

FIRST-LEVEL TRIGGER SYSTEMS FOR LHC EXPERIMENTS

I. Brawn, R.E. Carney, Y. Ermoline¹, J. Garvey, D. Grant, P. Jovanovic, I. McGill,
R. Harris, R. Staley, A. Watson

School of Physics and Space Research, University of Birmingham, UK

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

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

P. Böd, 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, G. Thompson
Queen Mary and Westfield College, University of London, UK

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

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

E. Gennari, A. Nisati, E. Petrolò, M. Torelli, S. Veneziano, L. Zanello
Intituto Nazionale di Fisica Nucleare and Universita La Sapienza, Rome, Italy

G. Appelquist, C. Bohm, S. Hellman, B. Hovander, N. Yamdagni, X. Zhao
Stockholm University, Sweden

(*) Spokesman

¹ Visitor from Institute for High Energy Physics, Russia, 142284, Moscow region, Protvino.

Summary

The RD27 project was approved in June 1992 with the following milestones:

- design outline of a first level trigger system
- detailed design studies and beam tests of a prototype calorimeter trigger
- detailed design studies of a muon trigger, including several detector options

Design studies have been performed for all parts of the level-1 trigger system, which consists of subtrigger processors associated with the calorimeters and muon detectors, a central trigger processor that combines the subtrigger results and makes the overall *yes/no* decision, and a trigger, timing and control distribution system which distributes the trigger decision to the front-end systems. Our design studies also include the interface to the level-2 trigger which makes use of *region-of-interest* information provided by the level-1 system. Our design studies are described briefly in the main body of this status report and in more detail in RD27 internal notes.

Detailed design studies have been made for the calorimeter trigger processor, where we are investigating two architectures based on bit-parallel and bit-serial digital processors respectively. In the case of the bit-parallel approach, we have constructed a first prototype processor and successfully tested it in real time at test beams in conjunction with the Accordion and TGT liquid-argon calorimeters. This prototype processor is based on an Application-Specific Integrated Circuit (ASIC) which was produced for the project. The beam tests were carried out at the full LHC clock speed of 40 MHz; tests were also made at the originally envisaged frequency of 67 MHz.

Our calorimeter trigger design studies result in rather compact processor systems occupying only a few electronic crates. We believe that such processors could be constructed based on technology that is currently available. However, more work is required to validate these designs, particularly in relation to data transmission into the processor where the total data rate is several hundred Gbytes per second. We plan to build further technical demonstrators to address this issue.

Extensive simulation work has been performed for the calorimeter trigger. The algorithms that are included in the processors were selected on the basis of these studies. Our efficiency and rate calculations were used in the letters of intent of ATLAS and CMS. Where comparison is possible, our predictions have been found to agree well with our test-beam measurements.

A detailed design study has been performed for a muon trigger based on Resistive Plate Chamber (RPC) detectors and a first demonstrator system is under construction in preparation for tests with RD5 later this year using large-area RPCs. The initial prototype is based on coincidence matrices made from expensive commercial GaAs products. We hope to develop a second prototype based on custom programmable coincidence-array circuits with LHC performance.

We are actively working on simulation studies to evaluate the physics requirements and the backgrounds for the muon trigger. The implications of the high rate of “random” hits in the muon detectors, due to the enormous flux of low-energy photons and neutrons in the experimental areas, are still being evaluated.

The possibility of making a level-1 muon trigger using the precision drift-chamber detectors is being studied. Conceptual design studies for various chamber geometries have been performed, taking into account the need to trigger at relatively low p_T ($p_T \approx 6$ GeV) for B physics analyses. We have also performed an initial investigation of the feasibility of bunch-crossing identification logic based on programmable coincidence arrays, extending ideas used by members of our collaboration in the forward-muon trigger of the H1 experiment at HERA.

Initial design studies have been made for the central trigger processor. This receives information from the subtrigger processors, allowing for differences in latency, and makes the overall level-1 decision. Facilities to prescale high-rate triggers are included, and extensive monitoring is provided with the provision of scalers to record trigger rates. We plan to make more detailed design studies and to perform modelling using VHDL. Several key components – a variable length pipeline, a scaler and a prescaler – will be studied in more detail, initially using a field-programmable gate array test bench.

In conjunction with RD12, we have studied a range of technologies for the optical transmission of timing, trigger and control signals from a single node at the central trigger processor to large numbers of front-end electronics locations. Using currently available components, we have demonstrated the capability of time-division multiplexing the required signals and broadcasting them over an entirely passive all-glass distribution network to 1024 receivers per laser transmitter. The recovered clock jitter is less than 100 ps rms, compatible with foreseen LHC experiment requirements. We plan to pursue this development using such higher-power lasers and new optoelectronic technologies as become available in the future.

In the next phase, in conjunction with RD12, RD16, CERN Microelectronics Group and interested users, we plan to develop a demonstrator timing receiver ASIC which could be incorporated in the front-end electronics of detector prototypes in future beam tests. In addition to the clock, the receiver would deliver a unique bunch crossing number synchronously with each level-1 trigger decision, with programmable compensation for detector, electronics, time-of-flight and signal propagation delays. The receiver ASIC would also be able to transmit to the front-end electronics separately deskewed broadcast and individually-addressed commands and data.

Initial studies have been made for the level-1/level-2 interface system, including first timing measurements. This work will be extended in the second year of the project.

A list of suggested milestones for the second year of the RD27 project and the funding request to CERN can be found in Sections 8 and 10 of this status report.

1. Introduction

The RD27 project is a broad study of level-1 trigger systems for LHC experiments. The work includes level-1 subtrigger processors based on calorimetry and muon detectors, the central trigger processor that combines the results from different subtriggers, the timing, trigger and control distribution system that delivers the trigger decision to the front-end electronics, and the interface to the level-2 trigger. Monte Carlo simulation plays an essential role in our studies, for algorithm optimisation (rate and efficiency calculations for background and signal processes) and for comparison with test-beam data (single particle simulation).

The challenges of triggering in LHC general-purpose experiments are formidable, especially at level-1 [1]. At high luminosity, the proton–proton interaction rate will be $\sim 10^9$ Hz, with a bunch-crossing rate of 40 MHz. The readout time for subdetectors is typically 10 μ s, limiting the rate into the level-2 system to 100 kHz. Hence, a very large rejection power is required from the level-1 trigger. The cross-sections for many of the physics processes that are sought at LHC are very small, so high trigger efficiency is mandatory, in some cases with rather complex signatures. Furthermore, LHC trigger systems should be flexible, allowing one to react to new physics or to unforeseen background conditions.

Issues that affect the design of the level-1 trigger system are the following:

- The readout of LHC detectors is pipelined because of the short (25 ns) bunch-crossing period. The length of these readout pipelines, which may be analogue or digital, depends on the level-1 trigger latency which should be as short as possible.
- The trigger must uniquely identify the bunch crossing containing the interaction of interest. This is especially challenging in the case of triggers based on detectors with a signal collection time longer than 25 ns.
- The calorimeter trigger must be able to operate at high luminosity in the presence of pileup from many overlapping events.
- Muon detectors will be exposed to a high flux of low-energy photons and neutrons which can produce “noise” hits; the muon trigger must be able to operate under these conditions.
- The physical size of LHC detectors and their large numbers of readout channels require a complex system for establishing and maintaining synchronisation, and for distributing the trigger decision to the front-end electronics which will, in most cases, be on the detector.
- Information from the level-1 trigger can be used to guide the level-2 trigger, for example by indicating the location of electron/photon candidates.
- It is essential to foresee extensive test and monitoring facilities, especially in the case of on-detector electronics.
- Facilities for maintaining the calibration and synchronisation of the trigger processors must be included in the design.
- Parameters such as transverse momentum (p_T) thresholds should be programmable, requiring a control system.

The design studies and demonstrator prototypes being developed in RD27 take into account all the above issues.

The RD27 collaboration maintains close contact with other R&D groups and with both ATLAS and CMS. We collaborate informally with RD11 (EAST) in the study of the interface to the level-2 trigger. Work on the timing, trigger and control distribution system is performed jointly with RD12. For the calorimeter trigger work, we have made beam tests with the RD3 Accordion calorimeter and with the RD33 TGT calorimeter; close contact is maintained with RD16 (FERMI). For the muon trigger, beam tests will be performed with RD5. Members of RD27 participate strongly in trigger-related activities within ATLAS. Members of CMS attend RD27 meetings as observers and RD27 work is presented to the CMS collaboration on request.

The RD27 project was approved in June 1992 [2] with the following milestones:

- design outline of a first level trigger system
- detailed design studies and beam tests of a prototype calorimeter trigger
- detailed design studies of a muon trigger, including several detector options

In this report we describe the work that we have done so far, and propose a continued programme of work for the coming year.

Outline of Level-1 Trigger System

In Figure 1 we show a block diagram of a level-1 trigger system based on calorimetry (electron, photon, jet, missing transverse energy triggers) and on muon detectors. The subtrigger processors work independently and in parallel to summarise the characteristics of the event:

- information on multiplicity of electrons/photons for several thresholds with various isolation requirements
- information on multiplicity of jets for several thresholds
- information on missing transverse energy for several thresholds
- information on multiplicity of muons for several thresholds

These data are transmitted to the central level-1 trigger processor that makes the overall *yes/no* trigger decision, based on combinations of trigger requirements, with the possibility of prescaling high-rate triggers. The central trigger processor is in turn connected to the timing, trigger and control distribution system that broadcasts the trigger decision to the front-end electronics, much of which is on the detector.

Region-of-interest information from the level-1 subtrigger processors is sent to the level-2 trigger system, which only needs to acquire and process data from regions of the detector in the vicinity of (electron, photon, jet or muon) candidates flagged by the level-1 system.

The trigger system is based on purpose-built digital hardware processors and makes extensive use of micro-electronics. The processing is pipelined and many operations are performed in parallel. Such processors have a fixed set of algorithms, but are programmable at the level of parameters. In practice considerable flexibility can be provided. The trigger processors themselves are synchronous with the LHC clock. However, we are considering using asynchronous transmission of zero-suppressed data to reduce bandwidth requirements where necessary, notably at the input to the calorimeter trigger processor.

In the following sections, we describe the functional blocks of Figure 1 in detail and comment on how the different functional blocks communicate with each other.

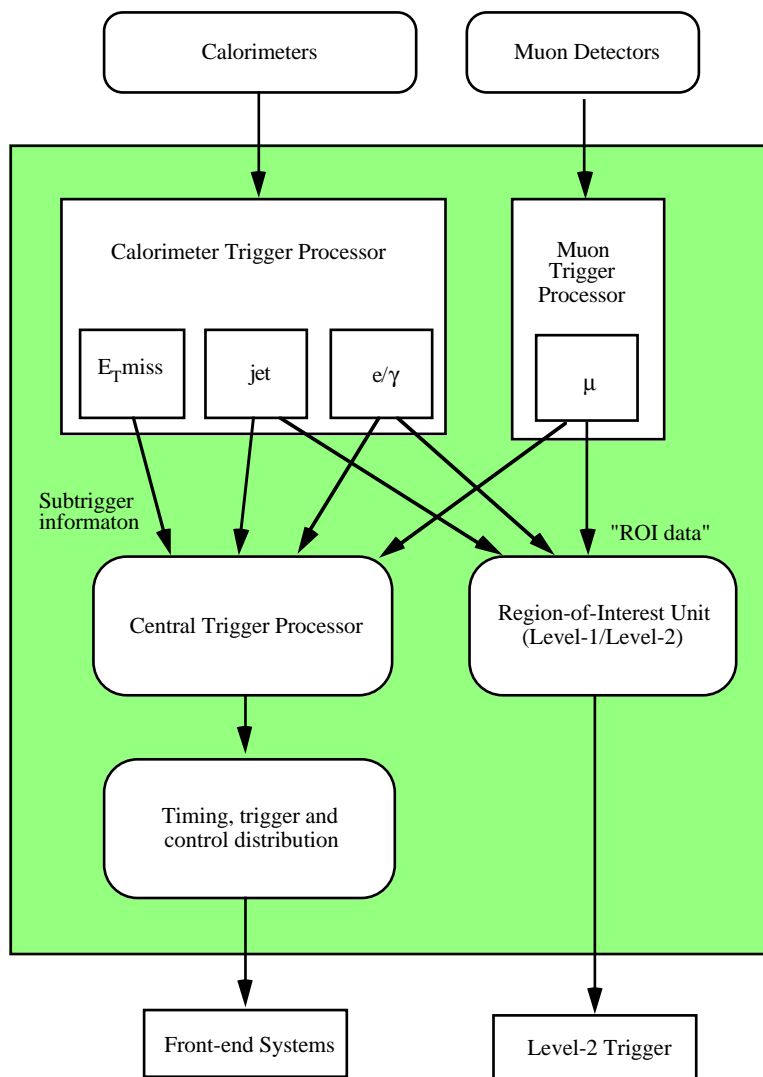


Figure 1: Block diagram of level-1 trigger system.

2. Calorimeter Trigger

Introduction

Several possibilities have been suggested for level-1 calorimeter triggering at future hadron colliders. ATLAS [3], CMS [4] and SDC [5] propose to use purpose-built digital processors, while GEM [6] considers a mixed analogue/digital system as its baseline option, with a programmable digital processor as an alternative.

In RD27 we are studying purpose-built, synchronous digital processors for the level-1 calorimeter trigger. Strong points of this approach are the following:

- It is possible to implement relatively complicated algorithms that optimise the trigger performance in terms of efficiency and rejection power. Many different sets of thresholds can be provided, and (optional) criteria such as electron/photon isolation can be implemented.

- A large degree of flexibility can be provided through programmable parameters (thresholds and control words).
- A programmable threshold can be applied at the cell level, suppressing contributions from electronic noise and pileup. The threshold can be increased for high-luminosity running (more pileup) or in case of unforeseen problems such as coherent noise.
- Ease of calibration – the trigger ADC system (if separate) can be cross-calibrated against the precision readout.
- Ease of monitoring – intermediate results from the calculation can be read out and checked against values calculated from the input data, requiring an exact match.
- Ease of testing and diagnostics – test patterns can be played through the processor, localising faults.
- The processing latency can be minimised using hardwired algorithms and fixed routing of the data by direct links.
- The interface between the digitisation system and the trigger processor can be optimised to cope with the extremely high bandwidth, for example by exploiting the low detector occupancy (using zero-suppression).
- It is relatively easy to equalise the phase of digital signals relative to one another, for example using programmable delays or FIFOs.
- Extensive use of Application Specific Integrated Circuits (ASICs) allows cost-effective implementations. Commercial CAD tools provide an easy route to a wide range of products from different vendors. Hence, one can exploit new technology as it becomes available and select among manufacturers to obtain the best price. Our preliminary studies suggest that ASIC products that are already available are adequate for our needs; we can exploit future developments to add functionality or reduce the system cost.

We are following two calorimeter trigger design schemes in RD27, based respectively on bit-parallel and bit-serial operations (addition, comparison, etc.), as described below. These design studies are complementary and substantial “cross-fertilisation” has already occurred. They could be combined in the future into a hybrid design; many features are common to the two designs as described below.

The calorimeter trigger processor acts on reduced granularity data from the electromagnetic and hadronic calorimeters (ECAL and HCAL), with a typical trigger cell size of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in pseudorapidity – azimuth space, and with just one sampling in each of the ECAL and HCAL. Simulation studies (described below) suggest that 8-bit or 9-bit linear data with a least count of about $E_T = 1$ GeV are adequate for the level-1 trigger calculation. These data can be obtained from an independent trigger ADC system, using analogue summation over trigger cells before the ADC, and a look-up table after the ADC to apply calibration constants. Alternatively, if a digital front-end system such as FERMI [7] is adopted for the calorimeter readout, the data for the level-1 trigger can be derived from the precision ADCs. Beam tests performed so far by RD27 have used dedicated trigger ADC systems. However, we are in close contact with the FERMI collaboration, who have included in their design a level-1 trigger sum that could be interfaced to either the bit-parallel or the bit-serial processor design.

The signals from calorimeters proposed for LHC experiments, such as the Accordion liquid argon calorimeter, are slow compared to the 25 ns bunch-crossing interval. Hence, if the

raw digitised signals are used, large energy deposits would be seen by the trigger processor for several consecutive bunch-crossings. This problem can be overcome by using signal processing techniques, which can be integrated with the ADC system. We have studied simple digital filter algorithms, applying them to data collected with the Accordion calorimeter. Analogue signal processing, as demonstrated by GEM [6], is an alternative under consideration.

The transmission of data into the trigger processor system is a major issue. For a pseudorapidity coverage of $|\eta| < 3$ and a trigger cell size of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in each of the ECAL and HCAL, one has ≈ 8000 trigger cells. Assuming 8-bit data with a sampling rate of 40 MHz, the total data rate is 320 Gbytes/s. Furthermore, the two-dimensional nature of cluster-finding algorithms requires that each processing element knows about its environment. Hence one has to fan-out data between processing elements, increasing the overall bandwidth. (While one could instead make use of additional connections between processing elements, we do not find this alternative attractive.) During the next year, RD27 plans to address this difficult problem, verifying our detailed design studies using demonstrator prototypes as described in detail below.

We have already shown that complex trigger algorithms can be implemented on ASICs. A first demonstrator trigger processor for isolated electrons has been built and tested with the Accordion and TGT liquid-argon calorimeters at test beams. Building on this experience, we have made detailed design studies for a bit-parallel and a bit-serial trigger processor system.

The rest of this section is organised as follows. Physics simulation studies, which were used to optimise and evaluate trigger algorithms, are presented¹. Next, we explain the differences between the bit-parallel and bit-serial processing architectures. We then describe the bit-parallel trigger processor demonstrator, the beam tests performed with RD3 (Accordion) and with RD33 (TGT), and present results of off-line analysis of the test beam data. Finally, we present the two detailed design studies that have been performed for calorimeter trigger processors, using bit-parallel and the bit-serial architectures, and propose a programme of work for the coming year.

Algorithms and simulation studies

The essential requirements for a first-level trigger system are that it should reduce the event rate in an LHC detector from the bunch-crossing frequency of 40 MHz with about 25 events per bunch crossing to a rate that can be accepted by the next stage in the trigger chain (at most ~ 100 kHz), while maintaining high efficiency for interesting physics processes. Physics simulations, that is the simulation of possible LHC events and detector response, are needed to determine whether or not a proposed trigger algorithm can meet these goals. Such simulations are used to optimise algorithms and their parameters, and to determine requirements such as the dynamic range for calorimeter signals used in the trigger. Members of our collaboration have performed extensive simulation studies that build on earlier work performed for the Aachen workshop [8] and within the LHC experimental collaborations [9]. Here we briefly describe a few examples of simulations performed, mainly in relation to the electron/photon trigger. A more detailed description of this work, which includes studies of jet and missing E_T triggers, can be found in Ref. [10].

¹Our predictions for trigger rates were used by both ATLAS and CMS in their letters of intent.

Physics processes simulated

For the electromagnetic (e.m.) cluster trigger, the efficiency of the algorithms was studied using events containing single electrons of different energies and rapidities. Minimum-bias events, generated using PYTHIA, were superimposed to simulate the effects of pileup. Real events are more complex than this, containing in addition to the electrons or photons an underlying event from the residual proton–proton system and possibly also hadronic jets. In order to study the performance of the e.m. cluster trigger in more realistic circumstances three examples of possible physics processes were simulated using PYTHIA:

- $H^0 \rightarrow \gamma\gamma$, $m_H = 120$ GeV
- $H^0 \rightarrow ZZ^* \rightarrow e^+e^-e^+e^-$, $m_H = 130$ GeV and 150 GeV
- Top events where the top quarks decayed to electrons

The rate for the first-level e.m. cluster trigger will be dominated by backgrounds from high- p_T jet events. Jets may mimic the signature for electron or photon production if they contain one or more high- p_T π^0 s or an early-showering high- p_T hadron. In order to study the background rates for this trigger, a large-statistics sample of jet events was simulated using PYTHIA.

Detector simulation

A number of detector models of different degrees of complexity and realism were used in these studies. The choice of model was determined by the degree of detail needed to adequately answer the question posed and by the computer resources available.

For many of the studies, simple models based on parameterisations of detector response were used. The four-vectors of final-state particles were smeared with the resolutions of a “typical” LHC detector, namely $\Delta E/E = 10\%/\sqrt{E}$ for electromagnetic particles and $\Delta E/E = 50\%/\sqrt{E}$ for hadrons. Longitudinal sharing of hadronic showers between e.m. and hadronic calorimeters was parameterised; no transverse shower width was simulated. Electronic effects and event pileup were included in the simulation. Such models were used for studies of jet trigger rates, the required ADC range in jet and missing- E_T triggers, and for fast studies of many effects in the electron/photon trigger.

More complex detector simulation was also used, based on the GEANT simulation package [11]; the detector model was that of ATLAS [12]. This simulated in detail the material in the tracking system, magnet and cryostat as well as the calorimeters², and also the effects of magnetic fields. This model was used for all efficiency studies in the electron/photon trigger, where these effects are most important. We have so far concentrated on the trigger for $|\eta| \approx 0$, but studies of the variation in jet rejection with angle are planned for the coming year.

E.M. cluster algorithms

In all of the algorithms studied here, it is assumed that the calorimeter trigger will work with a reduced detector granularity of $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ with one depth sampling in each of the ECAL and HCAL. This granularity was chosen on the basis of simulation studies. A threshold is applied to the E_T in each trigger cell and only the cells above this threshold

² To economise on computer time, a somewhat simplified calorimeter model was used in the simulation of the jet background to the electron/photon trigger.

contribute to the trigger. This reduces the sensitivity of the trigger to electronic noise and pileup.

The electron/photon trigger is based upon a localised deposit of energy in the ECAL. This cluster should ideally contain fully the energy of an e.m. shower while being small compared with the size of a jet core. Several cluster definitions have been studied (see below) to find the best compromise between sharpness of threshold and background trigger rate. Many of the physics processes sought at the LHC result in isolated leptons or photons, i.e. not associated with jets, while the dominant background to the electron/photon trigger comes from jets containing high- p_T π^0 s or early-showering hadrons. Thus a further trigger rate reduction can be achieved using isolation requirements.

Electromagnetic cluster definition

The figures of merit for the electromagnetic cluster algorithm are the sharpness of the trigger threshold and the trigger rate for a threshold which is efficient for the desired physics. The sharpness of the threshold is affected by the resolution of the calorimeter and the digitisation system for the level-1 trigger, the degree of containment of e.m. showers within the cluster, and by the pileup and electronic noise in the cluster window. Rate is primarily determined by the area of the cluster window. Three cluster definitions were considered:

- A threshold was applied to the E_T in a single e.m. trigger cell
- The E_T sums of pairs of trigger cells adjacent in azimuth or rapidity were compared with a threshold (Figure 2) – this is selected as our standard algorithm
- The summed E_T of a 2×2 trigger cell cluster was compared with a threshold

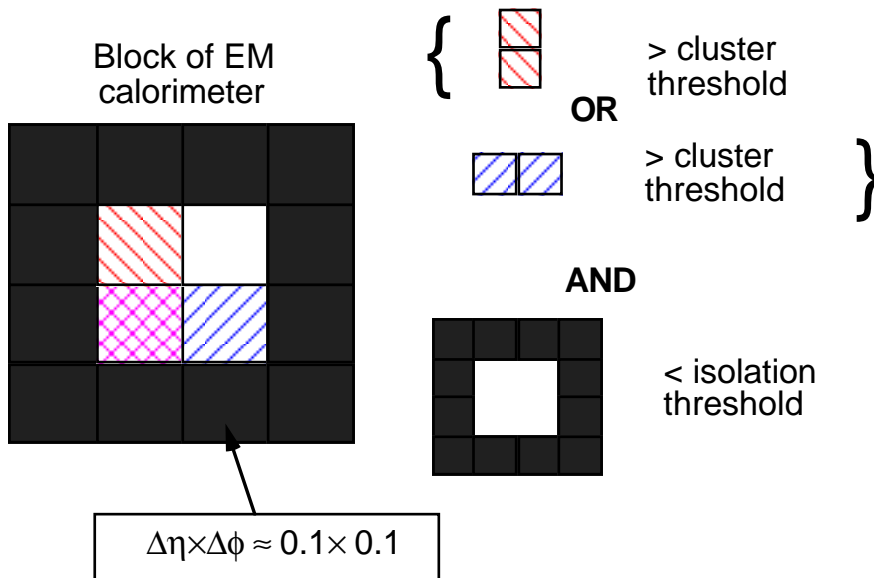


Figure 2: EM cluster algorithm.

The efficiency as a function of threshold for these three cluster definitions is shown in Figure 3 for 50 GeV electrons at $|\eta| \approx 0$. The single-cell algorithm exhibits a significantly softer threshold than the two-cell and four-cell clusters due to shower leakage at the edges of the cell. For this granularity the two-cell cluster threshold is almost as sharp as that of the four-cell cluster. These simulation results have been confirmed by test-beam measurements. The corresponding trigger rates are shown in Figure 4 as a function of threshold. Since the

two-cell algorithm gives a much sharper threshold than the single-cell one, and has a significantly lower rate than the four-cell one, it is taken as our standard algorithm. The question of the optimum trigger algorithm for large rapidity requires further study, but it is worth noting that the single-cell algorithm is particularly badly affected by the small physical size of the trigger cells at large rapidity.

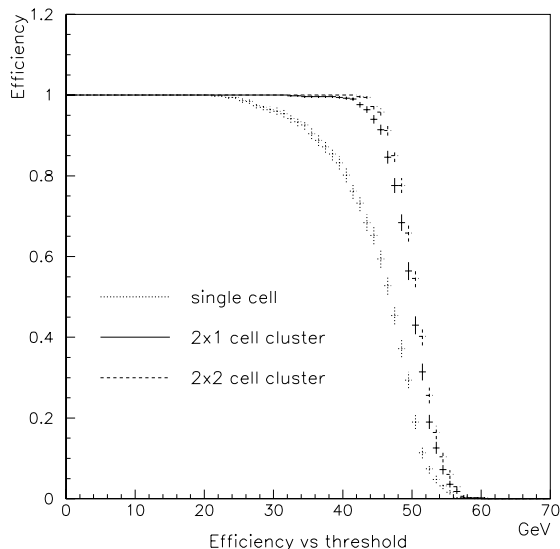


Figure 3: Efficiency for 50 GeV electrons at $|\eta| \approx 0$ versus cluster threshold for various cluster algorithms.

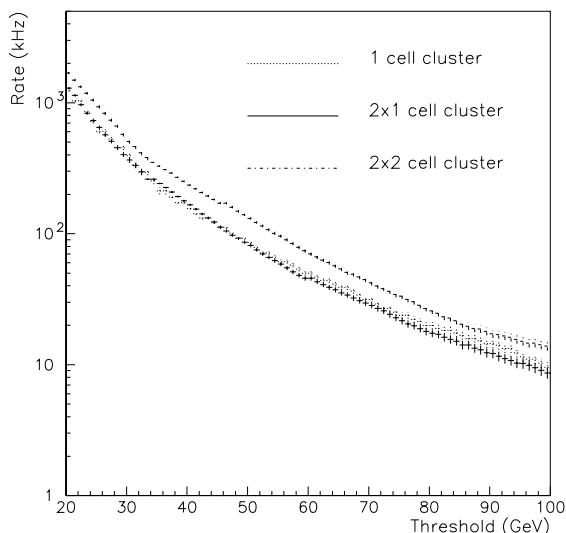


Figure 4: Trigger rate versus cluster threshold for various cluster algorithms for $|\eta| \approx 0$. Here the “threshold” corresponds to 95% efficiency for electrons.

Isolation requirements

The E_T sums in the ring of e.m. cells surrounding the cluster and in the hadronic window behind it provide extra rejection against jet backgrounds. The effectiveness of the isolation cuts is determined by the area of the isolation window and by the tightness of the cuts which can be applied while retaining high efficiency for isolated e.m. showers. The latter is

limited by electronic noise, pileup and by leakage of the shower into the isolation regions. Electronic noise and pileup are suppressed by applying an E_T threshold to the individual trigger cells, but this also reduces the sensitivity to low- p_T jet fragments and hence the effectiveness of the isolation requirement. This trigger cell threshold parameter is crucial to the performance of the isolation veto.

Table 1 shows the mean and rms E_T values for different-sized regions in the calorimeter due to electronic noise and two levels of pileup. The level of electronic noise used corresponds to 400 MeV energy in a trigger cell, which is approximately the level seen in tests with RD3 [13]. The effect of the cell E_T threshold in suppressing these noise contributions is shown in Table 2. In most of the studies described here a threshold of 1 GeV per trigger cell was applied.

In addition to inclusive e.m. cluster triggers, it is foreseen to trigger on pairs of electrons and photons at lower thresholds. Here it is reasonable to use less restrictive isolation requirements, since both the rejection power of the isolation cut and the loss in efficiency due to it are squared. We have checked, for example, that high (>95%) isolation-cut efficiency for $H^0 \rightarrow \gamma\gamma$ can be retained while obtaining an order-of-magnitude reduction in trigger rate, even in the presence of pileup at the highest LHC luminosity.

Estimated trigger rates

We have estimated trigger rates for the most difficult case, where the peak LHC luminosity is taken as $1.7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. We aim for an inclusive e.m. cluster threshold of ≈ 40 GeV and a trigger on pairs of clusters with thresholds of ≈ 20 GeV, as suggested in the ATLAS letter of intent [3]. The e.m. cluster trigger rates obtained are shown in Figure 5 for single-cluster and two-cluster triggers, with and without isolation requirements. It is worth noting that the isolation requirement reduces the background rate by an order of magnitude, while giving excellent efficiency for isolated electrons.

Noise Sources	Single em trigger cell	two-cell em cluster	em isolation ring of 12 em trigger cells	hadronic isolation 4×4 hadronic trigger cells
Electronic noise only	0.32 ± 0.24	0.51 ± 0.36	1.12 ± 0.95	1.23 ± 0.93
Plus 19 events pileup	0.36 ± 0.30	0.60 ± 0.46	1.49 ± 1.15	1.30 ± 1.01
Plus 38 events pileup	0.41 ± 0.35	0.68 ± 0.54	1.90 ± 1.40	1.34 ± 1.04

Table 1: Mean and rms noise levels (transverse energy in GeV) in different-sized regions of the calorimeters. The results are for $|\eta| \approx 0$ which gives the worst-case contribution from electronic noise.

Noise sources	No threshold	0.5 GeV	1.0 GeV	1.5 GeV
Electronic noise only	1.68 ± 1.26	1.98 ± 0.99	0.35 ± 0.54	0.06 ± 0.12
Plus 19 events pileup	2.10 ± 1.58	2.32 ± 1.22	0.56 ± 0.85	0.15 ± 0.50
Plus 38 events pileup	2.51 ± 1.84	2.66 ± 1.37	0.76 ± 1.03	0.24 ± 0.68

Table 2: Mean and rms noise levels (E_T GeV) in the e.m. plus hadronic isolation region as a function of the trigger cell threshold.




Figure 5: Trigger rate for $L = 1.7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ versus threshold with (lower curves) and without (upper curves) isolation. (a) Single e.m. cluster rates, (b) Two e.m. cluster rates.

ADC resolution and dynamic range

Together with the trigger cell granularity and occupancy, the resolution and dynamic range of the ADC system affect the input bandwidth to the first-level trigger. The requirement on the resolution is that it should not degrade the sharpness of the trigger thresholds significantly. For a threshold of 20 GeV a calorimeter with a “typical” performance of $\Delta E/E = 10\%/\sqrt{E} \oplus 1\%$ would have a resolution of 0.5 GeV. However, many currently-suggested designs include a preshower detector for γ/π^0 separation, and if the energy from the preshower detector is not included in the trigger this resolution is degraded by about a factor of two. Thus a least count of $E_T \approx 1$ GeV is well matched to the physics requirements.

A finer resolution is desirable in setting the trigger cell threshold. This is determined by a balance between suppressing noise in the electron/photon isolation region and retaining sensitivity of the isolation test to jet fragments. Since the contributions from electronic noise and pileup are typically in the region of 0.25–0.4 GeV per trigger cell, the ability to tune the cell threshold with a step size of 0.25–0.5 GeV is desirable. There is, however, no need for these least significant bits to be transmitted to the trigger processors.

The dynamic range requirements are determined by the resolution for high- E_T jets and missing E_T . Comparing with a system with infinite dynamic range, it is found that a maximum count of 255 GeV/cell leads to a slight softening (few percent loss of events) at jet or missing- E_T thresholds of 300 GeV, becoming significant (20% loss) at 400 GeV. A maximum count of 511 GeV/cell was found to be practically indistinguishable from infinite range.

Trigger cell occupancy

The trigger-cell occupancy has been studied using a detailed GEANT simulation of minimum-bias events. The number of events per beam crossing was conservatively taken to be 38. It was assumed that bunch-crossing identification will not be possible for low-energy signals, and signals from successive bunch crossings were added with weights 0.5, 1.0 and 0.5 respectively, corresponding approximately to the shape of pulses observed in beam tests. Electronic noise was added with a gaussian distribution of 0.35 GeV rms. Under these very conservative assumptions, the occupancy for an E_T threshold of 1 GeV was found to be 6%.

Hardware architectures

Systolic arrays

The most efficient way to implement synchronous level-1 triggers with fixed algorithms is by using systolic arrays (processing pipelines). Here the hardware is tailored to the algorithm in question, which is broken up into its most basic arithmetic and logical components. These are allocated to compute stations which are connected via data paths and delays so that all arguments arrive in the same clock period. The clock frequency must be chosen so that all basic calculations can finish in one cycle. Data entered at the input nodes at each clock period will thus flow through the structure in a wave-like motion, ending in one destination which provides the final result.

The clock frequency will determine the processing capacity of the systolic array while the number of pipeline steps will define the latency.

Bit-parallel or bit-serial data representations

In a bit-parallel systolic array, the data are entered, processed and presented in a parallel form. The data will flow uniformly through the computation structure. Successive data will be completely independent since the components are without memory.

With a bit-serial data representation the process is carried one step further. Here the data are inserted into the system one bit at a time. Corresponding bits will then reach the compute stations where they are processed at the same clock cycle, presenting the corresponding result bit after a certain latency. The data flow in a bit-serial processing pipeline is usually not memory-less since consecutive bits will influence each other (e.g. via carry bits). It is not uniform since a reset signal must accompany the first, least-significant bits, clearing the memory functions so that consecutive data will stay independent. Formally this means that the bit-serial processing pipeline no longer satisfies the definition of a systolic array.

Before entering a bit-serial processing pipeline the data must be zero-extended by adding zeros after the most significant bit so that the word length is equal to that of the most precise intermediate result (n bits). Otherwise bits will be lost as the words overflow their allotted time frame. This means that full operations will take at least n clock cycles instead of one for the parallel case. However, the basic bit-serial operations are in most cases faster than corresponding bit-parallel ones, allowing an increased clock rate (e.g. bit-serial addition is not hampered by carry propagation like bit-parallel is). The bit-serial speed gain may, in favourable cases, be a factor $\approx 2-4$ depending on the word width. Multiple bit-serial calculation blocks (farms) can then be used to obtain the required throughput.

A bit-serial farm with $n/2$ to $n/4$ pipelines must be used to achieve the same throughput as for a single bit-parallel processor. This increased circuit complexity is acceptable since the

gate counts of the basic bit-serial processing elements are much smaller than their bit-parallel counterparts. In the case of additions and comparisons, the reduction factor is approximately n . Thus, with algorithms predominantly composed of additions and comparisons, the overall gate count saving of bit-serial architectures is considerable.

The latency of each full operation depends on the clock rate and the type of operation performed. Additions can in both cases be cascaded with a latency of one clock period. This means a considerably shorter latency for bit-serial additions due to higher clock frequency. Bit-serial comparisons will require all bits before delivering the result, i.e. a latency of n clock periods, which is slower than for the bit-parallel case. To these latencies one has to add the overhead of serialisations and parallelisations. Thus the total latency effect depends very much on the algorithm.

Bit-parallel processor demonstrator

We have constructed and tested a first-prototype bit-parallel calorimeter trigger system. Its aim was to test the efficiency and rejection power of the e.m. cluster-finding algorithm described above, running at full LHC speed and using affordable present-day electronics, with real signals from prototype LHC calorimeters in test beams. The core of this system is a custom-made integrated circuit (ASIC) that implements the algorithm.

Description of the ASIC

The ASIC implements most of the features of our trigger algorithm. Referring to Figure 6, it carries out the operations needed for one e.m. trigger channel. This means that it must work on the 8-bit energies from 16 trigger channels. It forms separate energy sums with one horizontal and one vertical neighbour to find potential e.m. clusters. It also sums the energy in the outer 12 e.m. cells to examine isolation. A cluster is found if the vertical sum or the horizontal sum is greater than a programmable cluster threshold, and the sum of the outer 12 cells is less than a programmable isolation threshold. This logic is duplicated, so that two pairs of programmable threshold values allow for two different cluster conditions. In addition, the sum of the energies of all 16 input channels is formed.

In this first prototype chip, no attempt is made to handle hadronic trigger information.

To implement the algorithm on an ASIC, the logic was broken down into 15 ns pipelined steps since that was the initially proposed LHC bunch-crossing period. Further details of the ASIC design are given in [14]. The ASIC uses a 0.8 μm CMOS gate array and was produced by Fujitsu. The cluster-finding results are available within 7 clock cycles, and the 12-bit 16-cell total energy sums within 6 clock cycles. The package is a 179-pin ceramic pin grid array. Tests at the full design frequency of 67 MHz have been successful.

Description of processor hardware

The prototype trigger processor is designed to trigger on a 3×3 area of calorimeter trigger cells. It must first digitise the signals, then pass them to a cluster-finding module (CFM) containing nine ASICs. A clock module is needed to distribute clock signals to the ADCs and the cluster-finding module.

A block diagram of the system is shown in Figure 7, while the layout of the electronics is shown in Figure 8. The configuration shown is the one used for tests in conjunction with RD3, described below. Although VMEbus is used for computer interfacing and control, the trigger modules needed to be larger in size so we modified some crates originally developed for UA1. Modules for these crates are 9U high and 40 cm deep (1U = 1.75 inch). At the

rear the modules plug into two backplanes. The lower one is a printed circuit carrying TTL-compatible signals that uses 96-way connectors and provides a simple control and addressing system. This is organised by a crate controller module which is in turn controlled by a VME interface module. The upper backplane allows wire-wrap connections between modules and uses Teradyne connectors with 320 active pins and 80 grounds. We use this to transmit fast ECL signals. The wire-wrap backplane minimises the amount of cabling on the front panels of the various modules.




Figure 6: Calorimetry associated with one e.m. trigger cell.

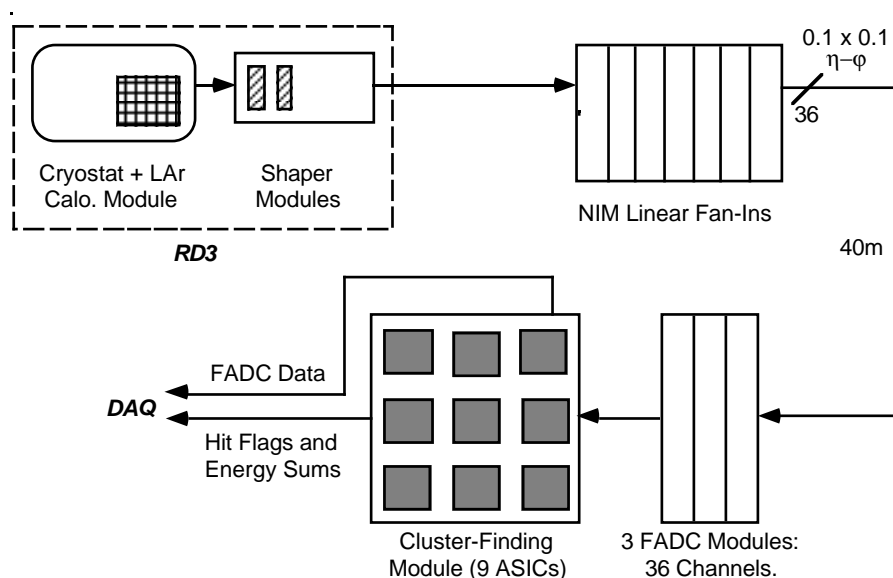


Figure 7: Block diagram of the calorimeter and trigger system.

Calorimeter signals are digitised by three flash-ADC modules (FADC), each having 12 channels. The FADC hybrids were of the type used in the ZEUS experiment at HERA, and are 8-bit devices that can run at up to 104 MHz. The modules are simple, with no on-board memory, since we rely on the input memories of the cluster-finding module to monitor and record the FADC data.




Figure 8: Schematic layout of the electronics.

The prototype trigger system is built around a single cluster-finding module, which incorporates nine cluster-finding ASICs. A block diagram of this module is given in Figure 9. The module was developed in order to attack the problems of using the ASICs under realistic conditions at full speed. Therefore, in addition to control facilities, it includes fast memories on the inputs for capturing and reading out the FADC input data and FIFOs on the outputs for reading out the hit results and energy sums from the ASICs. The memories on the inputs can also be loaded with test data that can then be clocked through the logic. The module has been successfully tested at speeds up to 70 MHz.

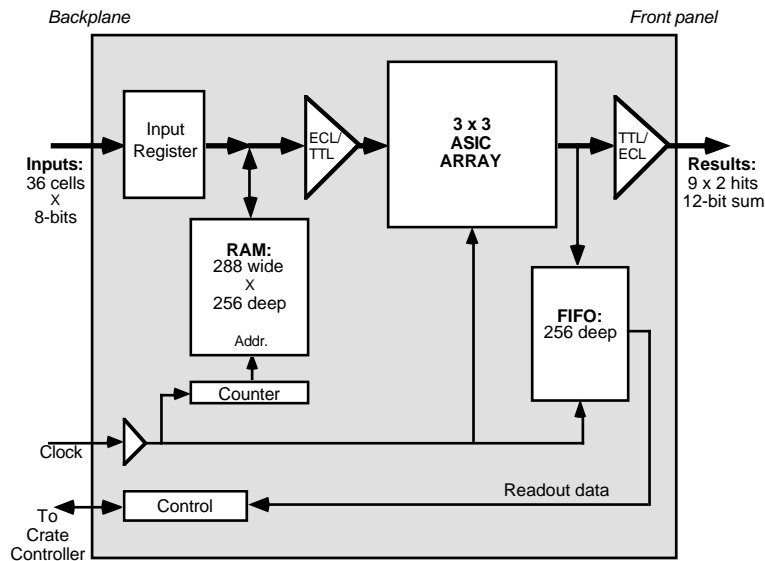


Figure 9: Block diagram of the cluster-finding module.

Synchronisation of modules within the trigger crate was under the control of a clock module, which distributed correctly phased clock signals to the three 12-channel FADCs and to the cluster-finding module. A system clock frequency of 40 MHz was used for most data collection, corresponding to the anticipated 25 ns bunch-crossing period of the LHC, although data were also taken at 67 MHz.

On each system clock, digital data from the 36 FADCs were transferred in parallel to the cluster-finding module and inserted into a pipeline running at the system clock frequency. Trigger bits resulting from the cluster-finding algorithm emerged seven clock cycles later from the ASICs and were available as prompt front-panel outputs.

Copies of the incoming FADC information, output trigger hits and ASIC energy sums were saved on the cluster-finding module in high-speed memories, allowing up to 256 time slices to be recorded. During data-taking, the system clock was normally free-running, and the memories scrolled continuously at the system clock frequency. At any instant, therefore, the memories contained a history of the preceding 256 FADC samples and the corresponding trigger algorithm analysis results from the nine ASICs. The phase of the event trigger (see below) with respect to the system clock was recorded using a CAMAC time-to-digital converter.

A fast event trigger signal was derived from a coincidence of beam scintillation counters. Gating and monitoring of event signals, and various other control functions, were done using a combination of standard NIM and CAMAC modules, with the CAMAC connected via a VME-based branch driver.

On receipt of the event signal, the clock module counted a further set number of system clocks before generating a stop signal to freeze the cluster-finding module memories. The event signal was also sent to a VME-based data acquisition system, running a variant of the CERN Spider suite [15] adapted for RD27 and running on an MVME167 (68040-based) processor. This selectively read out and recorded the memory contents of the cluster-finding module, and performed on-line analysis on a sample of the data. A copy of the data was transmitted in real time over a VME-VME link to the RD3 data acquisition system, where a duplicate recording was made together with complementary information from the RD3 calorimeter. On completion of data readout, the trigger crate was restored under software control to the initial free-running state.

An Apple Macintosh application developed specially for RD27 was used to set the contents of programmable registers on modules in the trigger crate. Parameters controlled in this way included the ASIC cluster and isolation thresholds, clock frequency and phase, and the number of system clocks preceding the stop signal generation. In addition, the Macintosh was used to run various diagnostic test programs.

Tests with the RD3 “Accordion” calorimeter

We have now had two periods of running together with the RD3 prototype liquid-argon Accordion calorimeter [16] in an SPS test beam in the North Area at CERN. In order to form trigger cells of the desired granularity, analogue signals from the calorimeter were added both laterally and in two $9 X_0$ depth samples using LeCroy 428F NIM linear mixers, as shown in Figure 7.

The full trigger processor system was not ready in time for our first tests in November 1992. We therefore chose simply to digitise the data using FADCs and record them for later analysis and off-line “playback” to the trigger processor. We digitised the data from a 4×4

region of the calorimeter using a 16-channel F1001 FADC module borrowed from the H1 experiment at HERA. The FADCs were 8-bit and could sample at rates up to 104 MHz. The data acquisition system (VME and CAMAC) was independent of RD3 and simply ran asynchronously. Control and monitoring were done by software running in a Macintosh, and data recording was done onto a hard disk. The analysis of these data is described in Ref. [17].

In April–May 1993 we ran with our prototype trigger processor installed. The configuration is shown in Figures 7 and 8 and was as described above. We recorded a total of 446 data runs, where a “run” means a unique combination of beam particle, beam energy, beam position in the calorimeter, and cluster energy and isolation thresholds in the trigger.

Analysis of test-beam data from the RD3 calorimeter

At the time of writing this report, data from the two test-beam periods with the RD3 calorimeter prototype are being analysed [17, 13]. As has been described above, the equipment used in the two periods differed greatly. For both periods we recorded the signals from an area of the e.m. calorimeter. These data enable us to study, with two different FADC systems, the expected performance of the trigger algorithm with real calorimeter data, including the effects of electronic noise. In addition, these data are being used for studies [18] of algorithms to identify the bunch-crossing to which a calorimeter signal belongs. In the second period we also have the results of the real-time processing of the calorimeter data performed by the demonstrator trigger system, which, combined with the stored FADC signals, enables us to study thoroughly the performance of the demonstrator hardware.

It should be noted that the analysis, particularly of the data from the second period, is still under way and results shown below should be considered as preliminary. We include here only a brief review of our initial studies.

In the beam tests the trigger system had its own free-running clock. Since the clock period was short compared with the pulse width, the signal from a single shower is sampled over several clock cycles. A typical pulse sampled at a frequency of 40 MHz is shown in Figure 10. The significant differences in the pulse heights measured in consecutive digitisations offer the possibility of unambiguously associating a calorimeter pulse with a single bunch crossing, as discussed in more detail below.




Figure 10: Pulse from the Accordion calorimeter sampled at 40 MHz by an FADC.

Pedestals were measured for each channel in each run, using the ADC values measured in digitisings prior to the arrival of the pulse. Typical pedestal values and rms widths of 10 and 0.4 counts respectively, were found for the F1001 system used in the first test-beam period, and 26 and 0.6 counts for the purpose-designed ADC system used in the second period. We have checked that, with both ADC systems, the pedestal values were stable over periods of several days. The observation of low rms values is important because the amount of noise seen in unoccupied channels has a considerable impact on the performance of the trigger, especially in the isolation veto.

Calibration of the ADCs was performed using various methods, using electron beam data of several different energies, test pulse data and cross calibration against the precision ADCs of RD3. A typical calibration constant close to our design aim of 1 GeV per count was obtained.

We have used electron beam data of different energies to measure the energy resolution of the trigger. Clusters are formed from the sum of energy in two adjacent cells, and the cluster energy distribution fitted with a gaussian. An example of the distribution observed is shown in Figure 11. The long tail below the peak, due to pion contamination of the beam, is ignored in the fit. The resolutions obtained from the two beam tests are summarised in Table 3, together with the expectation for an intrinsic calorimeter resolution of $10\%/\sqrt{E} \oplus 1\%$.

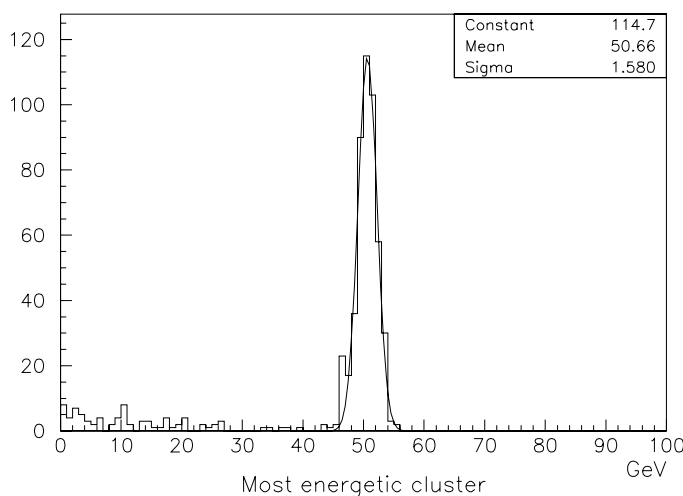


Figure 11: Distribution of energy measured by trigger in cluster window.

The two-cell cluster algorithm is proposed in order to maintain an acceptably sharp trigger threshold in the presence of sharing of energy between cells in the calorimeter. This sharing is illustrated in Figure 12, which shows the energies measured in two cells as the beam was scanned across the boundary between them. As can be seen, the sum of the two energies remains constant across the boundary.

Beam Energy	10 GeV	20 GeV	50 GeV	100 GeV
Nov. 1992	0.77 GeV	1.06 GeV	1.24 GeV	1.81 GeV
April 1993	1.54 GeV	1.65 GeV	1.53 GeV	
$10\%/\sqrt{E} \oplus 1\%$	0.33 GeV	0.49 GeV	0.87 GeV	1.41 GeV

Table 3: Cluster energy resolution of the trigger ADC system, compared with pure calorimeter resolution. The somewhat larger values for the 1993 data are due to the wider pedestal distribution of the FADCs.




Figure 12: Sharing of energy of shower between trigger cells (data from November 1992).

The beam spot in these tests was of about the size of a single calorimeter cell, whereas a trigger cell was the lateral sum of 4×4 calorimeter cells. In a real experiment the electrons will be evenly distributed over the trigger acceptance, so in order to study the trigger threshold sharpness one needs a number of data samples taken over a range of points within a trigger cell. Since a trigger cell contains four calorimeter cells in the central core, four at the corners and eight along the edges, a “representative” distribution can be obtained from one data point taken in the core, one in a corner and two along an edge. Such a group of four runs was recorded during the 1993 beam test with a 50 GeV electron beam. The data from the four runs were combined, and the efficiency of the cluster algorithm as a function of threshold computed. This is shown in Figure 13a, with the Monte Carlo expectation for a similar calorimeter geometry shown for comparison. Figure 13b shows the corresponding threshold curves for a trigger based on a single cell, which has much less sharp threshold behaviour. In both cases there is reasonable agreement between the data and the simulation.

The other element of the e.m. cluster algorithm is the isolation requirement. This is primarily designed to reject jets mimicking the electron/photon signature, and thus the data taken during these tests tell us little about its performance. However, the efficiency of the requirement for isolated electrons and its dependence on noise suppression can be studied. Figure 14 shows the efficiency versus threshold on the energy measured in the isolation ring with and without a trigger cell threshold, in events containing a 50 GeV electron cluster. Requiring that the pedestal-subtracted cell energy be greater than twice the pedestal rms is extremely effective in suppressing noise, giving an efficiency of 95% for a requirement that the isolation energy be less than 5 GeV. The corresponding distribution from the simulation is shown for comparison. This was calculated assuming a noise level equivalent to 360 MeV in a $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ trigger cell area in the calorimeter, which is consistent within 10% with the noise contribution measured in earlier RD3 tests and slightly lower than used in the simulation studies described above.




Figure 13: Efficiency versus threshold (a) for two-cell algorithm, (b) for single-cell algorithm. In each plot the test-beam result is shown as the solid line and the Monte Carlo result is shown as the dashed line.

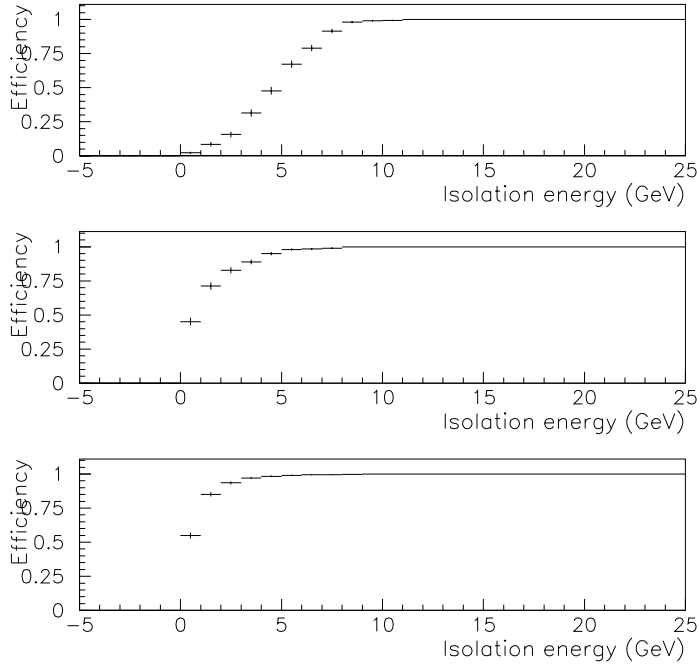


Figure 14: Measured and simulated efficiency versus isolation threshold: (a) Pedestal-subtracted data with cell threshold = 0, (b) Pedestal-subtracted data with cell threshold = 2σ , (c) Monte Carlo.

A final test of the trigger algorithm is its ability to distinguish electrons from pions. Cluster and isolation thresholds were chosen to achieve a high efficiency ($>95\%$) for electrons with energies of 10, 20, 30 and 50 GeV, and these were applied to data taken with a pion beam. The rejection power of these selections is summarised in Table 4. As expected, the e.m. isolation requirement contributes little to the rejection power against single particles. Nevertheless a significant rejection is obtained from the cluster threshold alone, and

hadronic isolation, not included in the demonstrator system, would increase this significantly.

Beam energy	Cluster threshold	Cluster plus Isolation
10 GeV	14 ± 4	15 ± 4
20 GeV	42 ± 9	45 ± 10
30 GeV	260 ± 130	340 ± 190

Table 4: Estimated rejection power of the trigger algorithm against pion beams of different energies. Trigger thresholds chosen to give 95% efficiency for electrons of the same energy. As the pion beam contains a significant muon contamination these numbers can only be taken as indicative.

During the beam test in April–May 1993, using our demonstrator system of nine cluster-finding ASICs, both the FADC data and the ASIC outputs were recorded. Thus it is possible to determine with precision whether the system behaves as expected, and to study in detail any problems which we encounter. In the following, a clock period of 25 ns was used unless stated otherwise.

Each cluster-finding ASIC accepts 16 8-bit energy values every clock cycle. It calculates the sum of all 16 inputs, the sums of the two cluster options (horizontal and vertical) and the 12-cell isolation sum. The 16-cell energy sum, which in a full system would form an input to the jet trigger, is available on the front panel of the trigger module after six clock cycles. The cluster and isolation sums are compared with thresholds, and if either of the two clusters is greater than the cluster threshold and the isolation sum is less than the isolation threshold a trigger hit is generated. There are two combinations of cluster and isolation threshold, so two different trigger bits may be set. The trigger hit results are available on the front panel of the trigger module after seven clock cycles. The demonstrator system includes no calibration, pedestal subtraction or zero suppression facilities, and so the raw ADC counts form the input to the algorithm.

During the design of the ASICs an error was made in the connection of the input signals, so that instead of vertical and horizontal sums being formed for the clusters, vertical and diagonal sums were instead formed. This effect was included in the simulation of the trigger algorithm in the study of the demonstrator performance. While it would have been straightforward to correct this hardware error by making a second prototype cluster board, we did not consider that the expense was justified.

We have compared the energy sums read out from the ASICs with the expectation from the FADC signals. In about 80% of events there is exact agreement between the ASIC results and expectation, with the ASIC result appearing correctly with a latency of six clock cycles. In almost all of the remaining events exact agreement is still found, but with the ASIC result appearing one cycle late in the readout pipeline. We are optimistic that this problem will be resolved for future beam tests by re-adjusting the clock phases used to strobe the readout pipelines. In Figure 15 we show the correlation between the real-time calculation and the prediction after correcting in the analysis for this readout problem.

To avoid confusion due to the readout timing problems described above, only events in which all ASIC energy sums appeared “in-time” were used in the following analysis.

In the e.m. cluster logic there are two requirements – that one of the two (horizontal/vertical) cluster windows lie above the cluster threshold and that the sum from the isolation ring lie below the isolation threshold. In order that these two operations could be tested separately, in many runs the isolation requirement was disabled. The affect of the

cluster threshold is illustrated in Figure 16. In this particular run, one cluster threshold was set at the value of 60 counts, only a few counts above the pedestal sum for two channels. In all the events examined the cluster threshold test worked perfectly.




Figure 15: ASIC output versus expected value after correcting for readout problem described in the text.




Figure 16: Energy in the cluster window for events with (solid line) and without (dashed line) a “hit” from the trigger processor.

We have also analysed runs in which one of the isolation thresholds was set close to the expected 12-trigger-cell pedestal sum. Choosing events and ASICs for which one of the clusters should be above the cluster threshold, one can compare the isolation sums in events for which the trigger hit was or was not set. The isolation requirement worked as expected as shown in Figure 17.

The ability of the trigger processor to separate electrons from pions and muons in real time can be demonstrated by using RD3 data tapes, which included a copy of the RD27 data. Figure 18a plots the energies observed in the RD3 e.m. and hadronic calorimeters against each other. Three distinct types of event can be seen – events with a large deposit of energy in the e.m. calorimeter and none in the hadronic calorimeter (electrons), events with little or

no energy in either calorimeter (muons), and events in which energy is shared between the two calorimeters (pions). Figure 18b shows the same distribution for events in which the RD27 e.m. cluster trigger was fired. The trigger threshold in this run, 200 counts, corresponds to approximately 150 GeV (after pedestal subtraction).




Figure 17: Energy in isolation window for events with (solid line) and without (dashed line) a “hit” from the trigger processor.

The pulses from many proposed LHC calorimeters are longer than the bunch-crossing period, as shown in Figure 10. The unambiguous association of such signals to a unique bunch crossing simplifies the trigger, since otherwise a large pulse would cause the trigger to be fired for a number of consecutive bunch crossings. It may also reduce the data rate to the trigger processors, since the signal need only be transmitted for one bunch crossing. Investigations are therefore being made into bunch-crossing identification using digital signal processing algorithms. These algorithms have been simulated in software and tested with the data recorded in the November 1992 and April 1993 beam tests, with both 25 ns and 15 ns sampling periods. The studies using the November 1992 data are described in detail in Ref. [18].

Algorithms studied include a peak-finding algorithm, a zero-crossing identifier, a constant-fraction discriminator and a deconvolution filter. Very encouraging results have been obtained, even for relatively small deposited energies. For example, using data from the November 1992 tests with a 25 ns bunch-crossing period, a simple peak-finding algorithm was able to correctly identify the bunch crossing in > 99 % of cases for energies as low as 4 GeV in a single trigger cell.

In addition to the off-line studies described above, a single-channel bunch-crossing identification module is currently under construction. This will be able to apply any of the above algorithms to calorimeter data in real-time. This demonstrator prototype makes extensive use of static RAM look-up tables.

Although analogue signal processing has not yet been investigated in RD27, it is still being considered as an alternative method of bunch-crossing identification.

Finally, to demonstrate real-time operation of the processor, in Figure 19 we show an oscilloscope photograph of calorimeter pulses at the input to the FADC system and the trigger processor output which follows after the expected latency, including cable delays.




Figure 18: Energy in ECAL versus HCAL (a) for all events and (b) for events with e.m. cluster “hit”.




Figure 19: Oscilloscope photograph showing calorimeter pulse at input to trigger FADC (upper trace) and the output of the trigger processor (lower trace).

Tests with the RD33 TGT calorimeter

A prototype TGT calorimeter [19] was installed in an SPS test beam in late June 1993. In parallel with the RD33 calorimeter performance measurements, a test of its features for triggering was performed by RD27.

The TGT prototype contains electronics in which readout cells are summed to form trigger towers of area $6 \times 5 \text{ cm}^2$. The trigger signals from the calorimeter are delivered to the external electronics separately for the five layers in depth, requiring further analogue

summation before digitisation. The total number of trigger towers is 4×4 , which is the minimum number needed to feed one of the cluster-finding ASICs including the isolation window.

We built VME boards which receive the five analogue signals from a TGT tower. Summing and shaping is performed before the signal is digitised by an 8-bit FADC sampling the analogue-sum signal at 40 MHz produced by a free-running clock. A fast memory stores the FADC digitisations in a 256-word circular buffer until the content is frozen by an external STOP signal. The memory serves as a debugging aid for the apparatus by showing the digitisation history back to triggering pulse and beyond. The complete FADC system consists of a VME crate housing eight FADC modules (two channels per module). Interface boards to the special backplane of the trigger processor demonstrator, on which signal distribution to the cluster-finding modules is performed, were also built for these tests.

A Macintosh software package was used for data acquisition in the RD33/RD27 joint test. All components of the trigger system were tested in the laboratory prior to installation at the test beam. Analogue trigger signals from the TGT were simulated with a pulse generator. Digitisation and transmission to the cluster finder showed satisfactory results.

The full trigger system was installed in the North Area at CERN for the TGT test-beam period at the end of June 1993. The analysis of the data which were collected is in progress.

Bit-parallel processor design study

We briefly describe here a preliminary, but fairly complete, design for a bit-parallel first-level calorimeter trigger processor; more detail is given in [20]. This design will, of course, evolve as our studies and prototype work progress and also as new technology becomes available. However, it should be stressed that the present design looks feasible and is based on technology available now or in the near future.

The processor consists of e.m. cluster logic, jet logic and missing- E_T logic. The e.m. cluster logic is the most demanding part of the system, and will consist of four crates of electronics. The remaining electronics will be housed in two further crates. The processor will use four different ASIC designs and seven different circuit board designs. Figure 20 shows a block diagram of the trigger processor system. The total latency of this processor will be under 500 ns.

Receiver module

A receiver module will take in digitised, serialised, zero-suppressed data at 160 Mbits/s from the ADC system. If the connections are optical, the receiver module must convert them to electrical signals. There is one receiver module per cluster processor module, so there would be 64 in total (see below).

E.M. cluster-finding ASIC

The cluster-finding ASIC will do the required processing for 16 e.m. trigger channels, including hadronic information. This means that it must examine the transverse energy in a total of 98 channels (7×7 trigger cells for each of the e.m. and hadronic calorimeters). To reduce the number of input/output (I/O) pins and the bandwidth, the information for each channel will be sent as a single, zero-suppressed 160 Mbits/s serial stream. Thus serial-to-parallel conversion and tag-matching are required. With derandomizing buffers at each end of the asynchronous link four words deep, the fraction of data which are lost can be reduced

to a negligible level³ for a zero-suppression threshold of $E_T = 1$ GeV even at the highest LHC luminosity; the fact that this threshold is programmable gives additional safety.

The ASIC will produce results for eight sets of threshold values (cluster threshold, isolation threshold) 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. A total of 256 ASICs would be needed to process the entire calorimeter.

The trigger latency could be reduced by either running the cluster-finding ASICs at 80 MHz (the first prototype already runs at more than 67 MHz) or by reducing the number of pipeline stages and so doing more processing per stage.

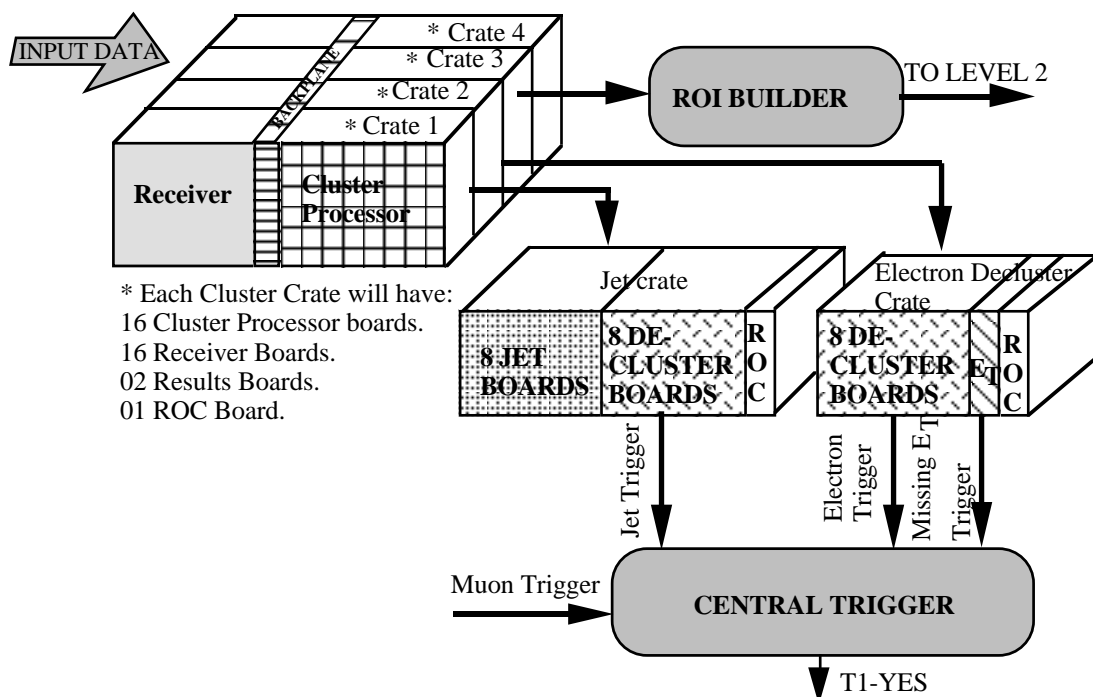


Figure 20: First-level calorimeter trigger system.

E.m. cluster processor module

The cluster processor module will contain four cluster-finding ASICs and one adder ASIC to combine the E_T sums from the ASICs, as well as look-up tables to convert E_T to its components, E_x and E_y . A total of 64 such modules would be required.

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 two's-complement results of this adding process will be sent to the missing- E_T module. A total of eight such modules would be required.

³ The fraction of data which are lost under the very pessimistic assumptions discussed in the simulation section is 0.2%.

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 module would 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. These areas 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 of the 13-bit sum areas, i.e. a total of nine jet windows. Approximately 30 of these 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, or “declustering”, 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 modules⁴ 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 and 16 receiver modules. The receiver modules and the cluster processor modules will be plugged back-to-back via the backplane connectors⁵. 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.

⁴ 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.

⁵ Such back-to-back connections are possible using commercial technology such as Molex Omnigrid 2.5 connectors.

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

Proposal for future work

We have identified several key areas of the bit-parallel processor that require further study before a full level-1 system could be designed and constructed. We therefore propose building a phase-2 bit-parallel demonstrator system. This would be based on a new, high-speed ASIC together with associated support modules. The aims are:

- To demonstrate our data sparsification scheme, using zero-suppression and asynchronous serial data transfer to reduce the number of interconnections required.
- To include a custom digital signal processor in order to implement and test bunch-crossing identification using various algorithms.
- To address the question of whether the high-speed links between the front-end digitisation electronics and the level-1 processors should be electrical or optical.

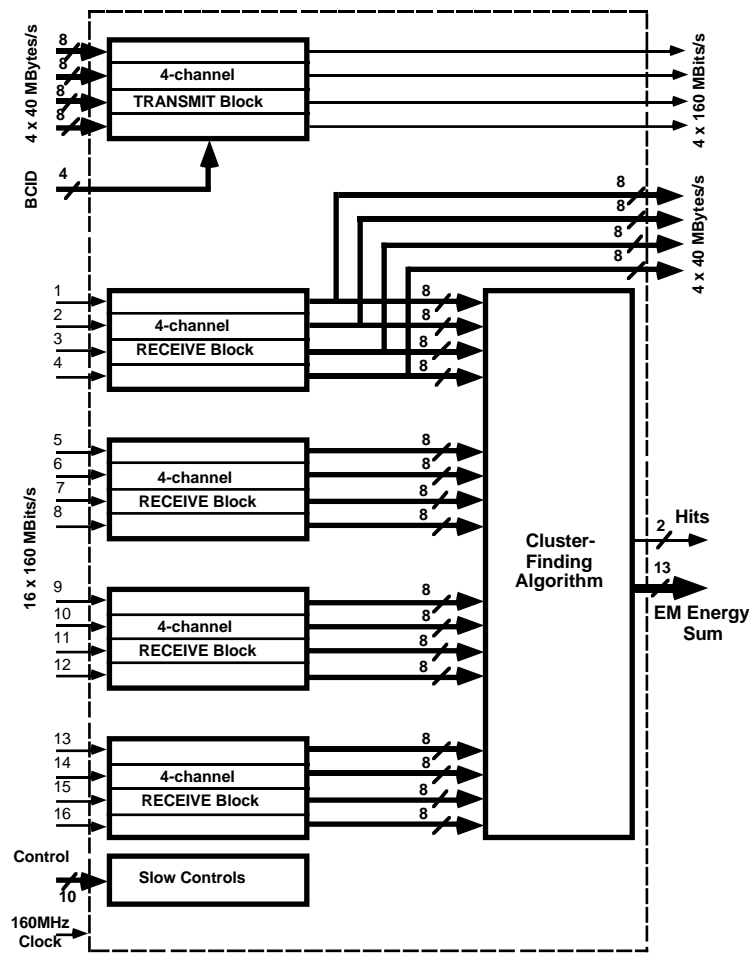


Figure 21: Phase-2 demonstrator ASIC.

The new ASIC would be designed explicitly as a test bench. We feel that our e.m. clustering algorithm has been adequately demonstrated by the existing prototype system, and we believe that putting more channels onto one ASIC in order to reduce the overall size of the system and its interconnections is a relatively straightforward task. Thus, in order to minimise the cost of the new ASIC we will stay with one channel per chip, and also put onto the same type of chip other functions that we want to test but which would not be located there in the final design. This is illustrated in Figure 21.

The critical functions that we would like to test are:

- “Transmit” block: receive and buffer (FIFO) 8-bit wide data at 40 Mbytes/s, perform zero suppression, tag, do parallel-to-serial conversion, and transmit serial data at 160 Mbits/s. This section would actually be part of the ADC system in an LHC experiment.
- “Receive” block: receive data at 160 Mbits/s, perform serial-to-parallel conversion, buffer (FIFO) and tag-match.
- “Algorithm” block: receive data as above for 16 e.m. trigger channels and perform cluster-processing for a 4×4 area as in the first prototype system.
- If gate counts permit, we might also include bunch-crossing identification logic, at least for one channel.

This ASIC requires about 63 inputs and 51 outputs. By using bi-directional pins and sharing pins (the transmit block and receive block would not be used at the same time on the same ASIC) the I/O pin count could be reduced by 36. We are considering implementing the ASIC on a 34k, 0.5 μm CMOS gate array.




Figure 22. Phase-2 demonstrator trigger processor system.

The demonstrator system would be similar to the existing one, which fully processes a 3×3 area of the calorimeter using nine ASICs. Figure 22 shows a block diagram of the new system. The receiver/transmitter (RT) module will have nine of the ASICs in transmit (T) mode, which will only use the transmit block of the ASICs. It will receive 8-bit parallel data from the 36 FADC channels and will transmit the processed data via a transmit interface module on 36 serial links at 160 Mbits/s/link; these links could be optical or electrical. The cluster processor module will receive the serial data via a receive interface

module and will have nine ASICs in receive (R) mode, which means that they will use the receive block and the cluster algorithm block to carry out the cluster-finding algorithm. The design of the interface modules will depend on whether the links are optical or electrical. Look-up tables may be included between the FADC system and the RT module to subtract the pedestals before zero suppression.

Bit-serial processor design study

Advantages and disadvantages of a bit-serial calorimeter trigger processor

The calorimeter trigger algorithms are such that most parts benefit from the use of a bit-serial data representation, both with respect to the gate count and the overall latency. However, some operations (e.g. table look-ups) are, for practical reasons, best performed in a bit-parallel mode. A hybrid solution with serial-to-parallel and parallel-to-serial conversions is therefore necessary. Since these conversions introduce latency, the system should be partitioned to minimise the number of transitions between representations.

Most of the electron identification processing occurs in the initial calculations of “environment sums” – two 2-cell e.m. cluster sums, one 12-cell e.m. sum and one 16-cell hadronic sum for each trigger cell. The results of these operations are placed in categories by comparisons to fixed limits. The compound category codes are then converted into a trigger cell classification, e.g. representing high-energy electrons with good isolation, high-energy electrons with medium isolation, etc. The final operation is to count the global number of different classification instances. These numbers are then fed to the central trigger processor for evaluation. Of these operations only the classification has to be done in bit-parallel mode, using table look-up.

The missing- E_T calculation is also made up mostly of additions. The logic can be arranged so that the bulk of this processing is in initial additions of transverse energies from trigger cells along the axial (η) direction, well suited for bit-serial operations. These partial sums can then be converted to bit-parallel representation for multiplication with $\sin\phi$ and $\cos\phi$ and for the final vectorial addition. The jet algorithm can similarly be factorized into a main part suitable for bit-serial operations and a bit-parallel part for the remainder.

A further advantage of the bit-serial approach results from the presence of point-by-point environment operations in the electron algorithm. Such operations favour the parallel processing of many trigger cells in each calculation unit. The efficiency of the implementation, measured as number of results per input connection, increases drastically from 1:32 for a single processing element per processing unit, to 1:2 in the optimal case (infinitely large calculation unit). A bit-serial architecture performs well in this respect because it naturally uses serial data input (allowing many inputs) and because of the high logic density that can be achieved with bit-serial processing.

Concentrating many trigger cell calculations to one module has the added advantage of allowing a large portion of the post-processing to be included as well. However, this concentration implies a high level of complexity. At the clock rates considered this is certainly not trivial, but it is within limits of present technology.

The major components of the trigger processor system

A calorimeter trigger processor system can be said to consist of three major system components:

- Calorimeter trigger preprocessors

- Data transmission
- Calorimeter trigger processor

The trigger preprocessors are responsible for merging high-granularity calorimeter data into trigger-cell data. These operations will probably be performed close to the calorimeter front-ends.

The data transmission component contains data encoding, data transmission and data decoding. The data encoding module will transform data into a physical format which is suitable for transmission. This may include serialisation, data compression and media conversion from electronic to optical signals. Optical transmission is for several reasons the method of choice. It is compact, allows high transmission rates and can be made sufficiently radiation tolerant if necessary. However, several practical problems remain to be solved before sufficiently inexpensive electro-optical components will be available.

Including data compression will reduce the required bandwidth which can be used to lower the transmission rate on each fibre or to reduce the number of fibres, in both cases significantly reducing cost. Disadvantages are increased complexity, latency and reduced fault tolerance. The viability of this approach remains to be determined.

Connectivity presents a problem due to the environment operations present in the electron identification algorithm. Thus each computation unit requires information from overlapping areas on the calorimeter surface. This problem can be taken care of by fanning out channels that are required in different computation units. The fan-out can be done before, during or after transmission. The latter case will require considerable communication between boards and crates, introducing additional complexity as well as potential reliability problems. Whether or not the fan-out can occur by splitting the transmission lines depends on drive capability.

The proposed bit-serial calorimeter processor system

A preliminary but fairly complete calorimeter processor design based on a bit-serial architecture has been suggested [21], as illustrated in Figure 23. By relying on new advanced electronic and electro-optic components, and related construction techniques, it is possible to considerably reduce the physical size of the trigger processor compared to more traditional designs. It is hoped that this size reduction will be accompanied by a significant reduction in cost. An important additional feature is that the design may allow expansion in several respects – higher spatial resolution, higher precision, more complex algorithms and more detailed jet calculations (e.g. di-jet mass cuts). It is important to evaluate the potential gain in physics output associated with such expansions and relate them to the increased cost.

The design is based on a combination of multi-chip modules (MCMs) with massive fibre-optic input (162 fibres), with several high-speed ASICs. Most of the latter use a bit-serial data representation which allows higher processing rates and more compact designs than are possible when using a parallel data representation.

A technical demonstrator is being constructed in which critical technical solutions such as fibre connectors on MCM substrates, optical I/O, and high-speed transfers on the substrate and out from the MCMs will be evaluated. The aim is that these results should eliminate most technical uncertainties involved in the construction of a functional demonstrator. A later step would then be a functional demonstrator that proves the feasibility of the bit-serial calorimeter processor as such by exercising all functional parts.

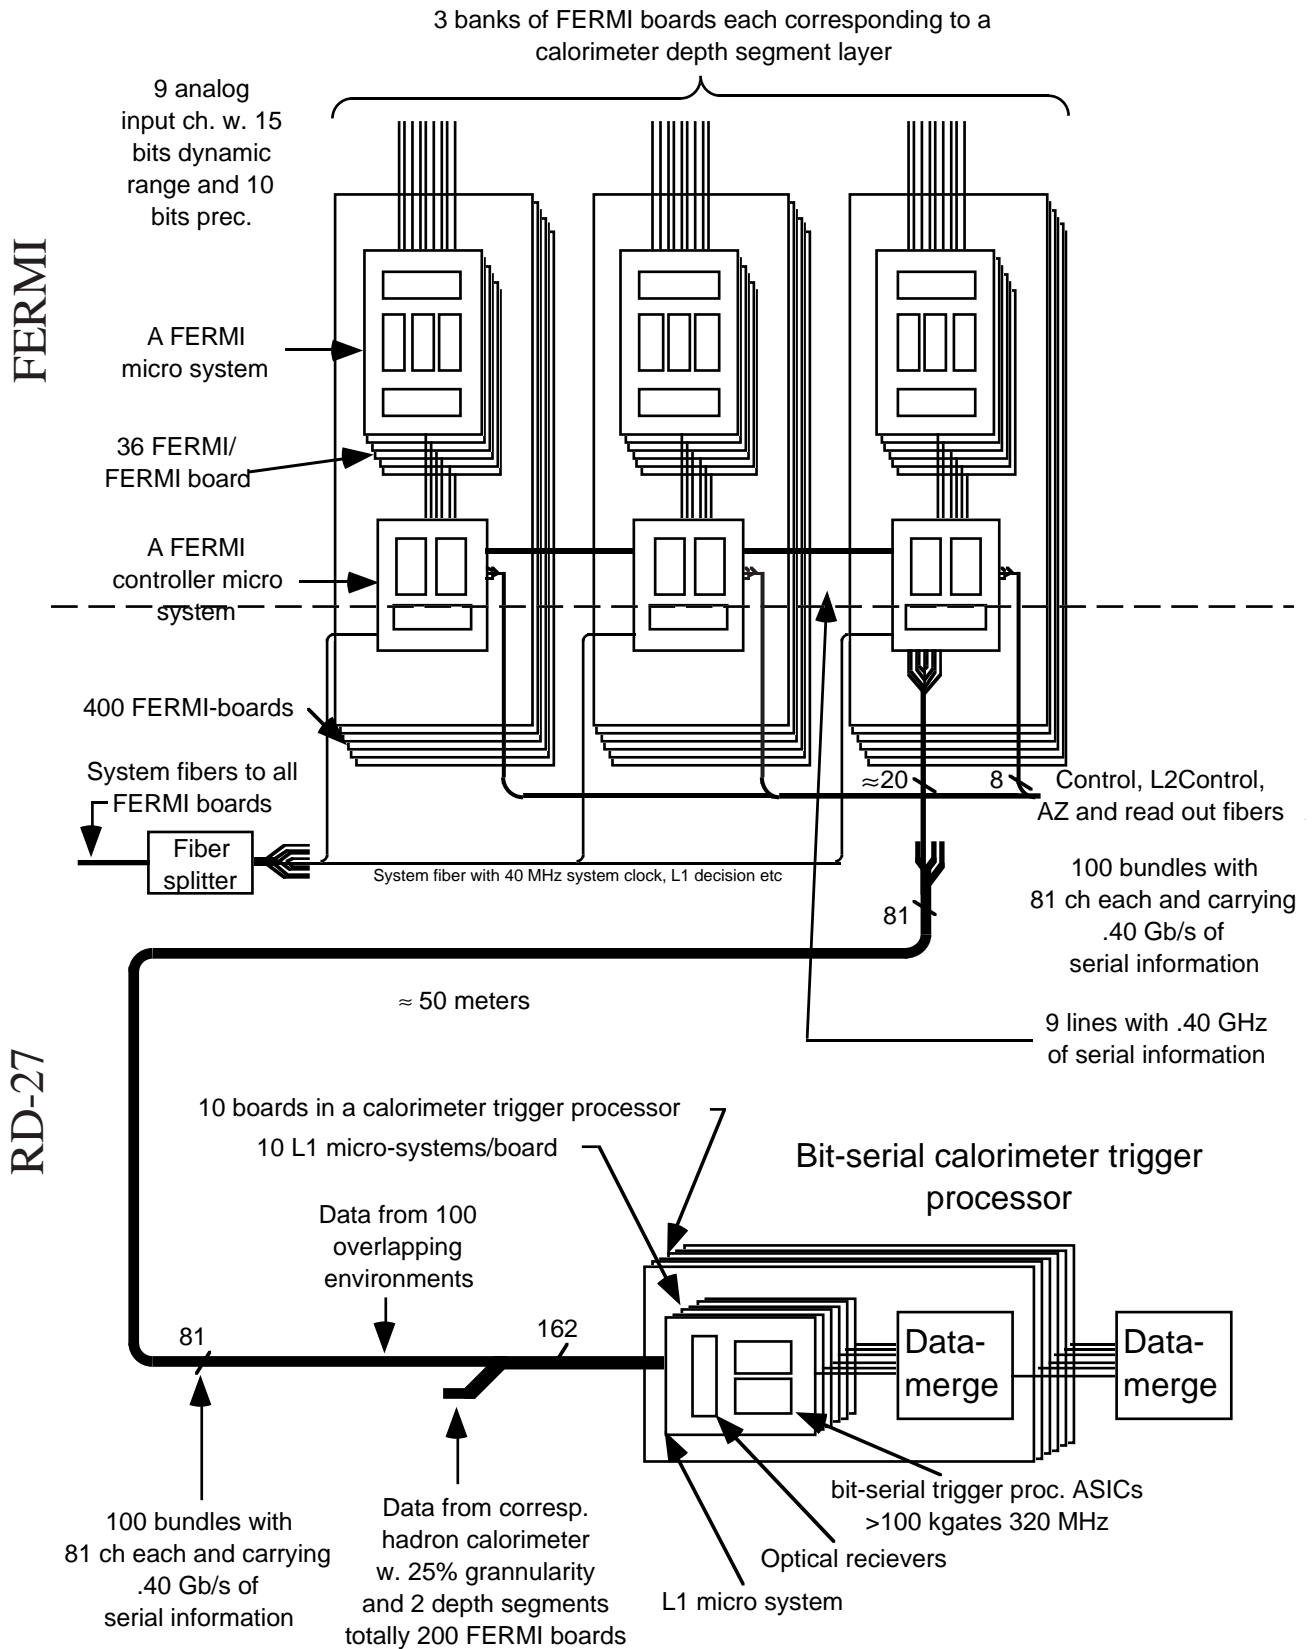


Figure 23: Bit-serial processor system overview.

The full system will contain two parts (see Figure 23). One part is the local first-level trigger preprocessors which will reside on boards together with the calorimeter front-end

electronics, for example with FERMI digital front-end modules [7]. This part may be located on the detector, which would require a high degree of radiation hardness and fault tolerance. These units will be implemented as MCMs with ASICs to sum front-end trigger signals and with light diodes or light modulators driving up to 20 fibres as outputs to the first-level trigger. Such an MCM might also include other front-end electronics so that all functions are kept in this comparatively safe environment. It may include front-end service functions, second-level trigger preprocessing and system clock distribution. This will, however, increase the number of fibre optic outputs. Such a multi-purpose MCM would serve as a general FERMI controller since it would take care of all auxiliary functions required by a farm of FERMI modules.

The second part is the calorimeter trigger processor itself, which will be housed in one create with 100 MCMs as the main processing components. These in turn will contain one GaAs ASIC as a receiver and a farm of two identical high-speed processing ASICs operating as local trigger processors. Most of the calorimeter trigger processor will therefore be implemented with three different ASIC types, including one multi-purpose ASIC for merging the output data streams.

The proposed bit-serial local calorimeter processor ASIC

The heart of the calorimeter trigger processor is the local calorimeter ASIC, illustrated in Figure 24, which consists of two parts. A fast part provides information for the central trigger processor, while a slow part is responsible for extracting region-of-interest information (ROI) for the level-2 trigger processor.

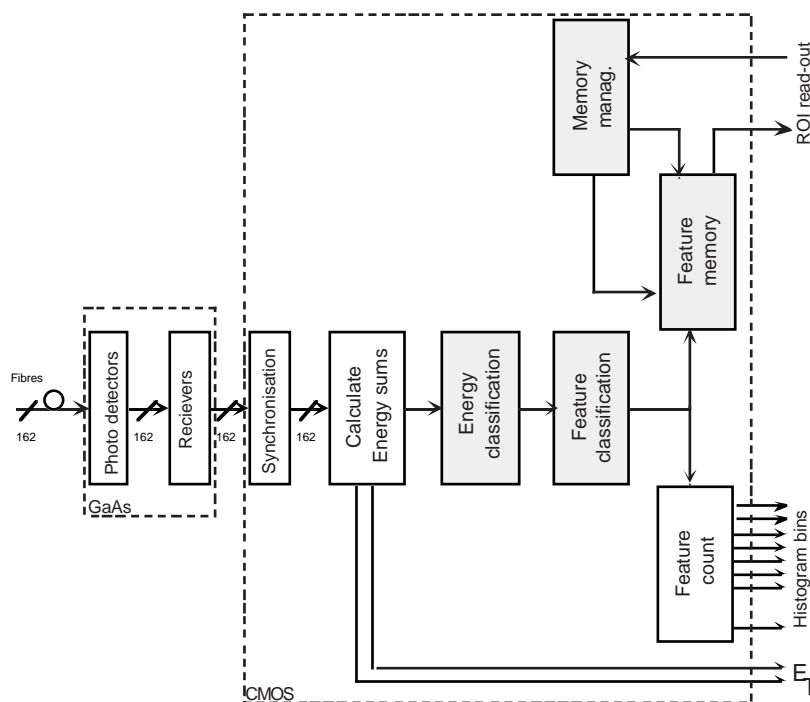


Figure 24: Block diagram of the local calorimeter ASICs, where all blocks delivering bit-parallel outputs are shaded.

The bit-serial data from the preprocessors arrive via optical fibres, photo detectors and receivers. The latter two items are located in a GaAs ASIC. The bit-serial data are then transmitted to a CMOS ASIC where the trigger cell data are summed into environments, transverse energy components, jet-supercells and clusters. These values are compressed by classifying them as belonging to one of 8 energy ranges (3 bits each). Relevant range

identifiers are then combined and used as indices to a feature classification look-up table. Finally the number of feature instances of different types are counted and transmitted as bit-serial data. These data streams, as well as the bit-serial transverse energy components, are merged with corresponding data from the other ASICs on the board and transmitted to a special data merge board to be combined with data from other boards. This information serves as part of the basis for the central first-level trigger decision.

Each feature classification is also stored in a temporary memory which allows region-of-interest information corresponding to affirmative level-1 decisions to be extracted at a later time.

The design has reached a relatively detailed level of description, with VHDL code covering most of the system. The construction has been verified at a high level of abstraction.

The bit-serial technical demonstrator

The technical demonstrator is aimed at evaluating critical design and hardware solutions in the system described above. It will consist of a simple MCM with a GaAs receiver and transmitter chip, together with a silicon ASIC as shown in Figure 25.

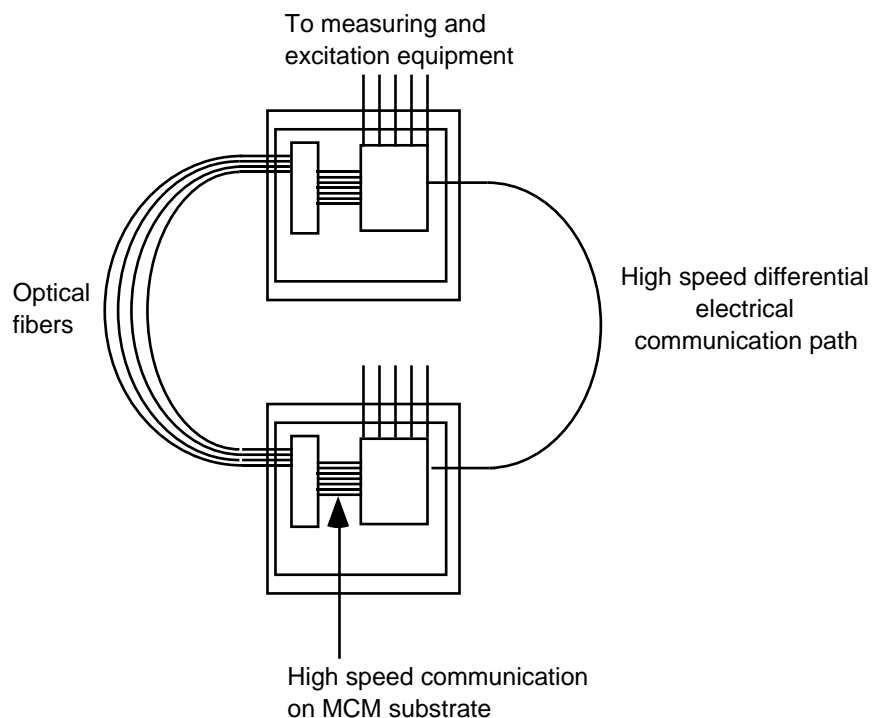


Figure 25: Two interconnected technical demonstrator MCMs.

The design aspects that will be evaluated are:

- High-speed optical communication between MCMs at the required rates
- Synchronisation of the optical signals
- High-speed communication via MCM substrate
- High-speed MCM to MCM communication via circuit board
- High-speed MCM to MCM communication via back plane
- High-speed bit-serial operations

- Comparison between LED and MQW (multiple quantum well) modulators as optical transmitters
- The use of fibre splitters to provide the required data fan-out
- Investigate general problems related to integrating complex high-speed circuitry

Independently of the technical demonstrator the following more general design issues will also be considered:

- Decide the optimum choice between using a few high-speed fibre links (≈ 10 GHz) or many moderate-speed fibre links (≈ 0.8 GHz)
- Consider the feasibility of an all-GaAs design

The development and evaluation of the technical demonstrator will be carried out in close collaboration with several other R&D groups, notably RD16 and RD23.

The technical demonstrator is an interdisciplinary project where groups from different universities, institutes and companies with the required expertise have come together in order to evaluate critical technical solutions. A major part of the ASIC design and the design of MCMs are handled by commercial companies (the Institute for Microelectronics in Stockholm, and SiCon and Swedish MicroSystems from Linköping).

3. Muon Trigger

Introduction

We are studying muon trigger systems for LHC and implementing a demonstrator prototype. The trigger system must be able to identify muon tracks, apply p_T cuts and assign the tracks to particular bunch crossings. The trigger must have good efficiency for muons above the chosen transverse momentum threshold and minimise the background rate. A multi-threshold capability is required since different p_T cuts will be used for single and multi-muon triggers.

The association of the track to a unique bunch crossing is an essential feature of the trigger system. It is achieved by means of a detector and signal processing system giving low overall jitter in time (well under 25 ns). The very large area of the envisaged muon detectors on which the trigger electronics will be distributed has to be taken into account in the design. In particular, one has to be able to control the relative phase of signals coming from different parts of the detector.

Extensive muon trigger simulation studies have been performed in the context of ATLAS by members of our collaboration. This work, which influences the design studies that we are performing, will continue during the coming year. Issues that are currently under study include the implications of the high rate of hits in the muon detectors due to low-energy photons and neutrons in the experimental area, and the special problems of triggers for low- p_T muons for B physics.

A first prototype demonstrator muon trigger is currently under construction and will be tested at the RD5 test beam later this year. In 1994, we plan to upgrade the system by developing a dedicated track-finding ASIC as described below.

Muon trigger simulation

Muon rates and background flux

In order to define the required performance of the muon trigger system it is important to evaluate the flux of particles impinging on the muon detectors. This flux is composed of the following:

- Semileptonic decays of heavy-flavour particles; leptonic decays of W and Z bosons; the Drell-Yan processes. The dominant contribution is due to beauty (and charm) decays [22].
- Secondary muons, produced by decays of charged pions and kaons in flight in the central cavity and punch-through from hadron showers in the calorimeter. These backgrounds are significant only at low transverse momenta ($p_T < 6$ GeV).
- Low-energy photons and neutrons present in the cavern, primarily due to interactions in the forward regions of the experiment. Recent preliminary estimates [23] indicate that the rate of hits in the muon chambers due to this background could be as high as $\sim 10 - 100$ Hz/cm², far higher than the rate due to other sources. We are currently studying the implications for the design of the muon trigger.
- Beam halo
- Cosmic-ray muons

Calculations indicate that the rate for an inclusive muon selection with $p_T > 20$ GeV and $|\eta| < 3$ is about 1 kHz assuming a perfectly sharp threshold and for a luminosity of 10^{34} cm⁻²s⁻¹; the corresponding rate for a dimuon selection is negligible. This threshold value is adequate for high-luminosity physics such as the intermediate-mass Higgs search. However, for B physics significantly lower thresholds will be needed to maximise the rate for CP violation studies. We are currently studying the feasibility of triggering with a threshold of about $p_T = 6$ GeV.

Algorithm

We have considered a model in which the detection of a track is performed using two stations of muon detectors separated by about two meters. Each station can be made from one or more layers of strip detectors, for example RPCs (resistive-plate chambers). In ATLAS this could correspond to detector planes at the midpoint and end of the barrel air-core toroid. In such a scheme, the trigger rate due to random hits induced by low-energy neutrons or photons can be controlled by requiring a local coincidence of hits in several detector layers in each station.

Simulation studies have been performed taking into account the spread of the interaction point, multiple scattering and magnetic deflection, as well as the strip width (3 cm). For each muon hit in the first plane, a search is made in the second plane within a pre-defined region as illustrated in Figure 26. Muons are deflected by multiple scattering in the shielding that precedes the muon detectors and by magnetic bending, low- p_T muons being deflected more than high- p_T ones. The size of the accepted region therefore defines a cut on the p_T of the muon.

Our simulation studies show that, after selecting the cone size to have good efficiency for muons of $p_T > 20$ GeV, the trigger rate is about four to five times higher than with a perfectly sharp cut [22].

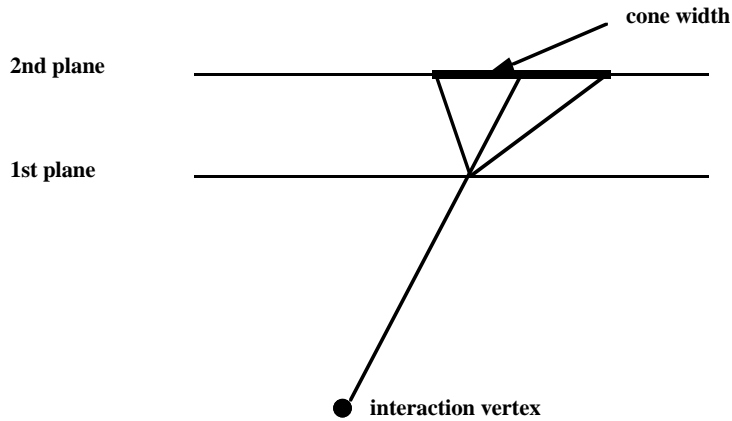


Figure 26: Track detection and p_T cutting algorithm.

System Design Studies

Track detection and p_T cut

The trigger algorithm described above can be implemented by means of coincidence matrices. Signals from strips in the first plane are put in coincidence with the “OR” of a number of strips in the second plane. Different regions in the second detector layer, corresponding to different cuts on p_T , can be programmed in the coincidence matrix as shown in Figure 27.




Figure 27: Multiple threshold coincidence array.

Members of our collaboration have successfully implemented a similar trigger in WA92. This experiment has been running for about two years, using a muon trigger based on two RPC walls. In WA92, the coincidence matrix is made using commercial Programmable Logic Devices, having a delay time of 12 ns and a jitter of 2 ns.

Our proposed algorithm for track finding and applying p_T cuts can in principle be used with any kind of wire or strip detector. It can be used in experiments where a dedicated detector is used for triggering (“stand-alone system”), or where the same detector is used for

triggering and precision measurement (“integrated system”). However, when used with drift-chambers, careful consideration has to be given to the problem of identifying the bunch crossing as discussed below.

Members of our collaboration provided the drift-chamber-based forward muon trigger [24] for the H1 experiment, which includes bunch-crossing identification logic. Building on this experience, we are studying the feasibility of making a level-1 muon trigger for LHC based on the precision measurement detectors. The short bunch-crossing period, the non-planar drift-cell geometries under consideration for the detectors, and the large number of channels significantly complicate the design of such a trigger for LHC. Furthermore, considering muons of relatively low p_T (≈ 6 GeV), the muons are deflected in the spectrometer by large angles. In contrast, in H1 the angle at which the muons impinge on the chambers is relatively well defined. In effect, the trigger logic for LHC would have to determine three parameters (bunch-crossing number, position and angle) in each drift-chamber superlayer, compared to two parameters (bunch-crossing number and position) in the H1 case. A preliminary investigation of the feasibility of building ASICs with the required functionality has been performed.

Unique bunch-crossing identification

A stand-alone muon trigger system can make use of a very fast detector, for example RPCs, with a time resolution much less than the bunch-crossing period. This makes the problem of bunch-crossing identification relatively straightforward, although allowance has to be made for the flight time of the muons and signal propagation delays within the muon trigger system.

We have studied a stand-alone trigger based on an RPC system, which has an intrinsic time resolution of about 3 ns, and where the length of the strips (2 m) corresponds to a variation in the propagation time of the signal of 10 ns, depending on the position of the muon track in the chamber. We are considering electronics which has a jitter of ≈ 2 ns so that the time resolution of the trigger system is well within the 25 ns given by the LHC bunch-crossing period.

Bunch-crossing identification is performed as follows. When a particle hits the first detector plane, a digital signal of about 20 ns duration is generated and sent to the x input of a coincidence matrix located on the second detector plane. The y input of the matrix receives a similar signal generated by the same particle hitting the second plane. The two signals arrive at the inputs of the matrix almost simultaneously (the flight time of the particle through the apparatus partially compensates for the time taken to transmit the signal generated on the first plane to the second one). If there is a coincidence in a programmed zone of the matrix, a trigger signal is generated.

System synchronisation

We have to take into account two fundamental aspects of the “synchronisation” of the system, the distribution of the machine clock over a very large area, and differences in time of flight of particles through the apparatus (this difference can be of the same order of the bunch-crossing period).

For the synchronisation of the system, we plan to make a logical subdivision of our apparatus into “zones”, so that the propagation time of the clock signal from one zone to the next is equal to one bunch-crossing period. In this way, the most remote zone will be the last to receive the clock, the zone next to this will receive the clock one bunch-crossing

earlier and so on. The differences in propagation delays are compensated by introducing at the input to each zone an electronic delay of an integer number of bunch crossings as illustrated in Figure 28⁶. To compensate for different times of flight of the particles through the apparatus an additional delay is added to the clock entering each zone. Programmable delays will be used adding flexibility to the system.

We have considered the need for a calibration system to determine the correct settings for the programmable delays discussed above. It is necessary to have the possibility of generating artificial triggers for calibrating the electronics. In addition, we include a time interpolator for each zone, synchronous with the machine clock and running at four times the machine frequency as shown in Figure 29. The value of the time interpolator for each trigger gives the exact time of arrival of the trigger pulse within the bunch crossing. It will be possible to correctly “time in” the trigger system by analysing the time-interpolator data and by adjusting the programmable delays.

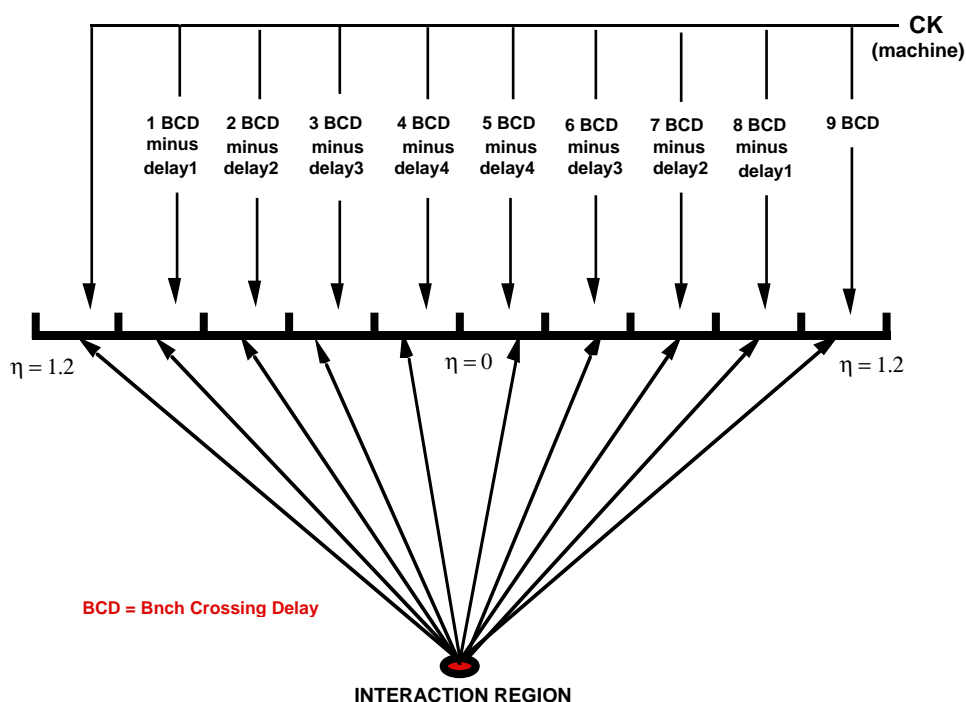


Figure 28: System synchronisation.

The Region of Interest

The trigger system contains information on the position of each muon, its p_T and the bunch-crossing from which it originated. This “Region of Interest” (ROI) information can be used by the level-2 trigger. In our design, the ROI corresponds to a region covered by a single coincidence matrix and its extension in η and ϕ will depend on the specific trigger implementation. For example, in the case of the proposed ATLAS stand-alone triggering system, where RPCs with strips of 2 m length and 3 cm pitch are used with a coincidence matrix of 32×64 , the zone covered by an ROI corresponds to an $\eta - \phi$ segmentation of about 0.1×0.1 .

⁶The board housing the electronics must also be time compensated, but this is a systematic compensation which is relatively easy to implement.

The ROI is the elementary unit of trigger processing in our design. In the case of ATLAS, data for $\sim 10^3$ ROIs will be generated. These will be arranged in groups of roughly 100, to form about ten “sectors” of trigger information. For each sector the system will give the number of triggered muons for every bunch-crossing, and their ROI information including p_T . The information coming from the different sectors is sent to the central muon processor that will generate the global muon trigger information that is passed to the level-1 central trigger processor.

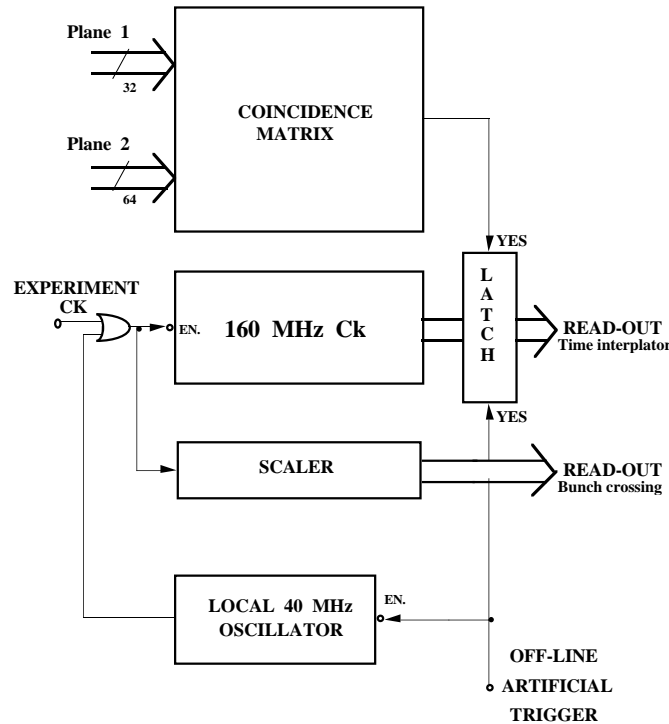


Figure 29: Timing calibration system.

Muon trigger demonstrator

A system to demonstrate the feasibility of the trigger described above is being implemented, based on large-area RPCs already installed in RD5. Construction of the trigger electronics is well advanced and we expect to be ready for the next RD5 run period in September 1993.

The working principle of the demonstrator is illustrated in Figure 30. The outputs of two planes of RPCs that are separated by 2 m are sent to a coincidence matrix, implemented with very fast GaAs commercial components. Whenever a trigger is validated by the coincidence matrix, its time of validation, measured by means of a counter running at four times the LHC frequency (160 MHz), is recorded in a memory. We also record the time of arrival of the RD5 beam trigger, generated from beam hodoscopes using very low jitter logic. The distribution of the time difference between the beam trigger and the muon trigger will thus be available. We will check that this distribution is narrower than the 25 ns LHC bunch-crossing period.

Our demonstrator system is being implemented on five boards, mounted on two VME support modules with interconnections between them. It is linked to the RD5 data-acquisition system by an optical link. The most demanding board contains the coincidence matrix (Figure 31) which is implemented using four GaAs cross-bar switches, each having the possibility of making 32 programmable connections between inputs and outputs. The

cross-bar is programmed to connect each of 32 strips from first plane to four outputs. In this way we map the first plane onto the second, by generating for each strip of the first plane a region of four strips in width. It is then sufficient to make a bit-to-bit coincidence of the cross-bar output with the strips of the second plane to obtain the programmed trigger. The other four boards of the demonstrator contain the cluster finding and p_T encoding, the time-interpolator counter, the memories, and the control logic.

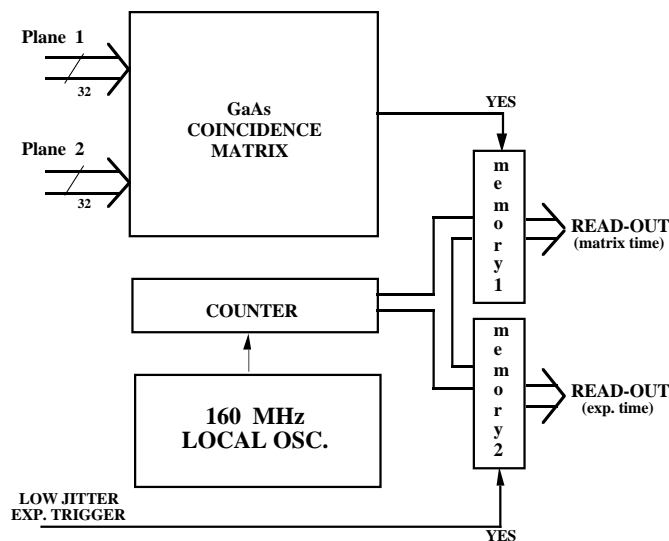


Figure 30: Bunch-crossing identification demonstrator system.




Figure 31: GaAs cross-bar coincidence matrix.

Chip Design Studies

The first RPC-based muon trigger demonstrator will use four commercial GaAs cross-bar switches which, with the addition of external logic, will allow hits in up to four trigger cells in the second (outer) plane to be put in coincidence with one in the first (inner) plane.

In the proposed LHC muon trigger, each inner cell will have to be put in coincidence with up to thirty-two cells in the outer plane, and the trigger will provide several momentum thresholds, using different coincidence patterns. Implementing such a muon trigger using discrete cross-bar switches would require an unacceptably large amount of electronics. We have therefore performed feasibility studies for an integrated coincidence array.

The typical characteristics of a coincidence array for LHC are the following: Thirty-two inputs from the first plane, 64 inputs from the second plane and three p_T thresholds; the skew on the propagation delay from the input to output should be minimised; the device should withstand the radiation levels on the detector.

The outputs from the device are an indication of which threshold level was satisfied and the input patterns that satisfied the trigger for the given threshold. The latter information can be used to more accurately determine the p_T of the muon in subsequent logic.

Our feasibility study investigated implementing such a device in several technologies of which the following two are the most promising:

- Fujitsu 0.5 μ m GaAs gate array
- TriQuint 0.7 μ m semi-custom GaAs

Our conclusion is that such a device is feasible. However the high development cost of large devices in these technologies does not allow the production of a demonstrator with all of the characteristics required for the final trigger. To reduce the cost of the demonstrator to an acceptable level the specification must be scaled down, but it must still be consistent with the aim of demonstrating the feasibility of a large, multi-threshold coincidence array suitable for LHC. Work is underway to investigate these trade-offs. For example new quotations are being sought from Fujitsu and TriQuint for arrays with the following characteristics: Eight inputs from the first plane, 16 inputs from the second plane and two p_T thresholds.

Subject to the obtaining the required funding, we hope that a demonstrator ASIC can be designed and manufactured for use in a test beam in summer 1994.

In parallel with the above studies for a track-finding ASIC, we have been investigating the feasibility of an ASIC for bunch-crossing identification using drift chambers.

4. Central Trigger Processor

Introduction

The level-1 central trigger processor (CTP) forms the final level-1 trigger decision from subtrigger data patterns (SDP) received from subtrigger (e.g. calorimeter or muon) processors. It makes this decision for each beam crossing, and may therefore be seen as a fully synchronous, 40 MHz processor. As the SDPs from different subtrigger processors are available at different times, the first task of the CTP is to equalise latencies so that all subtrigger data belonging to the same bunch crossing arrive in phase before a decision can take place. The second task of the CTP is, of course, represented by the decision logic itself, that determines which of the many different combinations of SDPs should be retained to form the final level-1 trigger decision. Another function of the CTP is to allow for filtering of high-rate and/or less interesting (but retained) combinations; this can be implemented by prescaling. For reasons of testability, it should be possible to monitor all parts of the CTP and also to count the frequency of occurrence of various features (input patterns, internal combinations, final output decision); this latter measurement can be done by means of scalers.

All these features are foreseen in the CTP architecture (Figure 32), and at a later stage will also be implemented, keeping in mind the following concerns:

- The CTP should introduce only a minimum of additional latency.

- The CTP should be fully programmable and allow for parameter entry, monitoring and testing.
- The CTP should be modular and permit some kind of scalability. It should be possible to operate several CTP “modules” in parallel in order to handle any reasonable number of SDPs. It should also be possible to put several CTP modules in series to give extra functionality at the expense of increased latency.

We plan to design the CTP with a top-down approach. This will be achieved by using the VHDL language for modelling and simulation at both high-level (architecture) and low-level (implementation of various parts). This approach will help us in the definition of the project, in the sharing of the work to be done at various stages, and in the final implementation of FPGAs and ASICs. Also, a VHDL description of the CTP could later be integrated in a wider environment and be used for the description and simulation of a complete trigger system.

Relatively complex central trigger processors are already in use in experiments at HERA and provide a useful starting point for LHC studies. The requirement of pipelined processing and pipelined readout is a significant complication compared to traditional experiments, and the short bunch-crossing period of LHC provides additional challenges.

General Description of the CTP

Current thinking on the CTP module, later on referred to as “CTP” for simplicity, is described in more detail by an internal RD27 note [25]. The block diagram of Figure 32 shows the proposed architecture for the CTP.

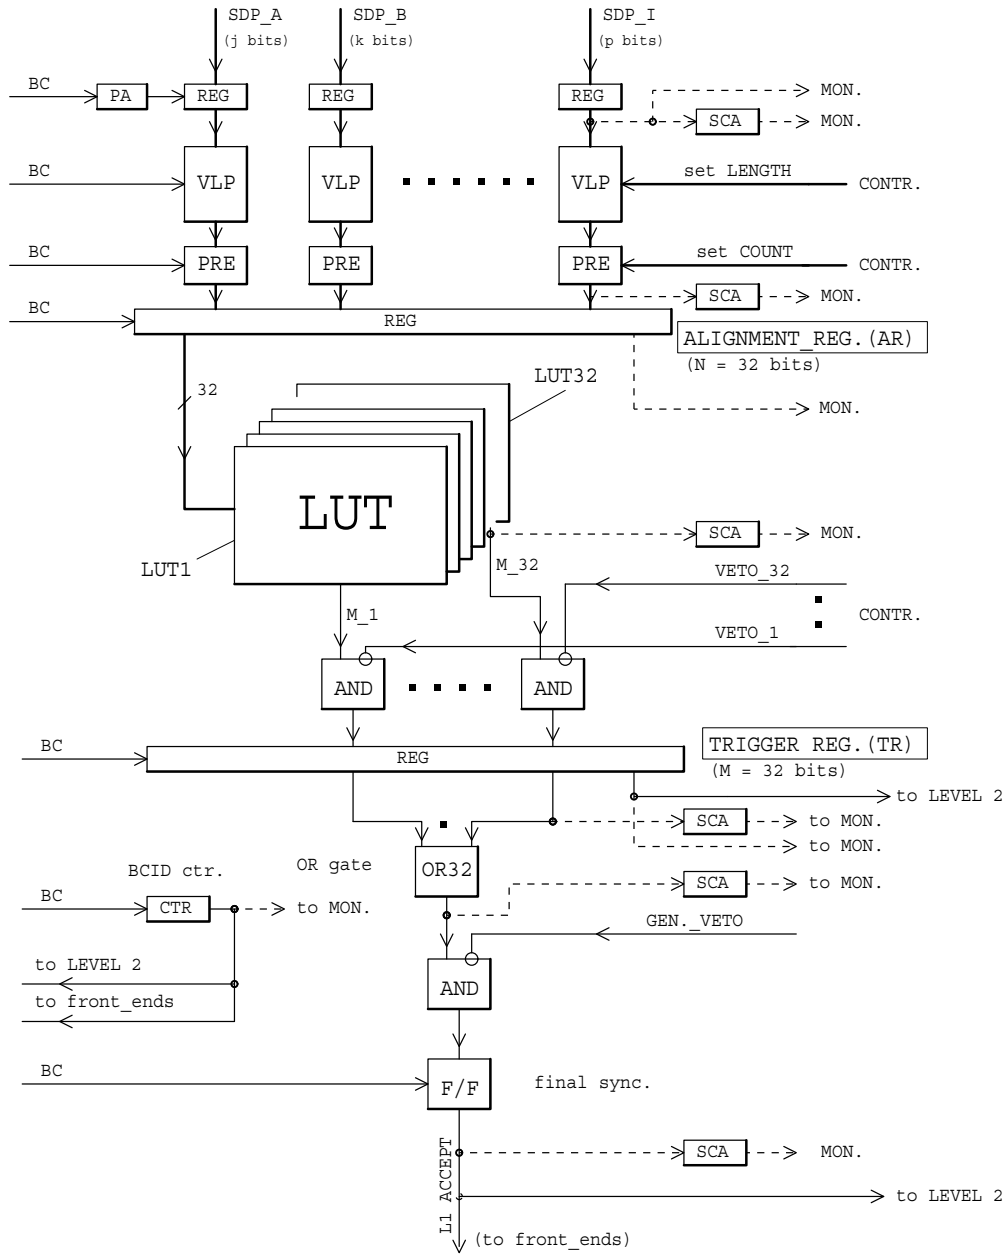
Synchronisation Circuits:

The data patterns (SDPs) coming from various subtrigger processors are for the moment assumed to be either 4 bits or 8 bits wide. The SDPs belonging to the same bunch crossing arrive at different times depending on a number of parameters such as the time-of-flight of the particles, the response time of the detectors, various electronics and cable delays, and the latency of the particular subtrigger processor. Therefore, these data patterns must be put in phase before any logic decision can take place. We plan to implement this synchronisation process by means of two circuits.

The first circuit is a “phase adjust” (PA) circuit which will tune the timing of the internal CTP clock (simply called “BC” in the diagram as it is related to the bunch crossing clock) to each individual SDP arrival time, so that the SDP data can be safely strobed into a register. It is currently felt that steps of the order of 2.5 ns should be adequate to perform this first synchronisation and that a maximum range of a little more than a bunch crossing period, e.g. 30 ns, would suffice. Such a programmable delay circuit [26] is presently being studied by RD12. Although the RD12 version is of higher performance than required for our purposes, it could perhaps be used in the CTP.

The second circuit is a variable length pipeline (VLP) which can be thought of as a FIFO with a variable, programmable depth. This circuit is required to add to the arrival time of the SDP a whole number of bunch crossing periods. A suggested maximum length for this pipeline is 24, allowing compensation of latency differences of up to 600 ns between arriving SDP patterns.

SUBTRIGGER DATA PATTERNS (from subtrigger processors)



AND.....AND gate
 BC.....BC_Clock
 BCID.....Beam crossing identification number
 CONTR.....Control system
 CTR.....Counter
 F/F.....Flip-flop
 LEVEL 2.....Level2 trigger processors
 LUT.....Look Up Table
 MON.....Monitoring system
 OR32.....32 input OR gate
 PA.....Phase adjust circuit
 PRE.....Prescaler
 REG.....Register
 SCA.....Scaler
 SDP.....Subtrigger data pattern
 GEN_VETO...Veto_Level 2
 VLP.....Variable Length Pipeline register

Figure 32: Block diagram of central trigger logic.

Decision Logic

After having passed through the synchronisation circuits, all subtrigger data patterns are in phase. They can then be strobed at an appropriate time into the alignment register (AR) which is N bits wide (e.g. $N = 32$), and be presented to the decision logic itself, which is simply represented in the diagram by a look-up table (LUT). The output of the LUT is a pattern of M bits (e.g. $M = 32$), where each bit indicates that a predetermined combination of SDPs (in coincidence, veto or “don’t care”) has been recognised, and that a candidate source signal for a level-1 trigger signal is present. These M trigger candidates are first ANDed with an equivalent number of individual veto signals generated by external control logic before being stored in the trigger register (TR). All outputs of the trigger register are ORed together, and the resulting signal may itself be ANDed with a global veto signal. The output of the AND is finally strobed at an appropriate time into the level-1-accept flip-flop.

It seems possible at this stage to implement all elements of the decision logic with commercially available parts such as RAMs, registers, gates, etc. However, the design of ASICs might still be interesting from the point of view of reliability, speed and space. It is worth noting that the system will be scalable allowing larger values for N and M than indicated above should this be necessary.

Prescalers and Scalers

Programmable prescalers are necessary to reduce trigger rates by filtering out most occurrences of selected subtrigger data patterns, allowing only a small percentage of them to be transferred to the decision logic. Although not shown in the block diagram of Figure 32, prescalers might also be required at the output of the trigger register to perform a rate reduction of the trigger source candidates. It is noted, however, that these can already be vetoed by individual control signals.

Scalers are an important ingredient of the CTP as they allow for real-time monitoring of the rates of various internal trigger patterns, which will probably prove to be essential for fine tuning of the level-1 trigger system. As shown in the block diagram, they are profusely used throughout all stages of the CTP. They are in principle read out and cleared once every LHC turn. Subsequent logic combines the scaler data, maintaining totals integrated over long periods of time.

Miscellaneous Logic

Apart from the above-mentioned circuits, some “glue logic” will be required in the CTP such as registers, gates, flip-flops, etc. Although not shown on the diagram, various obvious “hooks” have also to be included for the control and the readout of the CTP. Suitable receivers and drivers will be needed for fast and reliable signal transmission from the subtrigger processors, and to the level-2 trigger processor and the overall timing distribution system. Last but not least, as the CTP will be a key component in the trigger system, reliability will be of major concern, and testability aspects will also require extensive study.

Critical Components:

In the context of an LHC experiment, the CTP is “critical” in the sense that it fulfils an extremely crucial function. The CTP will be situated in the vicinity of the LHC machine, and therefore should be able to run at sustained data rates without any need for intervention. Speed and reliability are therefore two essential objectives of this work, cost arguments for

such a one-off system becoming comparatively negligible. With adequate circuit integration, we think that it should be possible to maximise both the speed and the reliability of the CTP.

Three critical components have been identified up to now for which an integrated solution seems appropriate – the variable length pipeline (VLP), the scaler and the prescaler.

It seems that a circuit having the functionality of a VLP is not presently commercially available, and a special study is therefore required. Figure 33 suggests a possible implementation of a cascadable 4-bit VLP slice ASIC.

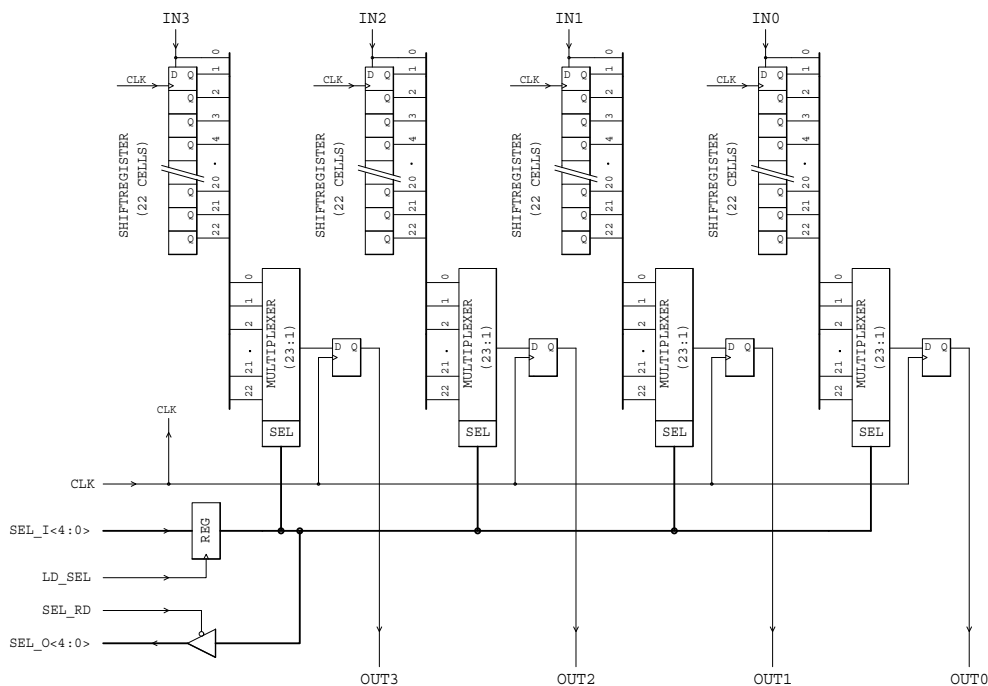


Figure 33: Four-bit variable length pipeline.

As both scalers and prescalers are required in large numbers, it seems imperative, in order to reduce interconnection problems which are a well-known potential cause for poor reliability and signal speed degradation, to integrate these circuits of the CTP. Two possible ASIC designs for a scaler circuit are shown in Figures 34 and 35: With the first option shown in Figure 34, all scaler values of the CTP have to be read out during a time window of $3.17 \mu\text{s}$, corresponding to the 127 missing bunches at the end of each LHC turn [27].

With the second option shown in Figure 35, the available time for readout is longer and actually equal to a whole LHC revolution period, i.e. $88.924 \mu\text{s}$; this is achieved by splitting the scaling from the readout functions in two different and independent stages, at the possible expense of a reduced number of scalers per chip compared to the first option.

It should be noted that with both scaler circuit options, sequential readout of all the scalers of the chip is foreseen for reasons of simplicity and connectivity. This readout process can easily be extended to several cascaded scaler chips. The scalers have a 12-bit dynamic range, allowing them to count any number of possible interesting bunch crossings within one LHC turn.

The architecture of a possible ASIC design for a prescaler is shown in Figure 36. For similar reasons as for the scaler circuit, the dynamic range of the prescaler has been fixed to

12 bits, although a lower number of bits may be sufficient. The prescale value can be reset once per turn, allowing full flexibility.

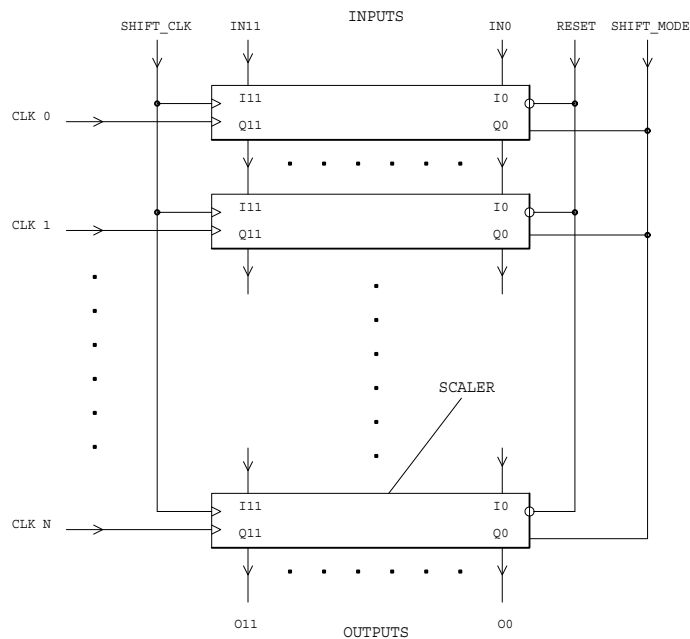


Figure 34: Scaler circuit (version 1)




Figure 35: Scaler circuit (version 2)




Figure 36: Prescaler

Future Plans

For the second year of this project, the objectives concerning the CTP have been defined as follows:

- Write a high-level VHDL specification of the CTP architecture.
- Write a low-level VHDL description of the three critical circuits – the VLP, the scaler and the prescaler.
- Implement these three critical circuits with FPGAs (e.g. XILINX) at a somewhat relaxed speed (e.g. 10 MHz).
- Build a test bench for these FPGA chip assemblies.

After the second year of this project and once the results of the implementation of the three ASIC circuits with FPGAs are available, semi- or even full-custom designs may be envisaged. By that time, new technologies will certainly offer more integration at higher speeds.

5. Timing, Trigger and Control Distribution System

Introduction

This study aims to develop a general strategy for the synchronisation of front-end electronics at LHC detectors and for the delivery of correctly phased clock, bunch-crossing identification and first-level trigger-accept signals to large numbers of channels. The current approach envisages the use of laser transmitters to broadcast the required signals over an entirely passive optical fibre distribution network [28].

In addition to broadcast signals such as the bunch-crossing clock, trigger-accept, bunch counter reset and certain test commands, the network would be used for the transmission of data to individually-addressed destinations in the timing receivers and the front-end electronics modules to which they are connected. In particular, these data could include the parameters by which system synchronisation is established by compensation of the different particle times-of-flight, detector and electronics delays and propagation times over the optical fibre paths of different lengths.

Methodology

It is anticipated that optical fibre will be deployed extensively in telecommunications subscriber loops within the next few years, spurred by the needs of future broadband residential and business services including multimedia workstation communications, high-speed colour fax, videophone, digital cable television distribution and video-on-demand. The cost of fibre-to-the-curb reached equality with copper in early 1993 and many experimental installations of fibre-to-the-home have already been made in Europe and the USA [29]. These developments give grounds for optimism that the basic components required for the implementation of an optical timing, trigger and control distribution system for large particle physics detectors could become available at affordable cost during the time frame of LHC preparation.

Furthermore, in comparison with more mature electronic transmission technologies, the rate of technical innovation in optoelectronic components and systems is high and still accelerating. Current research developments in such areas as laser gate arrays, circular-grating couplers, surface-emitting lasers, cube lasers, integrated optics, wavelength-division multiplexing, multifibre components, optical fibre amplifiers, semiconductor optical amplifiers, coherent heterodyne detectors, soliton generators or self electro-optic effect devices could at any time lead to new optoelectronic products offering system price-performance breakthroughs in this application.

In order to be open to such future developments, we have wherever possible refrained from making premature definitive technology choices which could rapidly become obsolete. However, we have found it essential to carry out some detailed design and development work in order to investigate critical technical issues and evaluate performance characteristics which are not specified by device manufacturers. But we have attempted to pursue our studies on a relatively broad front so that we have experience of the relevant system tradeoffs and are equipped to evaluate and embrace new developments as they appear.

We are, for example, continuing to study monomode and multimode single and multifibre solutions; 830 nm and 1300 nm operation; single high-power lasers and arrays of lower power devices; planar waveguide, gradient index lens and fused biconic taper couplers; silicon, InGaAs and germanium detectors in both normal and avalanche modes; a variety of network topologies, signal encoding and modulation methods and a range of technologies for timing recovery and programmable signal deskewing. We are also addressing some system aspects of the global synchronisation problem which are less dependent on the specific choice of technology for signal distribution and processing.

LHC machine interface

We are maintaining contact with our LHC machine colleagues and have contributed to the 15/25 ns bunch spacing selection studies carried out by the experiment collaborations.

While this parameter is now considered to be frozen, the LHC bunch structure, which will be subject to change at least during machine development, will be treated by the system synchronisation algorithms as variable.

CERN has already developed techniques for the transmission of stabilised radio-frequency phase references over long optical fibre links [30] and we do not anticipate that major new developments will be required to transmit the required timing signals to the central point of our distribution system in the experimental areas. Residual phase modulation at ejection from the SPS is expected to be about ± 56 ps and synchronisation errors between the SPS and LHC machines will result in a random (but constant for one batch) error at injection of about ± 100 ps. Phase oscillations will be generated when each new batch is injected and the maximum deviation, which occurs when the last batch is injected, is estimated to be 167 ps [31].

Due to the effect of transient beam loading, the equilibrium phase of the LHC bunches will not be constant along one orbit. Dedicated cavities or longitudinal kickers working on each beam separately will provide damping of the phase oscillations just after injection and may also be used to suppress any coupled bunch longitudinal instability during coast. The rms collision length for proton bunches in the LHC is expected to be 0.053 m, corresponding to an initial spread of about 180 ps relative to the equilibrium bunch phase.

At present, our general target for the allowable jitter introduced by the timing distribution system is a few hundred ps rms. Both the required resolution of the front-end electronics timing adjustment and the precision with which the timing can be tuned are a function of the time resolution of the associated subdetector. Current detector R&D suggests that in no case should a resolution of better than 100 ps be required and for many subdetectors a few ns is expected to be adequate.

Network topology

All the signals to be transmitted over each zone of the timing distribution system will be generated by an encoder at a single node in the vicinity of the central trigger processor and the basic topology of the network is a hybrid tree branching out from this point to the front-end electronics by the shortest practicable route. The configuration of this fanout and placement of the electrical-optical interfaces have significant system component implications.

At one extreme the encoder could modulate the output of a single high-power laser which is then fanned out optically by a hierarchy of tree couplers. This is an elegant configuration which features the simplicity and reliability of an entirely passive all-glass distribution network. CW lasers are available today which can generate over 2W at 1319 nm [32], sufficient to drive over 100K receivers. However, such lasers are expensive at present and they cannot be modulated directly at the required frequency. External modulators introduce a substantial insertion loss and must be used in an array to support the optical power level, so that in practice a small electrical fanout (from the encoder to the modulators) is still necessary.

The optical coupler hierarchy may be distributed or concentrated at the transmitting node. While the former configuration requires less fibre, the latter configuration could be appropriate where it is desired to use some fibres in multifibre cables or ribbon to be installed for data readout in the opposite direction. It is possible to fan out 1:1000 by a passive optical tree coupler the size of a cigarette pack.

At the other extreme a large number of low-power laser sources, equal to the number of destinations, could be installed at the transmitting node and a unique optical fibre provided to interconnect every transmitter-receiver pair. In this case the actual fanout is more onerous as it would be implemented electrically between the encoder and the many transmitters. This is an approach which could become feasible if laser gate arrays [33] or equivalent structures can be manufactured at low cost. The fabrication of arrays of laser diodes involves more process steps than for integrated circuits and problems of yield, crosstalk, thermal management and packaging have to be solved before they become commercially viable.

Between these extremes a range of network topologies could be implemented with different numbers of laser sources and the appropriate combinations of electrical and optical fanout. At present, we are working with modest directly-modulated laser diodes which are capable of driving 1000 receivers each and these must be used in arrays for larger numbers of destinations. But major technical progress may be expected in the coming years in the critical areas of high-power lasers, laser arrays and more efficient laser-fibre coupling [34]. We have therefore attempted to keep our present studies as general as possible to allow freedom of choice of the configuration which proves least expensive at the time of final procurement.

Timing signal parameters

In view of the large number of optoelectronic receivers which would be required in any configuration of the proposed optical broadcast system, it is essential that their cost be minimised. Low power consumption, small size, low mass and – for some detector locations – radiation hardness are also important considerations. Receiver power consumption is sensitive to data rate capability, while receiver costs are sensitive to both data rate capability and production volume. We have therefore implemented a scheme whereby the necessary information can be transmitted at a relatively low rate compatible with the photodetectors and PINFETs which are expected to be manufactured in the highest volumes and at the most competitive prices for the telecommunications market.

The proposed timing, trigger and control system would be able to deliver numerous broadcast and individually-addressed signals to the front-end electronics modules. At present, an important reduction in receiver cost, power consumption, size and mass is achieved by encoding these signals in such a way that they can be received by a single optoelectronic detector per destination.

Alternatively, the signals could be transmitted over a multifibre network, which would result in a small reduction in the receiver VLSI decoding logic. Considerable progress would be required in the development of ribbon fibre technology, photodetector arrays and low-power transimpedance amplifiers before a multifibre approach would actually be cheaper, smaller and lower power than a single-fibre one. But in the event that multifibre links are deployed extensively for non-broadcast tasks like data acquisition, dark fibres might be available for timing, trigger and control distribution at low incremental cost.

We are evaluating a number of signalling alternatives offering different tradeoffs between channel efficiency and synchronisation precision, and are currently using a scheme whereby two data channels are time-division multiplexed (TDM) and encoded biphasic mark at 160.32 MBaud (four times the LHC bunch-crossing rate). This is sufficiently close to the standard Sonet OC-3 (CCITT SDH STM-1) rate of 155.52 MBaud that an expanding range

of optoelectronic components produced in increasingly high volume for Sonet applications could prove appropriate.

The encoding system is self-clocking, i.e. the distribution network is only required to provide a single transmission path per destination and the timing reference is extracted from the modulated signal by clock-recovery circuitry in the receivers [35]. To implement a compatible twin-fibre variant, the data channels would still be time-division multiplexed but sent over a separate fibre from the clock.

A primary PLL phase-locks the encoder 160.32 MHz VCXO to a local clock generator which would normally receive the LHC machine signals, but which has an internal clock source to allow testing when they are not available. The PLL employs a high gain active loop filter with low offset drift which can track an input phase velocity range of 100 ppm with an rms error of less than 5 ps. In addition to driving the encoder, the clock generator provides bunch counter reset and 40.08 MHz clock outputs for the synchronisation of other systems such as the central trigger processor located in its vicinity.

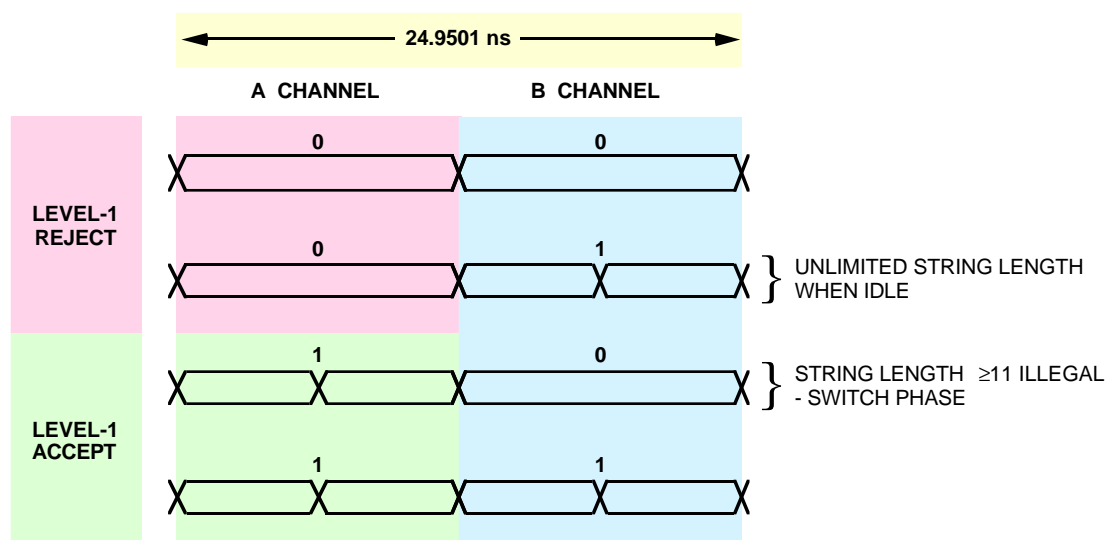


Figure 37: Timing signal encoding.


The four encoded symbols which can be transmitted in each bunch-crossing interval are shown in Figure 37. The DC offset is well bounded and timing reference signal transitions are generated at the start and finish of every such interval, which is not the case for alternative line codes such as non-return-to-zero (NRZ) inverted, scrambled-NRZ, enhanced (run-limited) NRZ, coded mark inversion, biphas level (Manchester), Miller (delay modulation, modified frequency modulation), Miller-squared or mB/nB group codes.

The “prompt” TDM A Channel is dedicated to the broadcasting of the first-level trigger-accept signal, delivering a one-bit decision for every bunch crossing. The total A Channel delay, from the trigger-accept input of the multiplexer through the encoder, RF amplifier, modulator, laser transmitter and fibre pigtail, optoelectronic receiver, amplifier, demodulator, demultiplexer, buffer and basic interconnections is less than two bunch-crossing intervals, to which must be added one bunch-crossing interval per 5m of optical fibre in the distribution path. Sequences of contiguous accept signals are allowable, so that the selection of the bunch crossings to be retained for any given event can be made by the central trigger processor [36].


The B Channel transmits broadcast 7-bit commands and individually-addressed 8-bit commands or data. The biphasic mark encoder is driven by a synchronous serializer which generates the frame format shown in Figure 38. The addressing scheme provides for up to 256 subaddresses associated with each of up to 32K timing receivers in each timing distribution group.

Since this channel may be shared by a number of command and data sources, the transmitter controller implements priority arbitration. Highest priority would normally be assigned to the timing calibration controller, to allow it to continuously scan all the timing receivers transmitting fine deskew adjustments to compensate for phase wander [37] due to changes of temperature, fibre tension, optical wavelength, signal amplitudes, voltage drifts and component aging.


Byte frame



Single-byte broadcast command



e.g. *Bunch counter reset:*



4-byte addressed command/data




Figure 38: B Channel data format.

The B Channel serializer will initially be implemented as a VMEbus slave module. When accessed by a MacVEE system, each group of up to 32K timing receivers will be directly mapped into just over 8 Mbytes of superslot address space [38] in a Macintosh computer. A single MOVE instruction referencing the appropriate address suffices to transfer data to a selected subaddress of any timing receiver in the group or to broadcast a command to all of them.

During a short interval at the end of the 3.17 μ s LHC extraction kicker gap in one of the beams, other transmissions are held off so that the bunch counter reset signal can always be broadcast with exactly the required phase. Although this signal arrives at the receivers at different times because of the different lengths of the optical fibre paths, it experiences the same propagation delay to any receiver as the trigger-accept signals and so does not require separate delay compensation.

The bunch counter reset signal is fully embedded in the data stream with the trigger-accepts. Although its duration spans 11 bunch-crossing intervals and its recognition requires a small decoding delay, its transmission is phased such that the timing receiver bunch counter outputs are cleared synchronously with the delivery to the front-end electronics of the trigger decision for bunch number 0.

With the two-channel encoding described, there is a fundamental ambiguity in the phase of the recovered clock. This is resolved automatically in the receivers by monitoring constraints on the data structure imposed by the B channel data format, which precludes strings of more than ten contiguous “10” symbols in the encoded data stream while “01” symbols, which represent the idle condition, are unrestricted. The phase corrector is continuously active, for cycle slips [39] are more sensitive to received signal-to-noise ratio than is the actual jitter of the recovered clock, and their rate of occurrence can determine the system operating margin required.

Laser transmitters

The project initially proposed [40] to extend the preliminary work undertaken in a test bench environment in RD12 [41] to the evaluation of high-power laser sources capable of broadcasting the required signals to a realistic number of front-end channels. In view of budget restrictions, the proposal to acquire a high-power CW Nd:YAG diode-pumped solid-state (DPSS) laser was abandoned in favour of work with less expensive laser diodes. While arrays of such diodes are required to equal the optical power output of a 1319 nm DPSS laser, the entry cost is lower because initial development work can be carried out using single diodes. As new devices are introduced the overall price-performance of laser diode solutions also appears increasingly promising.

Unlike laser diodes, DPSS lasers cannot be directly modulated at 160.32 MBaud. While external two-port interferometric Mach-Zehnder modulators could be used, their insertion loss is high (5-6 dB) and power-handling capacity limited. These problems can be circumvented by operating the DPSS laser as an optical amplifier of the modulated output of a second laser, but this would be an expensive approach. On the other hand when high-power laser diodes (which operate in multiple longitudinal modes) are directly modulated they exhibit mode-hopping, chirping, relaxation oscillations and pattern-sensitive turn-on delay. We are studying these characteristics as well as others, such as sensitivity to optical feedback noise [42] and temperature changes, which impact the stability of the timing signal which can be transmitted.

More efficient integrated optics modulators for broadcast applications could be based on three-port electro-optic directional couplers or X-switches since, with the encoding scheme described, the complementary output ports of these configurations can also be used to feed part of the distribution network. Such devices are normally polarization dependent, incompatible with lasers coupled to multimode fibre and tend to be subject to DC drift and thermal instability. At 1300 nm their power-handling capacity can exceed 100 mW but at 830 nm, at which much higher laser diode powers are available, light-induced index changes currently restrict them to about 5 mW. Modulator research being pursued by RD23 [43] is expected to lead to the development of devices at dramatically more attractive prices than those practised hitherto.

Laser gate arrays currently under development (e.g. by AT&T Bell Laboratories) have attractive characteristics for a topology requiring a large electric fanout. The basic laser gate is a 3-terminal “digital” Fabry-Perot device in which a small electroabsorptive modulator

[44] is integrated in the laser diode cavity to give it transistor-like amplifier characteristics. The absorber section saturates electrically, so that near 100% modulation is obtained with an uncritical direct ECL-level input. Monolithic arrays of 12 devices have been produced, butt-coupled to fibre ribbon with v -groove optical alignment, and device packaging is currently being evolved for increased array size and ease of manufacturing.

Our experiments have shown that with direct modulation of currently available laser diodes it is quite feasible to broadcast to groups of 1024 channels per transmitter through 100 m of 50/125 μm graded-index optical fibre and two levels of 1:32 passive optical tree coupler. For the two-channel system transmitting PRBS data over this network at 267.2 MBaud, an overall rms jitter from the LHC clock input to the remote recovered timing reference output of 35 to 100 ps has been achieved, according to the encoding system and type of optoelectronic receiver used. The contribution from primary PLL and encoder jitter is about 9 ps and from fibre dispersion about 12 ps at 1315 nm.

Photodetectors

Although it is foreseen that LHC detectors will have many optical links for data readout and monitoring purposes, the timing distribution network may be one of the few systems requiring large numbers of optoelectronic receivers located at or within the detector. The photodetectors used should have high optical signal sensitivity but low sensitivity to ionising and neutron irradiation. They should have fast rise times at a low reverse bias voltage, preferably only 3.3 V, and should ideally permit monolithic integration with the VLSI process to be used for the timing receiver logic. No such ideal devices have yet been fabricated.

Si photodiodes have been successfully integrated in a standard VLSI BiCMOS technology [45]. But to achieve high-speed operation at a low bias voltage, an intrinsic region of submicron depth has to be used, resulting in very low responsivity. Si diodes are not usable at 1300 nm because the photon energy at that wavelength is less than the bandgap energy.

InGaAs PIN diodes are superior to Si photodiodes in most technical characteristics including radiation hardness [46]. Even at 830 nm, at which the light absorption coefficient of InGaAs is relatively high, RD12 component tests have shown that some devices outperform Si because the bandgap of the material can have different values depending on its composition. Although the material is more expensive, InGaAs components are becoming increasingly competitive in price as they enter mainstream production.

Photodetector, preamplifier and timing receiver logic integration options are currently being studied by CERN ECP Microelectronics Group. While it may be possible to bond InGaAs PIN diodes rather directly to CMOS signal-processing circuits, optical fibre coupling issues tend to favour the use of external connectorized diodes or PINFETs closely coupled to the VLSI package. "Smart connectors", which integrate all the standard optoelectronic receiver functions, could well appear during the time frame of our study.

Timing receivers

While the current timing receiver prototypes employ ECLIPS logic, the specifications are being defined for a demonstrator ASIC which would have much lower power consumption and smaller physical size. The preparatory study is being carried out in collaboration with CERN ECP Microelectronics Group and Padua University/INFN (RD12). LPNHE (University of Paris VI & VII) members of RD16 are also participating with a view to developing a compatible local controller for FERMI modules. While GaAs technology has

been considered, the preferred approach is to implement a first design in conventional submicron CMOS which could be readily migrated to the standard radiation-hard silicon process which CERN hopes it will be possible to select for front-end electronics in due course.

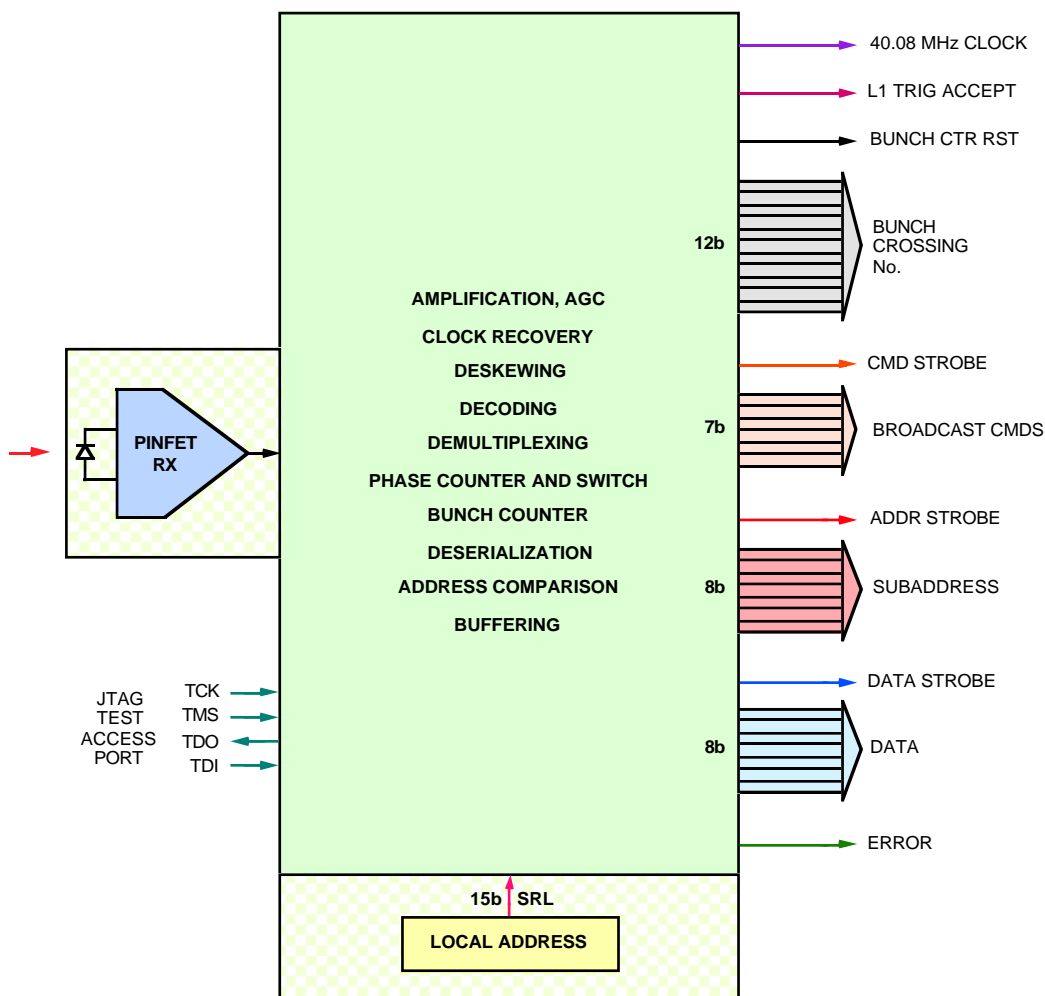


Figure 39: Timing receiver ASIC.

In view of the wide range of front-end electronics technical requirements, physical implementations and operating constraints it is somewhat improbable that the compromises inherent in a universal timing and control ASIC design would prove acceptable for all the different subdetectors of even one LHC experiment. But the development of a general-purpose demonstrator ASIC in the context of DRDC R&D would allow the technical aspects of timing, trigger and control distribution to be pursued further, to the benefit even of those groups who will prefer to develop their own custom solutions.

Figure 39 illustrates the signals which could be provided by such an ASIC from its single optical input. The receiver incorporates a multiphase clock generator [47], programmable delays and programmable-length shift registers for signal deskewing. These functions, as well as broadcast command generation and bunch counter reset, are controlled by the data transmitted over the B channel and subaddresses which are not used internally are made available for the selection of external destinations.

The bunch counter reset signal is also provided for external use, but a 12-bit counter is integrated on-chip which generates a unique bunch crossing number synchronously with the corresponding first-level trigger decision. All 3564 “potential” bunch crossings per orbit are numbered, since the actual LHC bunch structure can vary with the mode of operation.

It is assumed that the optical fibre path lengths in the timing distribution system will be dictated by installation convenience alone and no attempt will be made to cut them to particular propagation delay values. The bunch counter reset, and non-periodic signals such as the trigger decision and broadcast commands, will be deskewed by the timing receiver over a maximum 12-bit range; 4 bits for the number of bunch-crossing intervals and up to 8 bits for the phase within an interval.

The delay compensation range of 16 bunch-crossing intervals (total 399 ns) allows a substantial margin beyond the possible maximum variation due to differences in time-of-flight and optical fibre path length, so that the timing calibration controller will not have to deal with deskew register rollover. The periodic clock output, which is derived from a local charge pump PLL [48] with voltage-controlled delay elements [49], is of course only deskewed over the bunch-crossing interval.

The broadcast command outputs share the fine clock phase adjustment but have an independent deskew register for the number of bunch-crossing intervals. This is to allow them to be used for the generation of test and calibration signals having different delay compensations (excluding time-of-flight but including test signal latency, for example) without having to reload the deskew parameters used for normal running. The same deskew is applied to all 7 output bits, which may therefore be encoded where a large number of commands is required or used individually otherwise.

The local fanout of the 40.08 MHz clock, trigger-accept and other signals generated by the timing receiver ASIC will be implemented externally as the requirements for different front-end electronics systems can vary considerably. In cases where each timing receiver is associated with a group of many channels, adaptive clock manager [50] or inter-ASIC clock distribution techniques [51] may be used to compensate for differential skew introduced by active device or layout asymmetry in the clock fanout. It is assumed that such differential compensation will be stable and not require individual attention from the main timing calibration controller.

Some front-end electronics modules, such as FERMI, may require ADC control signals at multiples of 40.08 MHz, or multiphase clocks for internal logic sequencing purposes. Such modules would incorporate their own phase-locked ring oscillators to derive the required signals from the 40.08 MHz clock provided by the timing receiver ASIC.

Each timing receiver ASIC requires a unique 15-bit address. In the absence of a reliable non-volatile radiation-hard on-chip PROM technology, this could be provided by an external low-cost factory-lasered “silicon serial number” chip whose output is read into the ASIC address comparison register during initialisation. Such devices are available with zero standby power consumption.

Synchronisation adjustment

An important extension of our work will be the detailed study of control procedures for the adjustment of the synchronisation of the front-end channels by the tuning of the timing receiver deskew circuits. Different approaches are required for the adjustment of the coarse deskews (which delay the bunch counter resets affecting the bunch crossing number

assignment and delivery of the associated trigger decision) and for the fine deskews (which determine the phase of the delivered clock within the bunch crossing period).

For the coarse deskew adjustment, crosscorrelation of the complement of the SPS batch structure with the event occurrence pattern appears promising. The latter is a known function of the bunch structures of the two LHC beams, and their phasing relative to the interaction region, modulated by random event statistics. Since the computation only requires the accumulation of shifted and gated Boolean vectors, it is particularly straightforward to implement at high speed and could be performed for many channels in parallel.

The setup procedures for the global phase of the bunch counter reset relative to the beam orbits, and subsequently for all the coarse deskews, would be run before any data-taking, when we have free control of the trigger to modulate the data flow. The coarse deskews should require relatively little attention subsequently, except when cabling or equipment changes are made or a fine deskew is found to require tuning to one extremity of its range. The crosscorrelation synchronisation check could be maintained for monitoring purposes during normal data acquisition with a physics trigger.

Critical issues for this simple approach are the time required to accumulate adequate statistics for low-occupancy channels and the degradation of the “gap quality” caused by protons in RF buckets in the gaps (which in reality will not be completely empty) and bias introduced by any serious asymmetry in the populations of the bunches preceding and following them.

The adjustment of the fine deskews is a more challenging task demanding considerably more study. Since these deskews should compensate for phase wander which may be expected to occur slowly during normal running, it would be very desirable that the timing calibration system be able to perform as an extremum controller [52] effecting fine tuning adjustments continuously during data-taking, but without introducing significant spurious perturbations during optimisation. For the tracking detectors, high quality track fitting through multiple detector elements may be considered to be the ultimate criterion of front-end sampling phase precision because it is affected by all the variable components of delay.

Both the procedure and requirements for the fine tuning will be very dependent on the characteristics of each subdetector and its front-end signal processing electronics. It has been observed [53] that even calorimeters having a response which spreads over several bunch crossing intervals are sensitive to phase changes between particle arrival time and ADC sampling clock of 1 ns or less. On the other hand trainable FIR - order statistic hybrid filter algorithms which are being developed [54] for processing such signals have the property of compensating to a considerable degree for variations in phase between the detector signals and the sampling clock. Additional studies in collaboration with detector R&D groups and teams working on front-end electronics and their signal processing algorithms will be required to resolve these issues.

Simulation

Computer simulation of the timing system can be a useful tool for refining architectural concepts, evaluating the effects of parameter variations and analysing the performance of the synchronisation adjustment algorithms. For the preliminary phase we have employed flexible but unsophisticated tools such as Excel, Extend and LabVIEW for our conceptual studies. We are collaborating with members of CERN ECP Readout Architecture Group who are currently evaluating more powerful simulation tools such as Foresight, which it is

expected will eventually be able to generate VHDL and permit the comprehensive modelling of complete front-end data acquisition and timing systems.

Future plans

During the second year of the project we plan to pursue our evaluation of new laser diodes, laser gate arrays, or other appropriate optoelectronic devices which may come to market. We shall include devices coupled to monomode fibre which may be called for in areas of very high radiation level. If these trials are successful, we shall integrate them in the subsystem prototype being developed in RD12. We shall also continue our study of timing stability and synchronisation procedure issues, complete the conversion of our system from 15 to 25 ns operation and develop a first prototype of the B Channel VMEbus synchronous serializer.

We plan to advance our timing system simulation studies and refine our definitions of the external interfaces, transmission protocols and receiver architecture. We shall continue collaboration with our microelectronics colleagues and, if funding permits, proceed towards a first implementation of the timing receiver ASIC. We shall study the possibility of supplying a small-scale prototype timing, trigger and control distribution system to RD6 for their beam tests of an element of the TRD/Tracker equipped with ASIC electronics.

We shall strengthen our links with other front-end electronics teams and the LHC machine specialists studying bunch phase modulation and other synchronisation matters. In order to catalyse more discussion of timing issues among specialists in different R&D groups, an informal review document [55] will be produced which can be regularly expanded and updated as our ideas evolve.

6. Level-1 – Level-2 Interface

Introduction

In order to make substantial reductions in the trigger rate after level 1 (level-1 rate ~ 100 kHz) it is necessary to use the full granularity of the detector and complex algorithms. The first stage of the level-2 procedure, known as feature extraction, attempts to reduce basic data from a subdetector in the form of track hits or cluster pulse heights into a few words characterising the track or cluster. After reducing basic data to features, level 2 attempts to combine features into objects by associating the features from different sub-detectors and then to combine objects to form the event. Feature extraction is a local operation and the data input to the processor need be connected only locally. Global connection is needed to form the event trigger.

If level 1 is used to indicate the position of tracks or energy clusters and only these *regions of interest* (ROIs) are analysed at level 2 then considerable benefits in data transfer and processing rates accrue. In this case, level 2 may use either special or general purpose processors (or a combination of both). Transfers of data to “local processors” take place at a much lower rate (~ 1 kHz) and longer access and processing times (i.e. longer latency) can be accommodated. Alternatively, the information from ROIs may be routed to a few low latency feature extraction machines.

The use of ROIs involves the transfer of information at several thresholds from level 1 and not just information for regions above the trigger thresholds. As level 1 will require information at several thresholds for compound triggers, production of the ROI data for

Level 2 does not necessarily imply significant extra hardware inside level 1. This multi-threshold information may be particularly important for level-2 topology triggers.

In comparison with the use of ROIs, the analysis of all the data from the detector at full granularity and at a rate of up to 100 kHz is unlikely to be an economic proposition.

Level-1/level-2 communications problems addressed in RD27

Two models are proposed for level-2 operations – asynchronous and systolic – and, apart from a tighter time constraint for the systolic option, both require similar ROI information. In both cases, a level-1 trigger initiates the transfer of data from the front end of each sub-detector into a set of level-2 buffers where the data are stored until after the level-2 decision. In the “asynchronous” option, the data for feature extraction are transferred from the buffers to “local processors”; for the systolic option, the data are intercepted before the level-2 buffers and “routed” to the feature extractors.

In either case, a mechanism is needed to select the correct data for analysis, to transfer it to the feature extraction processors and to initiate the process.

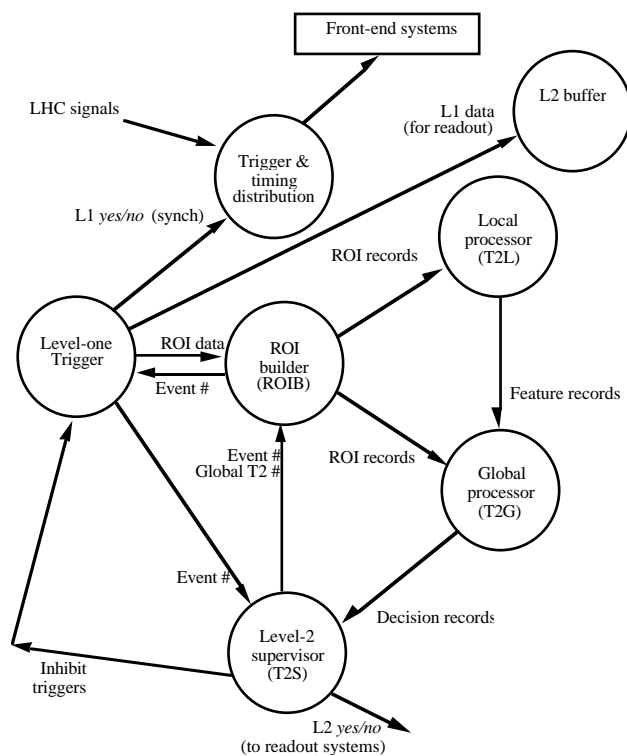


Figure 40: Outline of level-1/level-2 system.

As ROI information originates from different components of the level-1 trigger system, it must be collected, translated to suit the level-2 processors, and sent to the processors concerned. In addition, a summary of all ROIs has to be sent to the global level-2 processor allocated to handle a particular event. The major problem to be addressed is to determine the information required at level 2, to design a mechanism to generate that information from data available in the level-1 system (the ROI builder) and to distribute that information to the requisite places. This is illustrated in outline in Figure 40.

A more detailed description of the interactions between level-1 and level-2 are given in an RD27 note [56]. This shows that other problems on level-1/level-2 communications which have to be addressed at a later stage in the project are the allocation of level-2 global

processors, error reporting and handling in the level-2 system and a “throttle” control to prevent the overflow of buffers at level 2. At this stage we are concerned with the ROI builder as we consider it the most challenging of the inter-communication tasks.

ROI data definition for level 2

Information for the ROI builder comes from four sources: the muon system; the calorimeter system, the central trigger system and the level-2 trigger supervisor. These must be combined and processed in the ROI builder.

ROI information from the muon system

Muon information is sent on receipt of level-1 trigger accept and is buffered at the input to the ROI builder by a FIFO. The data format is fixed and compressed into a long word which contains information on up to four regions of interest. The information for an ROI is the $\eta - \phi$ of the track found, its momentum threshold pattern and sign. Additional information in the data block is the event number and bunch-crossing identifier (BCID) which is used to synchronise the data with the other sources and for error checking.

ROI information from calorimeter trigger systems

The information available directly from the calorimeter system is in the form of hit maps at different threshold levels. These maps will need to be declustered before combining with other information. The declustered information will be passed to the ROI builder together with the event number and BCID.

ROI information from central trigger system

This provides the event trigger mask together with the event number and BCID.

ROI information from the level-2 trigger supervisor

This provides the level-2 global processor number together with the event number and BCID.

Information required by level-2 local processors

The following information can be provided to each local processor which has to operate on a particular event: event identifier; global processor address; trigger type; level-1 feature type; information (threshold pattern, energy sum); $\eta - \phi$ and local processor address.

It is not clear whether all the information listed above is required by the level-2 local processors but at this stage we should aim to provide the full set.

Information required by level-2 global processors

The global processor can be provided with the same information as the local processors but compacted into one record with common information appearing only once.

Present demonstrator programme

In order to develop the ROI builder we intend to start with a basic demonstrator and to expand it towards the final system.

As a first stage we looked to see if a programmable machine could be suitable for the purpose. This required a data source and a VME-CPU where the data source provided data

in the appropriate form and the processor had to perform the necessary manipulations to format the data into ROIs suitable for transmission to level 2.

As the ROI data from the muon system is relatively well defined, the initial tests were made using these data. The basic parameters in the test are that triggers, and therefore muon ROI data, arrive randomly at a mean rate of 100 kHz and that each trigger can contain up to four ROIs. The additional information which must be correlated with the muon information is the trigger mask data from the central processor and the event global processor address from the level-2 trigger supervisor.

Using an arbitrary function generator (AFG) as the data source and an existing DSP board containing a Motorola 56001, measurements have been made of the time to convert muon data to ROI data. The muon input data were generated on a Monte Carlo basis within reasonable boundary values and downloaded to the function generator. The generator supplies data in the form of twelve 16-bit words to a FIFO at the front end of the DSP board at an input rate of 30 Mwords/s limited by the FIFO. A block diagram of the ROI builder demonstrator is shown in Figure 41.

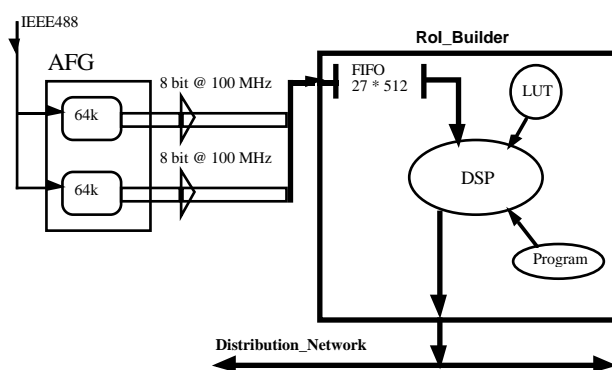


Figure 41: ROI builder demonstrator.

The CPU performs the following tasks:

- data consistency check and initiates error mechanism if mismatch
- copy global processor address to all ROIs
- copy trigger mask to all ROIs
- unpack muon ROI information
- form ROI records for local processors
- form ROI record for global processor

A program has been written in DSP56001 assembler code to perform these tasks. The start time of an event is taken as the arrival time of the first word at the input of the DSP board; the stop time is the arrival of the last word of the global processor ROI record in the VSB output register.

Measurements with a logic analyser show that the duration of the procedure is about 16 μ s somewhat in excess of the 10 μ s limit implied by a 100 kHz trigger rate. Nevertheless, this first result indicates to us that the task of ROI building could be handled with programmable devices where the flexibility of a programmable algorithm could be an advantage.

Future work

The next stage is to implement ROI building for a bit-parallel calorimeter trigger where initially we shall use declustered data from the calorimeter. Following this we shall develop a declustering algorithm to operate directly on data from the clustering crates in the bit-parallel system.

At the same time we shall investigate up-to-date processor hardware on new demonstrator modules. An architecture including direct data links opens the possibility of parallelisation where the task allows it. We intend to develop a design adequate to master ROI building in a full experiment, i.e. combining the calorimeter and muon systems and implementing the connection to the level-2 processor system.

7. Conclusions

During the first year of the RD27 project we have developed a detailed outline design of a first-level trigger system, including its interaction with the level-2 trigger and the front-end systems.

We have made detailed design studies for two calorimeter trigger processors, using physics simulation to evaluate algorithms. A first-prototype calorimeter trigger processor has been successfully evaluated in real time in test beams with the Accordion and TGT calorimeter prototypes. The design studies suggest that the calorimeter trigger electronics can be made reasonably compact, occupying only a few electronic crates. However, further development work is required to master technical issues such as data transmission into the trigger processor.

We have made detailed design studies for a muon trigger based on RPC detectors and a first-prototype trigger system is in preparation for beam tests with RD5 later this year. We are also studying the feasibility of a level-1 muon trigger based on drift chambers.

We have made an outline design for the central trigger processor which we plan to study in more detail in the coming year. In addition to modelling using VHDL, we will evaluate designs for key components using field-programmable gate arrays. An implementation using custom circuits may be appropriate at a later stage.

Jointly with RD12 we have studied systems for timing, trigger and control distribution and demonstrated the capability of time-division multiplexing the required signals and broadcasting them optically to 1024 receivers per laser transmitter with a recovered clock jitter of less than 100 ps. We plan to extend this work and, in conjunction with CERN ECP Microelectronics Group, to develop a demonstrator timing receiver ASIC which could be used in detector R&D projects.

Suggested milestones for the second year of the RD27 project are listed below.

8. Suggested Milestones

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

- more detailed design studies for a first-level trigger system
- further design studies of calorimeter triggers and beam tests of prototypes; detailed studies of data-transfer techniques using high-speed links and zero suppression

- further design studies of muon triggers, addressing the issue of bunch-crossing identification in a large-area system; beam tests of a muon trigger prototype
- continued development of system components for timing, trigger and control distribution (with RD12). Specification of ASIC implementation of timing receiver in collaboration with CERN ECP/MIC and Padua University/INFN; if funding permits, fabrication of demonstrator chips
- design studies for a central level-1 trigger processor and development of demonstrator prototypes based on field-programmable gate arrays
- design studies for the level-1 – level-2 trigger interface and development of demonstrator prototypes

9. 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	•				•	
Stockholm	•			•		•

Areas of responsibility of groups.

10. Funding

In order to continue our initial programme of work we will require funding from CERN at the level of 100 kSFR in 1994. This covers the CERN activities both on the level-1 central trigger processor development, and on component and system prototyping for timing, trigger and control distribution. In making this request, we have taken into account the fact that work on the central trigger processor will move beyond the design phase during 1994, and require the construction of demonstrator prototypes.

In addition, for the development of a timing, trigger and control receiver ASIC, which will be done jointly with RD12 and the ECP microelectronics group, we require 100 kSFR.

11. Publications from RD27

- N. Ellis, Level-1 and Level-2 triggering at LHC. Proc. CHEP92, CERN 92-07 pp. 51-56.
- N. Ellis, I. Fensome, J. Garvey, P. Jovanovic, R. Staley, A. Watson, E. Eisenhandler, C.N.P. Gee, A. Gillman, R. Hatley, V. Perera, S. Quinton, 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, Conf. Record IEEE Nuclear Science Symposium, Orlando, USA, 25-31 October 1992, pp. 492-494.
- N. Ellis, Level-1 and level-2 triggering at LHC. Presented at 3rd Int. Conf. on Calorimetry in HEP, Corpus Christi, Texas, USA, 29 September - 2 October 1992.
- E. Eisenhandler, N. Gee, A. Gillman, V. Perera, S. Quinton, N. Ellis, I. Femsome, J. Garvey, P. Jovanovic, R. Staley, A. Watson, 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 : "Use of GaAs circuits for first level muon triggering at LHC", presented by S. Veneziano at the 5th Topical Seminar on Experimental Apparatus For High Energy Particle Physics and Astrophysics, S. Miniato, Tuscany, 26-30 April, 1993, to be published in NIM.

12. RD27 notes

- RD27 note 1, October 1992: N. Ellis, Level-1 and Level-2 triggering at LHC.
- RD27 note 2, October 1992: J. Garvey et al., A calorimeter-based level-one electromagnetic cluster trigger for LHC.
- RD27 note 3, December 1992: B.G. Taylor, Multichannel optical fibre distribution system for LHC detector timing and control signals.
- RD27 note 4, October 1992: N. Ellis, Level-1 and Level-2 triggering at LHC.
- RD27 note 5, November 1992: V. Perera et al, First level calorimeter trigger system for the LHC.
- RD27 note 6, January 1993: C. Bohm et al, A bit-serial first-level calorimeter trigger for an LHC detector.
- RD27 note 7, March 1993: N. Ellis, Ideas for a local/global level-2 trigger system.
- RD27 note 8, April 1993: V. Perera, A first-level calorimeter trigger processor for the Large Hadron Collider.
- RD27 note 9, June 1993: J-P. Vanuxem and H. Wendler, Central trigger processor for the first-level trigger for LHC experiments.
- RD27 note 10, June 1993: V. Perera, Crate/Readout Controller (ROC).
- RD27 note 11, July 1993: I. Brawn, Preliminary investigations into bunch-crossing identification for the level-1 trigger.
- RD27 note 12, August 1993: R.E. Carney, The Level-1 Bit-Parallel Electromagnetic Trigger Processor. CERN Beam Tests - November 1992.
- RD27 note 13, August 1993: A. Watson, Physics Simulation Studies of the First-Level Electron/photon Trigger.

Acknowledgements

We express our thanks to ATLAS and CMS, and to other R&D groups with whom we have held useful discussions throughout the first year of the RD27 project. In particular, we thank RD11 (EAST) and RD16 (FERMI) for discussions on the interaction between the level-1 trigger with other parts of the trigger and DAQ system. We thank RD12 for providing us with test-bench modules and for their collaboration on the timing, trigger and control distribution system. We are particularly grateful to RD3 and to RD33 for allowing us to perform tests of our prototype processor using signals from their calorimeters, and for the extensive help we received for the tests.

We are grateful to Scott Kolya from the H1 experiment for providing us with the F1001 FADC module, which was used in the November 1992 tests with the Accordion calorimeter.

References

- [1] RD27 proposal, CERN/DRDC/92-17.
- [2] Minutes of Research Board 30 June 1992. Minutes of DRDC meeting 4 June 1992.
- [3] ATLAS letter of intent, CERN/LHCC/92-4.
- [4] CMS letter of intent, CERN/LHCC/92-3.
- [5] SDC technical design report.
- [6] GEM technical design report.
- [7] FERMI status report, CERN/DRDC/93-21.
- [8] ECFA LHC workshop, Aachen, October 1990, CERN 90-10.
- [9] ATLAS internal note, DAQ-NO-005.
- [10] Watson, A., "Calorimeter trigger simulation studies", RD27 note 13.
- [11] GEANT, CERN program library long writeup W5013.
- [12] ATLAS detector model – see [3, 9].
- [13] "Analysis of test-beam data (April 1993)", RD27 note in preparation.
- [14] Perera, V., "Data on the LHC Electromagnetic Cluster-Finding ASIC (RAL 114)", RAL, Feb 1992.
- [15] "Spider writeup", CN division, CERN.
- [16] RD3 proposal, CERN/DRDC/90-31.
- [17] Carney, R.E., "Analysis of test-beam data (November 1992)", RD27 note 12.
- [18] Brawn, I., "Preliminary investigations into bunch-crossing identification for the level-1 trigger", RD27 note 11.
- [19] RD33 proposal, CERN/DRDC/93-2.
- [20] Perera, V., "A first-level calorimeter trigger processor for the Large Hadron Collider", RD27 note 8.
- [21] Bohm, C. et al, "A bit-serial first-level calorimeter trigger for an LHC detector", RD27 note 6.

- [22] Nisati, A., “Level-1 muon trigger in the ATLAS experiment”, ATLAS note in preparation.
- [23] Ferrari, A., Private communication.
- [24] Ahmed, T. et al, “A pipelined first-level forward-muon drift-chamber trigger for H1”, Proc. CHEP92, CERN 92-07 , pp. 218-221.
- Dowdell, J. et al, “Fast pipelined trigger processors for the forward muon chambers of the H1 experiment on the Hadron Electron Ring Accelerator (HERA) at DESY, Hamburg”, presented at the IEEE Nuclear Science Symposium, Santa Fe, New Mexico, USA, November 2-9, 1991.
- [25] Vanuxem, J.-P. and Wendler, H., “Central trigger processor for the first-level trigger for LHC experiments”, RD27 note 9.
- [26] Cittolin, S., Private communication, June 1993.
- [27] Potter, K., Private communication.
- [28] Taylor, B.G., “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.
- [29] Shumate, P.W., “Fiber to the Home”, Tutorial presented at IEEE International Conference on Communications (ICC '93), Geneva, Switzerland, 23-26 May 1993.
- [30] Peschardt, E. and Sladen, J.P.H., “Transmission of a Stabilised RF Phase Reference over a Monomode Fibre-optic Link”, Electron. Lett., Vol. 22, 1986, pp. 868-869.
- [31] Boussard, D. and Onillon, E., “Damping of Phase Errors at Injection in the LHC”, presented at 1993 Particle Accelerator Conference, Washington DC, USA, 17-20 May 1993, CERN SL/93-20 (DI).
- [32] Baer, T.M. et al, “Performance of Diode-Pumped Nd:YAG and Nd:YLF Lasers in a Tightly Folded Resonator Configuration”, IEEE J. Quantum Electronics, Vol. 28, 1992, pp. 1131-1138.
- [33] Nordin, R.A. et al, “A Systems Perspective on Digital Interconnection Technology”, IEEE/OSA J. Lightwave Technology, Vol. 10, 1992, pp. 811-827.
- [34] Bell, T.E. et al, “Innovations: Parallel, circular beam shines from solid-state laser”, IEEE Spectrum, Vol. 30, 1993, p. 8.
- [35] Takasaki, Y., “Digital Transmission Design and Jitter Analysis”, Artech House, 1991, Chapter 4 - Clock Recovery.
- [36] Vanuxem, J.-P. and Wendler, H., “Central Trigger Processor for the First-Level Trigger for LHC Experiments”, RD27 note 9, June 1993.
- [37] Trischitta, P.R. and Varma, E.L., “Jitter in Digital Transmission Systems”, Artech House, 1989, Chapter 8.1 - Wander in Fiber Optic Transmission Systems.
- [38] Taylor, B.G., “Developing for the Macintosh NuBus”, Proc. Eurobus/89 Conference, London, 5-6 September 1989, pp. 143-175.
- [39] Ascheid, G. and Meyr, H., “Cycle Slips in Phase-Locked Loops: A Tutorial Survey”, IEEE Trans. Communications, Vol. 30, 1982, pp. 2228-2241.

- [40] Ellis, N. et al, "First-Level Trigger Systems for LHC Experiments", CERN/DRDC 92-17, 11 March 1992.
- [41] Cittolin, S. et al, "Status Report on the RD12 Project", CERN/DRDC 93-22, 5 May 1993.
- [42] Langley, L.N. and Shore, K.A., "The Effect of External Optical Feedback on Timing Jitter in Modulated Laser Diodes", IEEE/OSA J. Lightwave Technology, Vol. 11, 1993, pp. 434-441.
- [43] Stefanini, G. et al, "Optoelectronic Analogue Signal Transfer for LHC Detectors", CERN/DRDC 91-41, 9 October 1991.
- [44] O'Gorman, J. et al, "Dynamic and Static Response of Multielectrode Lasers", Appl. Phys. Lett., Vol. 57, 1990, pp. 968-970.
- [45] Lim, P. J-W., "A 3.3-V Monolithic Photodetector/CMOS Preamplifier for 531 Mb/s Optical Data Link Applications", to be published in Conf. Record IEEE International Solid-State Circuits Conference, San Francisco, USA, 24-26 February 1993.
- [46] Leskovar, B., "Radiation Effects on Optical Data Transmission Systems", IEEE Trans. Nuclear Science, Vol. 36, 1989, pp. 543-551.
- [47] Meng, T.H., "Synchronization Design for Digital Systems", Kluwer, 1991, Chapter 6.4.1 - Mesochronous Interconnect.
- [48] Meyr, H. and Ascheid, G., "Synchronization in Digital Communications", Vol. 1, John Wiley, 1990, Chapter 2.7 - Charge Pump Phase-locked Loops.
- [49] Christiansen, J. et al, "16-Channel TDC Macro", CERN/ECP-MIC Report, 1993.
- [50] Goodenough, F., "Analog Techniques Deskew Clocks, Multiply Sines", Electronic Design, Vol. 41, 13 May 1993, pp. 39-48.
- [51] "ASIC Clock Distribution using a Phase-Locked-Loop (PLL)", Motorola Semiconductor Application Note AN1509/D, December 1992.
- [52] Wellstead, P.E. and Scotson, P.G., "Self-tuning Extremum Control", IEE Proc., Pt. D, Vol. 137, 1990, pp. 165-175.
- [53] Watson, A., "Preliminary results from test with RD-3" presentation transparencies of RD-27 Meeting, CERN, 7 June 1993.
- [54] Inkinen, S.J. and Niittylahti, J., "Trainable FIR - Order Statistic Hybrid Filters", submitted to IEEE Trans. Circuits and Systems-II: Analog and Digital Signal Processing, April 1993; updated version 1.2, 26 July 1993.
- [55] Taylor, B.G., "Timing, Trigger and Control Distribution for LHC Detectors", working review document, CERN/ECP-RA.
- [56] Ellis, N., "Ideas for a local/global level-2 trigger system", RD27 note 7.