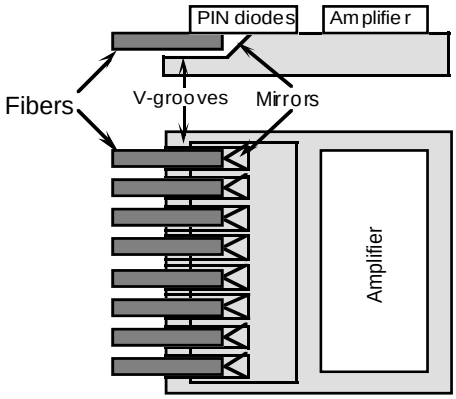

Figure C4. Opto/electric conversion unit with V-groove mounting of fibres.

The supervisor board

The supervisor board contains eight merger ASICs for merging serial feature-count data and two for combining total-energy components. The result is reported to the central trigger processor.

The board also contains logic for driving the daisy-chain read-out of RoI data and for post-processing RoI information (the RoI filter). This part will be implemented using a

programmable gate array. The RoI filter will remove multiple versions of the same jet RoIs from the borderline cases. It will also be able to remove RoIs not required by the second-level trigger processor for a given class of central first-level decisions.

Other essential parts are the memory management unit and the diagnostic supervisor. The former provides insert addresses for each bunch crossing. It will also report the address corresponding to each trigger accepted by the central trigger processor (the extract address). Since a memory position is consumed every bunch crossing, locations corresponding to first-level accepts must be rapidly returned to the memory management unit.