RD27 note 36
ATLAS DAQ-NO-32
16 January 1995

AN R&D PROGRAMME FOR ALTERNATIVE TECHNOLOGIES FOR THE ATLAS LEVEL 1 CALORIMETER TRIGGER

G. Appelquist, C. Bohm, M. Engström, S. Hellman, S-O. Holmgren, E. Johansson, N. Yamdagni, X. Zhao
Department of Physics, University of Stockholm, S-113 85 Stockholm, Sweden

R. Sundblad, A. Ödmark
Sicon, Linköping

P. Bodö, H. Elderstig, H. Hentzell, S. Lindgren, M. Tober
Microsystems, IMC AB, Linköping-Kista

H. Johansson, C. Svensson, J-R. Yuan
Department of Physics and Measurement Technology, University of Linköping

M. Mohktari
Department of Electronics, Royal Institute of Technology, Stockholm

N. Ellis
CERN, Geneva ,Switzerland

ABSTRACT

This note describes a first-level calorimeter trigger processor designed to take 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). The 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 ASICs 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. Remaining degrees of freedom can be used to enhance the ambition level of the trigger design

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 to benefit from improved performance and/or reduced cost.

CONTENT

INTRODUCTION

One of the many challenges when building detectors for the Large Hadron Collider (LHC) is to construct a first-level trigger able to efficiently select potentially interesting events while maintaining sufficient discrimination power against unwanted background events. Since the full data flow through the detector must be buffered during the processing in the first-level trigger, it is essential for the overall detector design to keep the first-level decision latency at a minimum. The present note will describe how the requirements can be met with modern design techniques using high speed data transmission via optical fibres, Application Specific Integrated Circuits (ASICs) and Multi-Chip Module (MCM) packaging technologies in a largely bit-serial systolic processor design. It will also be shown how these techniques suggest compact system designs capable of large degrees of flexibility. These degrees of freedom can, for example, be used to improve the precision in the calculated results. 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.

This note starts by reviewing some design issues that influence the construction of the trigger, then investigates the consequence of these when combined with the trigger algorithms in question. Finally, a trigger design based on these ideas and a demonstrator are presented. Some of the design considerations described here have already been taken into account in the trigger design presented in the ATLAS technical proposal [1,2]. However, in the technical-proposal design, it was decided to use only technology that is or will soon be commercially available, limiting the system speed and density. The design presented here relies on improved component performances made possible by recent technological developments, the viability and reliability of which have to be demonstrated. Depending on the outcome of demonstrator tests, the trigger design for ATLAS may be modified to benefit from the improved technology, wherever it can be shown to bring advantages in terms of cost or performance.

GENERAL DESIGN CONSIDERATIONS

When building complex high performance information processing systems there are several general design principles that should be considered.

Connectivity

The interconnections responsible for propagating information between the processing units are often an important limiting factor, being both costly and vulnerable. When implementing a complex system it is therefore important to partition the system so that inter-module communication is minimised and if possible located at a level where it is easiest to implement.

In the following system hierarchy:

- ASICs
- MCMs
- Circuit boards
- Crates
- Systems of crates (racks)
- System of racks

it is increasingly difficult to implement interconnections with a given high information transfer capacity. This is due to increased cost and power consumption, as well as reduced reliability and more severe bandwidth limitations. For example, MCMs and ASICs normally lack contacts and solder joints that reduce reliability. They also require less power to drive internal signals and can more easily accept high bit rates.

Environment operations

Environment operations, i.e. multi-dimensional mappings where a source point and its environment map into a destination point, play an important role in many applications such as convolutions in image analysis and trigger processing in particle physics. When implementing such operations in a distributed system with many computation units, the input connectivity is of considerable importance.

Size of processing units

If a one point processing unit requires a certain amount of environment data as input, a unit serving a large number of adjacent points will not require proportionally more inputs, in cases when the points share environment information. It is, therefore, desirable to maximise the size of the processing units, since large processing units will reduce the total connectivity requirements.

Input fan-out

Decentralised processing of environment operations also implies that the different processing units must be able to share information, since points on the boundary of one processing unit require access to environment data belonging to neighbouring units. This exchange can be achieved in several ways. Data with several destinations can be fanned-out at the source or they may be shared between the destination units (see fig. 1 *a*, *b*). The former method is best for synchronous systems since it is easier to achieve minimum deviations in arrival time of the different data inputs (skew), by making all signal paths of equal length. The disadvantage is, however, that an increased number of transmission lines is required.

There are also combined solutions: one may defer the fan-out to specially designed units near the destinations (fig. 1*c*). This, however, introduces new system components. Another possibility is a mixed solution with both source and destination fan-out (fig. 1*d*), where destination fan-out is used in situations where it is easy to implement, e.g. on a circuit board but not between boards or in a crate but not between crates.

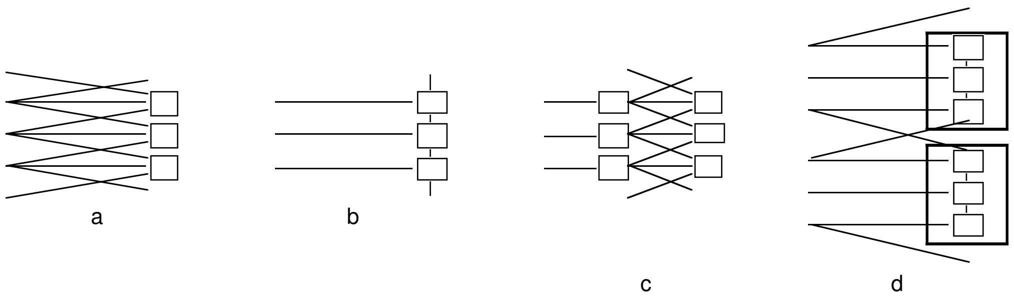

Figure 1 Different fan-out methods: a) source fan-out, b) destination fan-out, c) intermediate fan-out and d) mixed fan-out.

Processing architectures

Different processing architectures can be considered when implementing high performance systems. One possibility would be a farm of fast processors allocated to process data from successive bunch crossings on a “round robin” basis. A sufficient number of processors would then be able to serve data from each bunch crossing. In this case the total latency equals the processing time.

Another possibility is to use a data driven systolic array - a processing pipeline where data are transferred in a wave-like motion from stage to stage, successively performing elementary operations as they move along. Here the latency is the duration of the data flow through the processor.

The latter alternative has definite advantages when applied to fast fixed algorithm computations, since it implies more efficient use of the hardware. All computation stages operates continuously and simultaneously. However, mixed solutions can prove more advantageous when the input rate is high.

Data representation

A systolic array can be adapted to different data representations ranging from purely parallel to purely serial data as well as intermediate combinations of these. The optimum choice depends upon circumstances and the type of operations performed. Table look-ups favour a parallel representation which for pure additions bit-serial representation is more optimal. An n-bit parallel ripple carry adder implementation is roughly n times larger than a bit-serial adder. A bit-serial adder, on the other hand, will need n cycles to complete a full addition. Thus at this point, the number of operations per unit time per unit silicon area is the same. However, ripple carry propagation severely limits the operation speed of the parallel adder. Introducing carry look-ahead or carry-save methods in the parallel adder design will improve the speed but at the expense of larger silicon area requirements. A bit-serial adder, on the other hand, can operate at high clock rates without architectural modifications. On balance it is found that a bit-serial representation allow a much improved processing capacity per unit area and unit time.

An additional bit-serial advantage is due to the fact that in both methods the addition results are delayed one clock cycle with respect to input data. The higher bit-serial clock rate thus implies reduced latency for bit-serial additions. Compact high-speed bit-serial multiplications can, in a similar way, be performed with a latency equal to the multiplicand precision, e.g. the multiplication of two 8 bit numbers will produce a bit-serial result with a latency of 8 (short) clock cycles.

I/O-limitations

Independent of the data representation it is always desirable to reduce the number of I/O-pins. The most obvious way is to increase the input data rates to allow each pin to handle more than one bit of data. The time duration for entering serial data at low clock rates may easily be such that several pins must be used to achieve the required data throughput, combining the information into one channel inside the module if higher rates are required there. Increasing the rate will simplify the situation up to the point where one pin per data channel is reached. Rates beyond this point, however, will require multiplexing to allow many channels to share input pins. Parallel inputs, on the other hand, will require multiplexing to achieve any input rate improvement.

At rates where several serial inputs are required to carry the information into one channel it may be convenient to assign successive data values to the different lines on a "round robin" basis. If n inputs share the task each input will receive every nth data value. If, as is often the case, the different values are inputs to independent calculations then the processing can be decoupled into a farm of n independent processors (fig. 2). The independence means that, other than deriving the data from the same source and reporting the result to the same destination, no interconnections are required, implying an efficient partitioning of the task. If these processors are implemented as separate ASICs, a multi-chip module may well be a convenient carrier. A parallel on a still smaller scale would be to implement the processors as macros in one single ASIC.

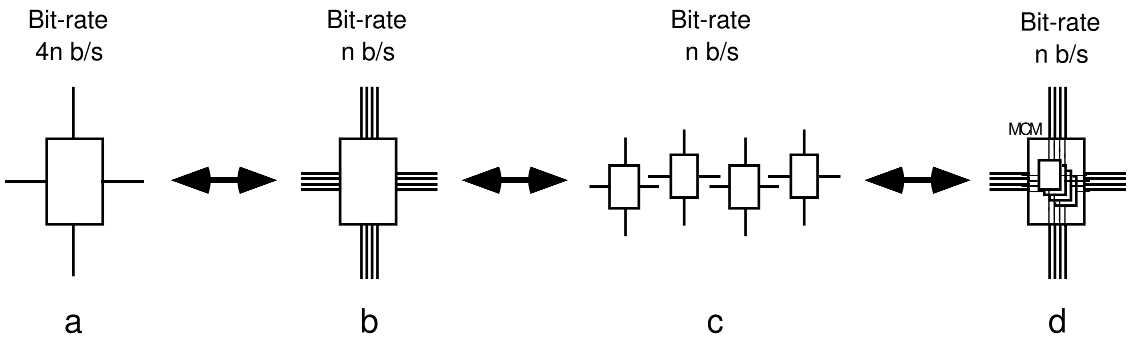

Figure 2. Serial I/O architectures. If high speed serial I/Os (a) are replaced with lower bit-rate I/Os multiple I/O-lines are required (b), which may be decoupled if the different inputs lines carry independent data sets (c). This farm of processing units can also be implemented with multi-chip modules (d)

ALGORITHMS

The purpose of the first-level trigger is to recognise, count and measure features in the collision data that are characteristic of events of special interest [1]. The calorimeter trigger is responsible for the recognition of electrons/photons, jets and indirectly (via missing transverse energy calculations) undetected neutral particles such as neutrinos. Electrons/photons and jets are identified via information about the energy distribution in the electromagnetic calorimeter and hadronic calorimeter. Missing total energy is indicated by a non-zero vector sum of the energy depositions recorded in all calorimeters. Since these tasks do not require high precision calculations, neither in angle nor in magnitude, much is gained by merging calorimeter data into a course grid of trigger towers in azimuth and pseudo-rapidity (i.e. along the beam pipe). A 64x64 trigger cell matrix is chosen in the ATLAS design.

Electron/photons that are not immersed in jets are recognised in the electromagnetic calorimeter as isolated energy depositions, with no energy deposited in the surrounding hadronic calorimeter. The recognition is accomplished by applying (convoluting) the environment operation kernels in *(a)* and *(b)* in figure 3 to each trigger cell in the electromagnetic calorimeter. The four kernels in *(a)* will provide energy sums with the neighbours below, to the right, above and to the left. The maximum of these four values can be chosen as a more accurate representation of the prospective electron/photon energy than the value of the trigger cell itself. Summing adjacent pairs of cells will eliminate problems caused by electrons/photons near the border region between two trigger towers [2,3]. The four kernels *(b)* will provide a measure of the isolation. Two of these alternatives are consistent with the cluster choice (fig. 3). The lower of these values must be sufficiently low to indicate an isolated electron/photon. By applying the *(c)* kernel corresponding to the isolation kernel to the corresponding hadronic trigger towers it will be possible to ascertain whether or not there is a leakage into the hadronic calorimeter.

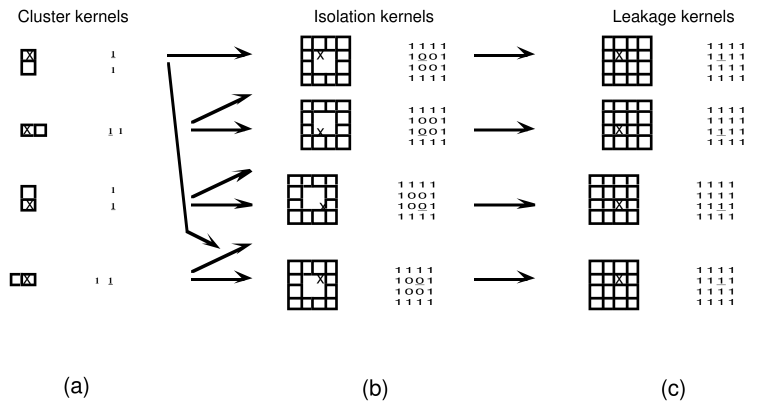

Figure 3. Environment operation kernel relevant to electron/photon and jet identification and their interrelations. The reference cell is marked with x or with underscore.

The above described method is an elaboration of a method used in previous tests [1,2] based on 4x4 environments around the reference cell. The improved method, however, which was mainly chosen for symmetry and simplicity of declustering (see below), requires access to 5x5 environments.

With these criteria a single incident particle will often produce clusters of cells classified as electron/photon candidates. When the electron/photon candidates have been identified and stored in a matrix representation of the calorimeter surface they must be declustered. Declustering will remove all but one identification so that the underlying events can be properly counted. There are several algorithms to perform this correction. Selecting corners with specific orientations (e.g. lower left hand corners) is such an algorithm. A simple and efficient declustering method, which can be used together with the electron/photon identification method described above, is to demand that the reference trigger cell should be a maximum in the immediate neighbourhood, i.e. in the 3x3 area surrounding the reference cell. This method has the additional advantage that it identifies a centrally located cell as a representative of the cluster which is preferred by the second level trigger.

Jets can be identified and their energies estimated by applying one of the kernels from figure 3*(c)* as a sliding window on both electromagnetic and hadronic data, adding the cell values to estimate the jet energies. This operation has a smoothing effect on the stochastic energy distribution. Jets will be efficiently identified by observing the local maxima of these sums, which will also provide estimates of the jet energies. Another possibility is to build a binary jet-identification map by applying a threshold to the sums. Such a map can, after declustering, be used to estimate the number of jets.

As mentioned above, unobserved neutral particles such as neutrinos are identified by observing missing transverse energy in a global vector sum of all recorded energies.

All the above computations are of similar form: first local operations, in two cases followed by declustering, then global merging of partial results via additions. Declustering is an intermediate process which can be performed either locally or globally using different algorithms. For global declustering a central unit needs access to all cluster data, which greatly increases the system connectivity.

The "local maximum" declustering method described above is indeed local. However, when applied to jet declustering it requires access to 6x6 environments. Another method, which can be used with jets, is to achieve local declustering by applying a threshold and then counting all convex boundary corners and dividing the total result by 4. This may lead to an overestimate of the number of jets. An algorithm where concave corners are counted negatively [4] will improve the performance but this time the result will be a lower limit, since internal holes in a cluster will be counted as negative objects. Using both methods simultaneously will provide an interval. Simulation studies will have to be performed to evaluate the efficiency of these methods. However, since the above methods themselves are environment operations, they require, apart from the cells to be declustered, access to an extra row and column from neighbouring units.

TRIGGER IMPLEMENTATION

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 [5], 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 figure 4 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.

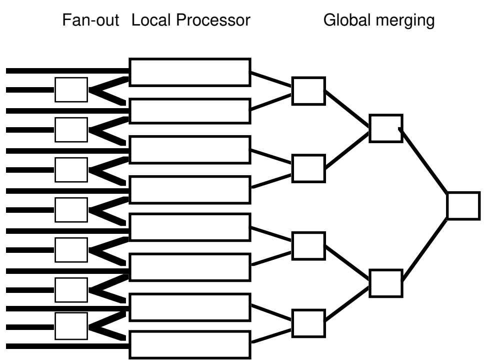

Figure 4. A possible partitioning of the trigger system.

The different 4x4 environment kernels (fig. 3) 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. Table look-ups still requires parallel data representation.

System layout

The trigger system now proposed is based on 128 large processing ASICs, each performing trigger calculations on 4x8 blocks of trigger cells out of the total 64x64 matrix (fig. 5). To allow for 5x5 environment operations around each trigger cell within the block, each ASIC will need information from (4+4)x(8+4) (i.e. 8x12) cells. This means that 8x12x2 (192) input channels are needed, where a factor 2 has been included to account for the two calorimeters. Each such channel would need to carry 400 Mbit/s to allow for 8 serial data bits/trigger cell, plus 2 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. Multiplexing data two ways to 800 Mbit/sec reduces the number of fibres by a factor 2. However, since differential inputs are required for safe transmission between the MCM and the processing ASIC, the number of inputs to the ASICs remains at 192.

If each processor board can support 8 ASICs these boards may be organised to serve an entire ring around the calorimeter with a width of 4 trigger cells. In order to give each ASIC access to the extended 8x12 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 2, 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 fibres splitters. The additional fan-out that will be required on the boards will occur after the opto/electric conversion.

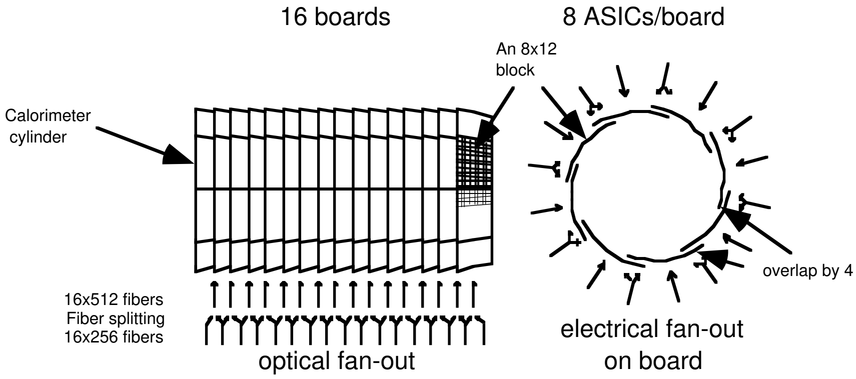

Figure 5. Physical partitioning of the trigger and the corresponding input data fan-out

A different system partitioning dividing the calorimeter surface into blocks of 16x16 trigger cells would be more optimal because it would reduce the amount of information sharing requiring optical fan-out. In fact, this would reduce the number of input fibres to each processor board from 512 to 400. It would also reduce the need for optical fan-out from 100% to only 36% of the fibres. However, the fan-out would be less uniform, optical fan-out of 1, 2 and 4 would be required. This requires sufficient light in the fibres to allow for the 4-way splitting, at least in the fibres that will be fanned out 4-ways. The same argument applies to the electrical fan-out on the boards. With this solution there would be a trade off between accepting an excessive amount of optical and electrical signal power or accepting an inhomogeneous solution with many different types of opto/electric drivers.

A functional description

Data sent via fibres are thus initially grouped into mutually exclusive 4x8 blocks (fig. 6). These are then expanded to overlapping 8x12 blocks by information sharing and converted into electrical signals (not shown in the figure).

Electron/photon identification

Figure 6 shows the electron/photon identification part where initially the kernels from *(a)*, *(b)* and *(c)* in fig. 3 are evaluated with electromagnetic (*a*, *b*) and hadronic data (*c*) for trigger cells in the original 4x8 environment. The required data can be extracted from a pool of all possible results from operations with the kernels in figure 3. There are 45 different ways to apply the isolation and leakage kernels. However, out of the 88 possible horizontal pairs (fig. 3*a*) and the 84 possible vertical pairs only 36 and 40, respectively, belong to the 4x8 region. For each trigger cell in the 4x8 environment the maximum of the four cluster pair sums is representative of the electron/photon energy associated with the trigger cell. The lower of the two isolation sums corresponding to the chosen cluster pair is selected. The hadronic leakage sum corresponding to the isolation sum is also selected. After this process the three sums are converted into compressed codes by comparing them to programmable energy levels. The cluster energy is compared with 7 levels placing the value in one of 8 different intervals, the identity of which can be specified by 3 bits. The other two sums are similarly compared with 3 levels, giving 2-bit codes as results. An additional calculation delivers a flag signalling whether the trigger cell is a maximum within its immediate environment (3x3). This is for declustering purpose. Thus the outcomes of the initial calculations are expressed in 8 bit codes, one for each trigger cell in the initial 4x8 block.

The 32 8-bit codes are then fed to a programmable look-up table (LUT) where they are translated into a declustered 8-bit feature code. This code summarises the calculations by placing the trigger cell into independent physics related categories like: high energy well isolated electron/photon, medium energy..., medium energy poorly isolated electron/photon....or none of the above.

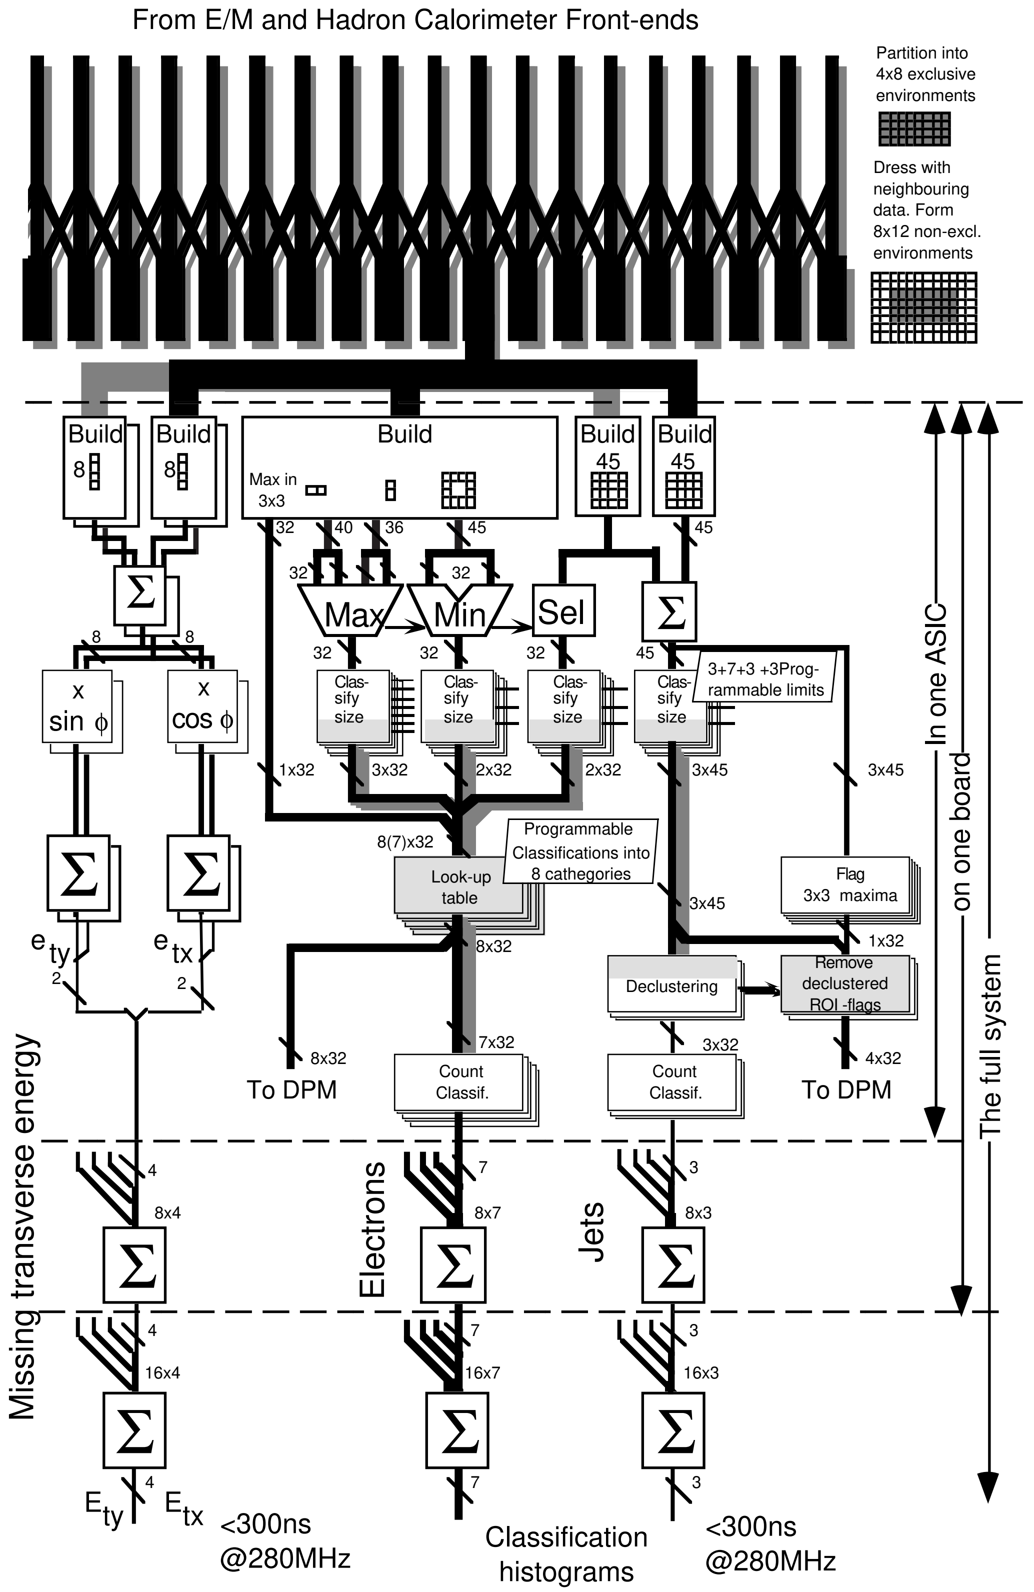

Figure 6. The computation flow for electron/photon and jet detection and for E_T evaluation. All blocks except look-up table and part of the magnitude classification are based on bit-serial implementation. (Bit-parallel grey and bit-serial white). The final additions occur on a supervisory board.

32 LUTs of the size 256x8 are required. However, since they have identical content they may in principle be realised as one 32-port 256x8 RAM. The practical solution is somewhere in-between. Eight 4-port memories is the choice in the present design.

The final step is to count the number of instances of each category, first locally, then on the board and ultimately, over the system. All operations described except for the memory look-up are performed bit-serially. This also simplifies the merging of data outside the ASICs.

Jet identification

The inclusive 4x4 kernel (fig. 3*(c)*) operating on the sum of electromagnetic and hadronic data provides estimates of possible jet energies, subsequently compared with 7 programmable levels to obtain a compressed 3-bit magnitude code, as in the electron/photon identification case. The declustering will utilise an extra row and column of the data to determine whether a positive (concave) or negative (convex) corner has been encountered in the present cell position. This procedure is repeated for 7 levels (0-6). The total number of corners (positive minus negative) are counted for each level over the ASIC(4x8), over the board (4x64) and finally over the entire system (64x64).

The magnitude codes are also temporarily stored in a dual port memory, together with a flag signalling whether the cell is a maximum within its 3x3 environment (2x3 or 2x2 if the cell is situated on the border or in a corner). The local maxima are later used to derive jet RoI coordinates (see below).

Missing transverse energy

The total transverse energy is estimated by summing input data corresponding to the same azimuthal angle from both the electromagnetic and hadronic calorimeters and then multiplying by the appropriate sine and cosine (with 10-bits precision). This addition is performed in groups of 4 since the extra environment data are not involved here. The two energy components are then added on the ASIC, on the board and finally between boards. At 280 MHz only seven bits can be treated for each 25 ns interval. The circuitry performing the total missing transverse energy sum will have to be duplicated, to be used alternatively in order to cope with the precision requirements.

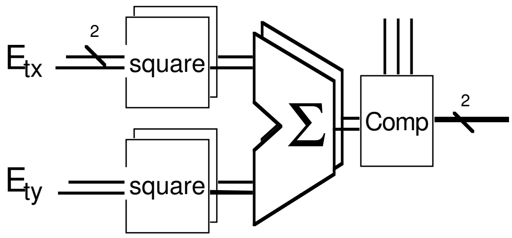

Figure 7. Final calculation of missing E_T.

The total E_T components are then squared, added and compared with programmable levels to determine compressed codes expressing the magnitude of the total momentum.

RoI processing

Data from electrons/photons and jets as well as results from missing E_T calculations will be stored in local dual port memories after an insert address has been supplied (every 25 ns) by a memory management unit [6] located on the supervisor board. The address locations will be protected against overwriting until a first-level reject has been issued by the central first-level trigger or until the protection has been removed by the second-level trigger after reading out all relevant information.

RoI information is read out using a daisy chain mechanism after an address pointer has been supplied, at the request of the second-level trigger. Only those memory positions are read which have the read-out flag set. The read-out flag is the maximum flag in the jet case. In the electron/photon case it is the 8th bit from the programmable LUT.

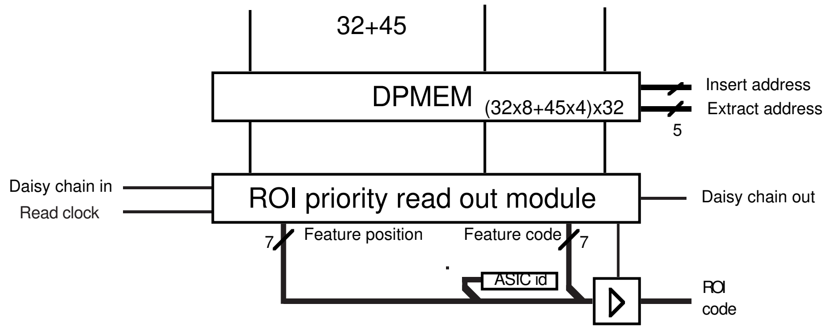

Figure 8. RoI read out unit.

The RoI sequence is intercepted by a RoI filter located on the supervisor board. The purpose of this unit is to filter away RoIs that are irrelevant given the actual trigger decision type. It is also necessary to remove jet maxima that occurred on the border of the 4x8 blocks if they were not reported by both (or in case of corners all four) units.

A System Description

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. 9). 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 8 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.

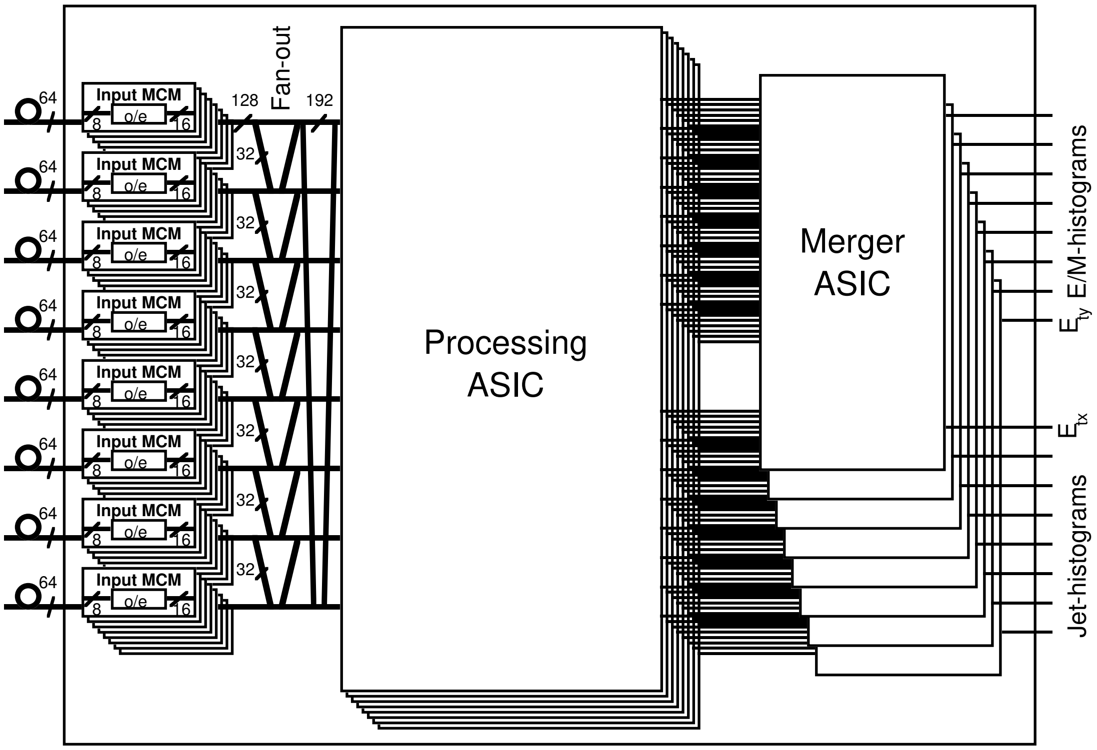

Figure 9. Processing board.

The opto/electric conversion is designed on a silicon MCM substrate (fig. 10) which contains V-grooves for retaining 8 fibres. The light from each fibre is reflected in a 54.7^{o} 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. 10) and the 8 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 an other. Approximate calculations indicate that the power dissipation into the MCMs, processing and merger ASICs would be about 1, 20 and 5 W respectively. This would add up to a total power dissipation of about 300W/board, and about 5kW for the entire system.

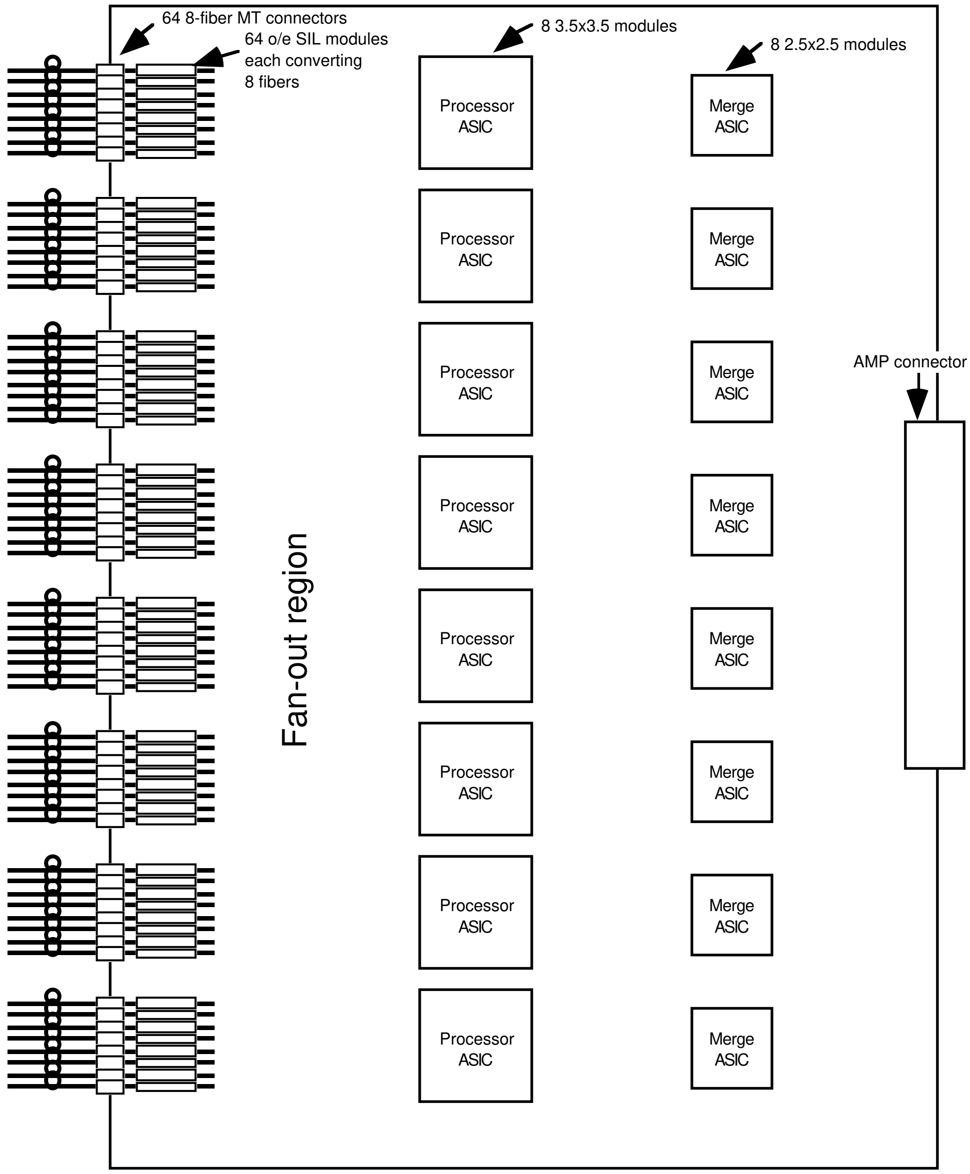

Figure 10. Processing board layout.

The processing ASICs will probably be mounted in 500-600 pin BGA (ball grid array) packages. Such a package has been used for the 620 Power PC chip which dissipates 30W of power. It is essential that the path from the I/O pin to the receiver pad on the ASIC is short especially for the unterminated inputs. There are several connectors with high frequency capabilities available (e.g. from AMP) that will be acceptable for the frequencies in question (300MHz)

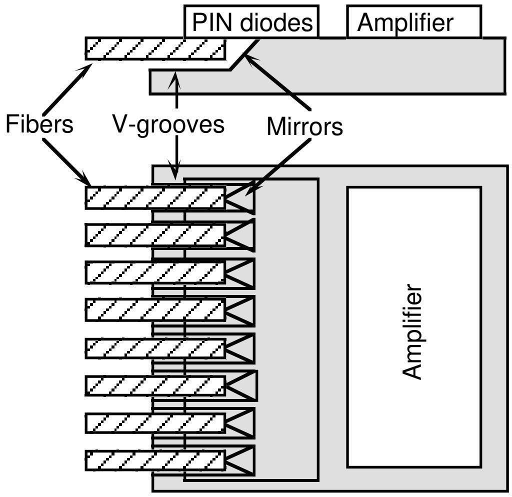

Figure 11. Opto/electric conversion unit with V-groove mounting of fibres.

The processing ASIC

The processing ASIC performs most of the processing required by the system. Data enter from the opto/electric converter via an electrical fan-out, received in a differential receiver and delayed with a suitable programmable delay of up to one bit in duration. This analogue delay will be chosen so that the subsequent sampling unit will sense the signal at an optimal position for proper bit alignment. The sampling is controlled by a local 800 MHz clock generated via a PLL from the 40 MHz system clock

Data are then demultiplexed into five 160 Mbit/s bit-streams and synchronised by programmable barrel shifters and word delays that can delay data by up to 3 25 ns clock cycles. This will ensure word and context alignment. After demultiplexing and synchronisation the data is fed to a (64 event diagnostic buffer) diagnostic memory and a serial to parallel converter, whose outputs alternately feed two outputs at the increased speed of 280 Mbit/s. This speed is more optimal for the CMOS process. A rate of 280 MHz will also permit a sufficient number of trailing significant zeroes in the bit-stream (data are sent with the LSB first) to accommodate larger numbers generated later in the adder tree.

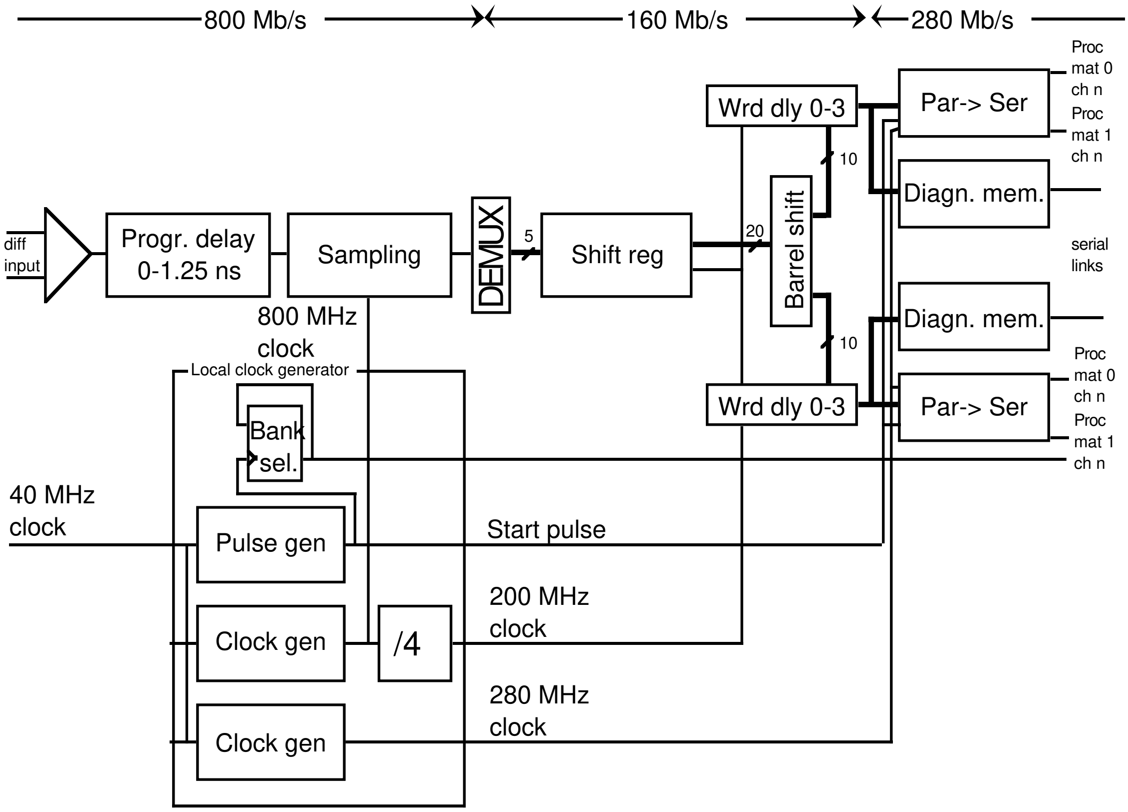

Figure 12. Input module for the processing ASIC.

The input stage also includes an extraction circuit for the pulse detect flag, located in the parallel to serial module, which is the first bit in the data frame. Only data words with even parity and the pulse detect flag set will be propagated into the processing ASIC. Parity violations will be reported to the diagnostic processor.

The two different data paths are fed to two bit-serial processing matrices operating at 280 MHz. While serial data enter from one side serial constant memories provide matching bit-streams from the other side. The 8x12 matrix consists of 96 basically identical units that feed their bit-serial results to three destinations. One destination has multiplication and summation units to produce transverse energy components and the others provide electromagnetic and jet post-processing. The electromagnetic results are parallelised and sent to look-up tables to be translated into feature codes. The abundance of these codes in the 4x8 environment of the ASIC is then determined by counting. The jet results are fed directly to counting modules. Both electromagnetic and jet results will also be stored in a dual port memory for possible future RoI processing. All results are reported via differential lines at 280 MHz to insure satisfactory noise immunity. Since this rate is not sufficient to transfer missing-energy data with the required precision, four lines are allocated for that purpose.

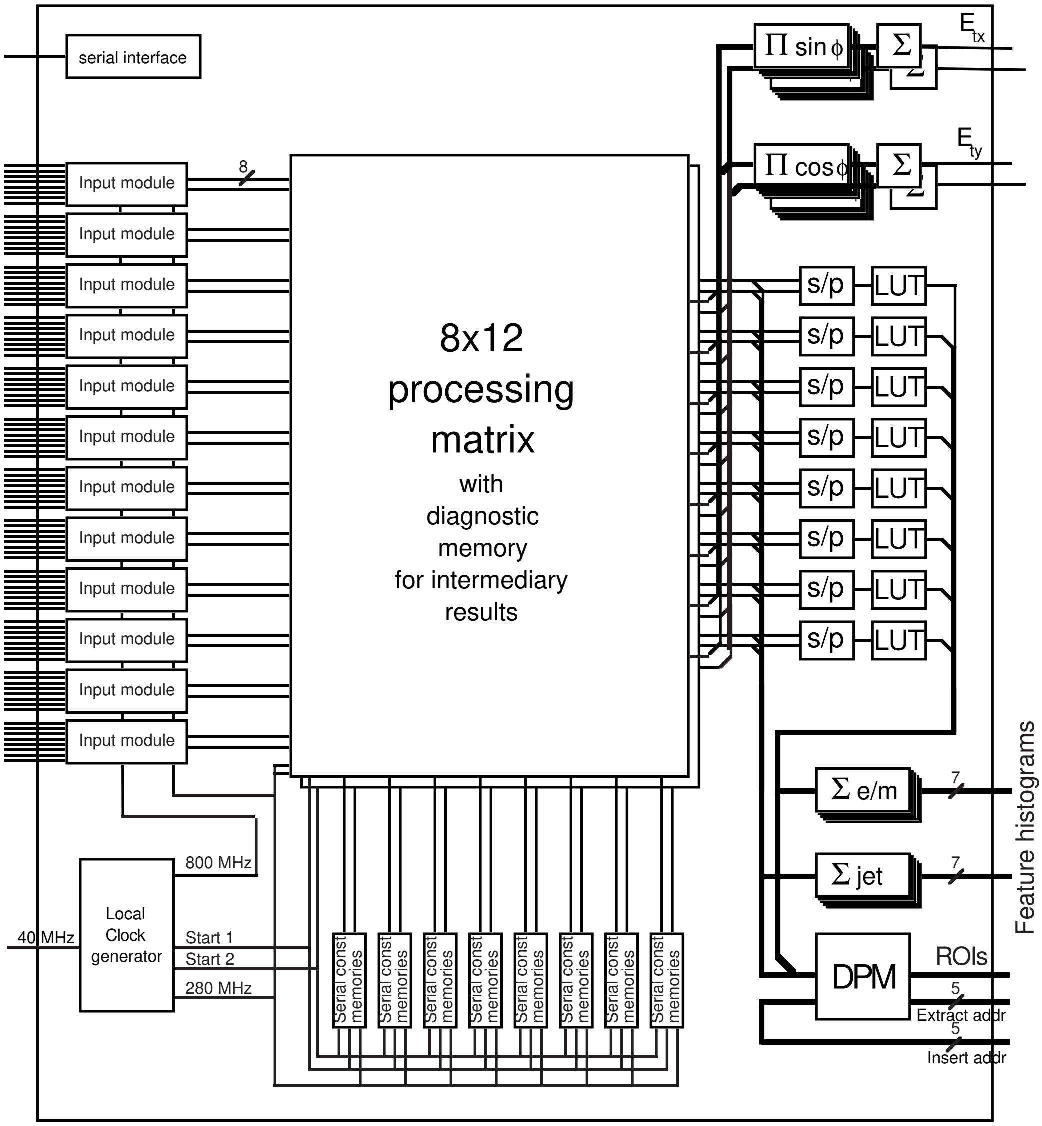

Figure 13. Processing ASIC.

The merger ASIC

All merging of bit-serial results outside the processing ASIC occurs in a merger ASIC (fig. 14) which can be used in two different modes: to implement two independent adder trees or two adder trees whose result are squared, added and finally compared with 7 programmable levels in order to produce a 3-bit result-code. The first mode is used for merging results on the processing and supervisor boards while the second mode will be used for calculating the total missing transverse energy. The overflow monitor supervices all additions so that overflow is properly reported. The design of this ASIC will be greatly simplified by the fact that it shares large portions of the processing ASIC layout. There is also a controller end of the serial link which terminates in an external parallel port. This port is meant as an access point for the diagnostic processor. This processor is implemented in conventional logic.

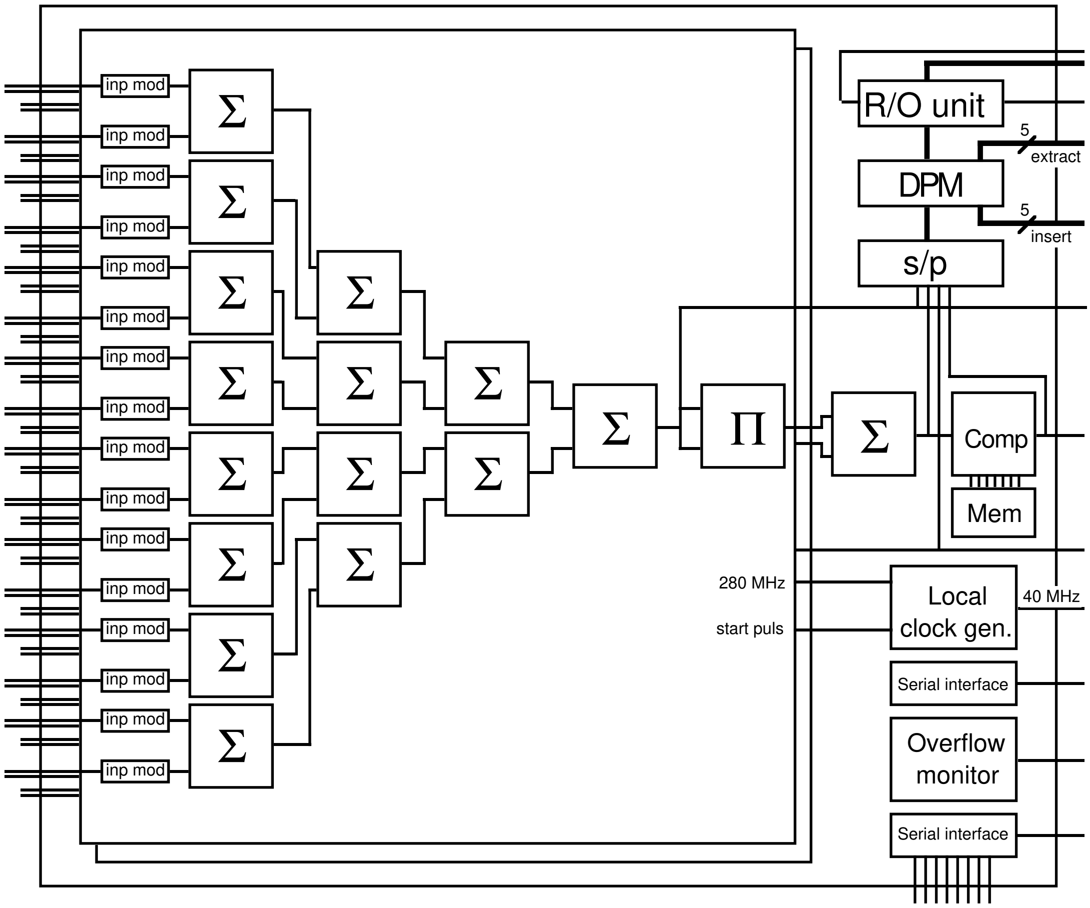

Figure 14. Merger ASIC.

The merger ASIC inputs will also be equipped with input modules (fig. 15) to take care of the re-synchronisation of data after transmission.

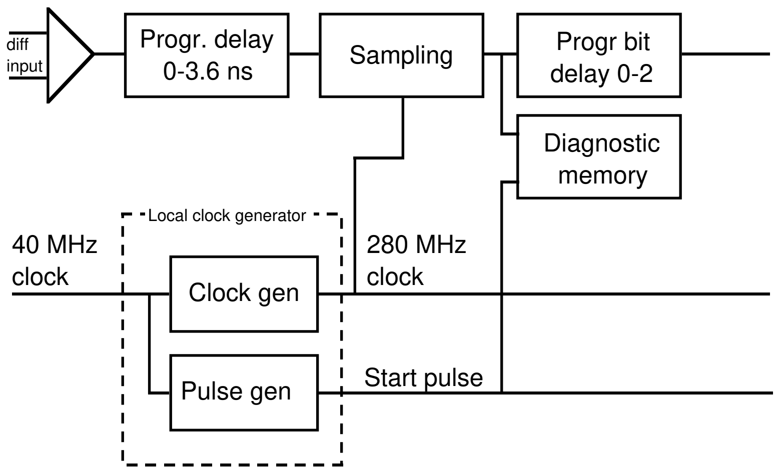

Figure 15. Input stage to merger ASIC. Only one local clock generator is required for each ASIC.

The supervisor board

The supervisor board contains 8 merger ASICs for merging serial feature count data and two for combining total energy components. The result is reported to the central first-level trigger.

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 [6] 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.

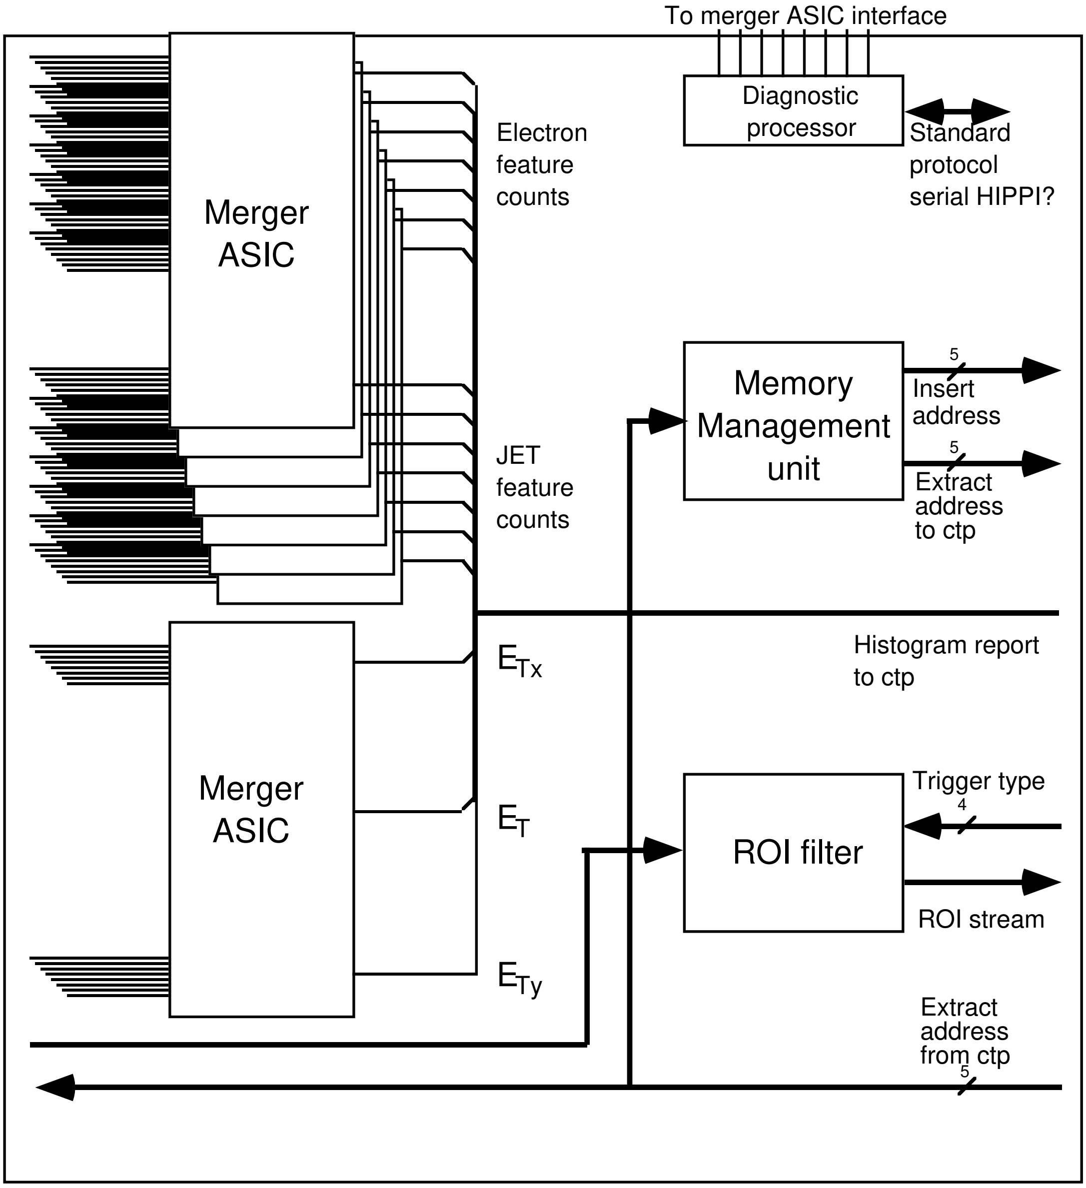

Figure 16. Supervisor board.

Trigger processor programming

The local diagnostic processor is responsible for programming all programmable registers and for scanning the diagnostic memories to verify proper system function. The diagnostic processor will communicate with the system via an internal high speed (280 Mbit/s) serial bus. The destination address, specified in the message header, is composed of two parts: the ASIC address and the internal address. The individual ASIC address will be defined by strapping dedicated input pins.

The trigger processor controller is responsible for interpreting external control information and translating it into local read and write operations. The controller will be built using a programmable gate array together with a local serial interface contained in the merger ASIC (fig. 14). The local diagnostic processor will, in turn, report to a remote diagnostic processor via a serial link.

Diagnostic facilities

The system will be monitored via spy registers capable of intercepting input, intermediate and final data. Input data are stored in the input modules, intermediate data in the processing matrix of the processing ASIC and the final data in the dual port memory. The storage will commence at the request of the controller with a suitable delay so that all results will correspond to the same input data. The diagnostic memory in the input modules will be extended to also serve as a possible source of data. The dual use of the DPM means that extensive storage cannot occur during normal operation since it would reduce the RoI buffer. The considerable silicon area required for memory implementation limits the size of the diagnostic memories, especially since the process chosen for this design is optimised for logic. However, if large diagnostic memories are required these may be realised as separate units using the same input modules and input layout but equipped with memories rather than logic. The ideal position for such a module would be on top of the processing ASIC, in a "piggy-back" configuration. This is clearly not possible for many reasons. Placing them on the reverse side of the board opposite to their processor ASIC partner is a possible solution.

The system will also include a diagnostic feature which runs the two processing matrices in parallel comparing their results to detect deviations. The possibility of using concurrent hardware detection of all operations via 3-code is under investigation.

THE DEMONSTRATOR PROJECT

The feasibility of the suggested design depends on some key technical issues, such as the possibility to bring a large number of .8 Gb/s fibres to an MCM, the reliability of bit-serial operations at the required clock rates and the performance of the opto/electric components. The process technology chosen for the processor is 0.5 μ BICMOS, with bipolar circuitry being used for clock generation and input stages. The processor itself will rely on CMOS. Until now, a number of smaller demonstration projects have already been initiated to prove the technical solutions. One such project addressed the question of multiple fibre optic inputs to MCMs. Other projects study high speed bit-serial operations and high frequency serial to parallel conversion.

At present a functional demonstrator is being designed. A detailed VHDL description down to a Register Transfer Level will serve as 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. However, the built-in diagnostic memories described in appendix B will be of smaller than the final system but of an adequate size for the test. 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. Although the prototype ASICs will be built in BiCMOS an alternative GaAs version of the trigger will also be investigated. The cost and possible advantages associated with GaAs [7] will be studied by simulating a GaAs version of the trigger. The fact that gate arrays are used rather than full custom design will make the simulation reliability sufficient for the purpose of evaluation.

Technical description

The aim of the demonstrator is thus to verify the feasibility of the compact processor design. This will be accomplished by designing and building a prototype to test the performance of the key components. These will be implemented to fit the full system as it is configured today. The diagnostic memories, however, will be kept large enough to be adequate for test purposes. but small enough to fit the design. More efficient implementations in the future will, due to technology scaling, release extra space for the memory. The demonstrator will thus be equipped with memory for storing 2 samples of input data, totalling 384 bytes and two bytes for storing intermediate results in the processing matrix (2x2x8x12=384 bytes). When working off-line the entire DPM can be used to store results, i.e. 32 slots, but under normal operation only a small portion is available.

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 processing ASIC as presented above and its testing with simulated data, recorded data and then finally on-line data. The prototype ASIC will be designed in BiCMOS but a GaAs alternative will be carefully investigated. The test will to the largest possible degree utilise facilities already developed for an alternative design [2]. 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 [5].The demonstrator board will be designed to be read out via VME in a way compatible with the ATLAS test-beam environment.

In paralell 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. Wether this stage will also include manufacturing and tests of the merger ASICs is largely a question of funding

CONCLUSION

A bit-serial first-level trigger processor has been presented after a careful system analysis based on the properties of available new technologies. This approach leads to a system partitioning, which better exploits the advantages of new technologies. It is more compact, more flexible, has a short latency, and should also be more reliable due to the smaller number of components. However, it is essential that the design principles are verified in extensive demonstrator tests.

The design uses the most advanced technologies available today. In addition the system has been designed to be able to profit from improvements in data transfer and data processing speed which are likely to occur in the near future. However, it is recognised that when designing and constructing the final system for ATLAS one will have to use technologies which, at the time, are well tested and mature.

ACKNOWLEDGEMENTS

The authors would like to acknowledge helpful discussions and useful suggestions from members of the RD-27 collaboration for development of first-level triggers for LHC detectors, of which the present project is a part.

The support from the Swedish national Science foundation, NFR, is also gratefully acknowledged.

REFERENCES

1. *ATLAS Technical proposa*l CERN (1994)
2. I.B. Brawn et.al, *The level-1 calorimeter trigger system for ATLAS Technical proposal* RD-27 CERN (1994)
3. N.Ellis, J.Garvey, CERN 90-10, vol3, p. 80 (1990).
4. M.J. Haney and G.D Gollin, *Cluster counting for E799/E832 at Fermilab*. Conf. report IEEE Nuclear Science Symposium, Orlando USA, October 1992.
5. RD-16 Status Report, CERN/DRDC 93-21, (1993)
   C. Bohm et. al. CERN *RD-27 note 6*, (1993).
6. G. Appelquist and C. Bohm, *A memory controller for FERMI*. FERMI note#5. CERN 1991
7 J Östberg et. al. *Design of a 1 GIPS peak performance processor using GaAs technology*. Submitted to High performance computer architectures 1994

APPENDIX A SUMMARY OF SYSTEM PROPERTIES

The compact calorimeter trigger operates on merged electromagnetic and hadronic calorimeter data from a mesh of 64x64 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 environments 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. 8 independent classes are foreseen. The global occupancy of these classes is counted and reported to the central first-level trigger. 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 4x4 kernel to reduce statistical effects followed by a 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 7 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 8 programmable ranges, i.e. compressed into a 3-bit code for transfer to the central first-level trigger. 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.

APPENDIX B TABLES

Latency budget:

Latency along the different calculation paths (clock cycles/frequency):

INPUT STAGE	
Input	1/800
prog delay	1/800
sampling	1/800
allign 400	
prog delay	1/400
pulse flag decode	1/400
allign 280	
total	3/280
bit-serial fifo	1/280
total	4/280

ELECTROMAGNETIC	
build sums	4/280
classify size	12/280
allign 280	
total	3/40
max,min	1/40
look-up	1/40
total	35/280
count classifications	5/280
transfer	6/280
merge	4/280
transfer	6/280
merge	4/280
allign 40	
total	8/40

ET	
build sums	3/280
add	1/280
x sine	12/280
transfer	6/280
merge	4/280
transfer	6/280
merge	4/280
square	12/280
classify	12/280
allign 40	
total	10/40

JET	
build sums	4/280
add	1/280
classify size	12/280
decluster	7/280
count classifications	5/280
transfer	6/280
merge	4/280
transfer	6/280
merge	4/280
allign 40	
total	8/40

Interconnections

Entity	Nr	Rate	Data frame	Type
From FE to CalTP:				
Trigger cell data	4096	800 Mbit/s	2x10b/25ns	optical
From CalTP to CTP:				
E/m multiplicities	7	280 Mbit/s	6b/25ns	diff
Jet multiplicities	7	280 Mbit/s	6b/25ns	diff
E_T	1	280 Mbit/s	3b/25ns	diff
Insert address	5	40 MHz		
Extract address	5	≈100 kHz		
From L2TP to CalTP				
Strobe	1	≈100 kHz		
Extract address	5	≈100 kHz		
Trigger type	4	≈100 kHz		
From CalTP to L2TP (RoIs)				
RoI data	8	?		
RoI strobe	1	?		

From FE to processor board				
Trigger data	512	800 Mbit/s	10b/25ns	fibers
From processor to supervisor board				
E/m multiplicities	7	280 Mbit/s	6b/25ns	diff
Jet multiplicities	7	280 Mbit/s	6b/25ns	diff
E_T	4	280 Mbit/s	14b/50ns	diff
Insert address	5	40 MHz		
RoI data	1	280 Mbit/s	?	diff
From supervisor to processor board				
Extract address	5	≈100 kHz		
Internal serial bus	2	280 MHz	?	diff
Local clock	1	40 MHz		
From system to supervisor board				
System clock	1	40 MHz		?
From supervisor board to control				
External serial bus	1	40 MHz	?	?

Registers

Type	Nr	bits
Processor ASICx128		
Analogue delay	96	8
Digital delay	192	5
Diagnostic memory	384	8
Intermediate data	484	8
Cluster levels	96x7	12
Isolation levels	96x3	12
Leakage levels	96x3	12
Mode register	?	8
Status registers	?	8
LUTs	32x256	8
DPM	32	32x8+45x4
Sine & Cosine	2	12
Merger ASICx(128+8)		
Analogue delay	32	8
Digital delay	32	3
Diagnostic memory	32	8
E-levels	7	12
DPM	32	39

Timetable

<u>Processor and merge ASICs</u>		
high level specification	Dec.	94
layout	June	95
production	Okt.	95
possible redesign	Dec.	95
Test of processor module with simulated data	fall	95
Beam test of processor module	winter	96
<u>O/e MCM</u>		
design	fall	94
preproduction and test	spring	95
redesign	mid	95
production	fall	95
test of o/e mod	Dec.	95
Test of partial system (processor ASIC & o/e)	mid	96