FIRST-LEVEL TRIGGER SYSTEMS FOR LHC EXPERIMENTS

N. Ellis*), I. Fensome, J. Garvey, P. Jovanovic, R. Staley, A. Watson

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

B.G. Taylor, J.–P. Vanuxem

CERN, Geneva, Switzerland

P. Hanke, M. Keller, E.-E. Kluge, K. Schmitt, M. Wunsch

Institut für Hochenergiephysik der Universität Heidelberg, Germany

J. Fent, W. Froechtenicht, C. Kiesling, H. Oberlack, P. Schacht

Max-Planck-Institut für Physik, Munich, Germany 1)

E. Eisenhandler, M. Landon, G. Thompson

Queen Mary and Westfield College, University of London, UK

V. Perera, S. Quinton

Rutherford Appleton Laboratory, UK

J. Carter, B. Green, J. Strong

Royal Holloway and Bedford New College, University of London, UK

A. Nisati, E. Petrolo, M. Torelli, S. Veneziano, L. Zanello

Intituto Nazionale di Fisica Nucleare and Universita La Sapienza, Rome, Italy

G. Appelquist, C. Bohm, B. Hovander, N. Yamdagni

Stockholm University, Sweden

(*) Spokesman

(1) Subject to approval by the institute

Abstract

We propose to carry out a broad-based programme of R&D on level-1 trigger systems for LHC experiments. We will consider the overall level-1 system which coordinates different subtriggers and which interacts with the front-end electronics and with the level-2 system. Careful attention will be paid to systems aspects and problems of synchronization within the pipelined processor system. Trigger algorithms for selecting events with high-p_T electrons, photons, muons, jets, and large missing E_T will be evaluated by physics simulation studies. We will study possible implementations of such trigger algorithms in fast electronics by making conceptual design studies and using behavioural simulation models. For critical areas, more detailed design studies will be made, and prototypes of some key elements will be constructed and tested.

The proposed R&D project builds on existing studies and will complement other R&D projects already funded by the DRDC.

Introduction

Triggering at the LHC will be an extremely difficult task given the very high luminosity of the machine and the short bunch-crossing period. The cross-section for high-p_T jet production, which is a background to many interesting physics processes, is large, while that for new physics is very small. Hence, the trigger must combine unprecedented rejection power against high-rate processes while retaining good efficiency for the events of interest.

The time between bunch crossings at the LHC will be only 15 ns, compared to 96 ns at HERA, 3.8 μs at the S$\bar{p}$pS and 22 μs (11 μs for future running) at LEP. Although triggering at HERA is addressing many of the problems associated with having a short time between bunch crossings, the problems at LHC are much more severe. Whereas at HERA almost all bunch crossings are empty, for high luminosity at LHC each bunch crossing will contain a large number of interactions. The rate at which data will have to be processed by the level-1 trigger will therefore be extremely high. A realistic number for the calorimeter trigger is ≈ 5000 Gbits/s as described later. Clearly, the trigger systems for LHC experiments will have to be much faster and more powerful than anything currently in existence.

Several R&D projects have been funded by the DRDC in the area of front-end electronics, triggering and data acquisition [1]. However, none of these fully addresses the question of level-1 triggering, which in our view is one of the most challenging problems of building experiments for the LHC.

We propose to perform a broad-based study of level-1 triggering for the LHC which will complement existing activities. Some members of our collaboration are also involved in related work within other R&D projects. We see this as being mutually beneficial since it will help to avoid duplication and ensure that there is an efficient exchange of information between different R&D groups.

Overview of multi-level trigger architecture

It is generally accepted [2] that the trigger systems for LHC experiments will consist of several trigger levels, as indicated in Fig. 1. The level-1 trigger will be synchronous and have a fixed latency (decision time) of at most a few microseconds. The input rate to the level-1 trigger system is the 67 MHz bunch-crossing frequency of the LHC machine, and at high luminosity each event will contain many interactions. It is envisaged that the level-1 trigger rate will be in the range $10^4 - 10^5$ Hz, chosen to match the capabilities of planned level-2 trigger systems [3]. During the level-1 latency, the data from *all* detector channels will be stored in pipeline memories. Given the very large number of channels expected for precision tracking detectors ($10^6 - 10^7$ channels

in some cases), the cost of the pipeline memory electronics is likely to be very substantial. Every effort should therefore be made to minimize the level-1 latency and hence the required depth of the pipeline memories. Whilst the level-1 trigger will probably be based on data from only a subset of detectors, its decision must be distributed to the front-end electronics of all detector elements.

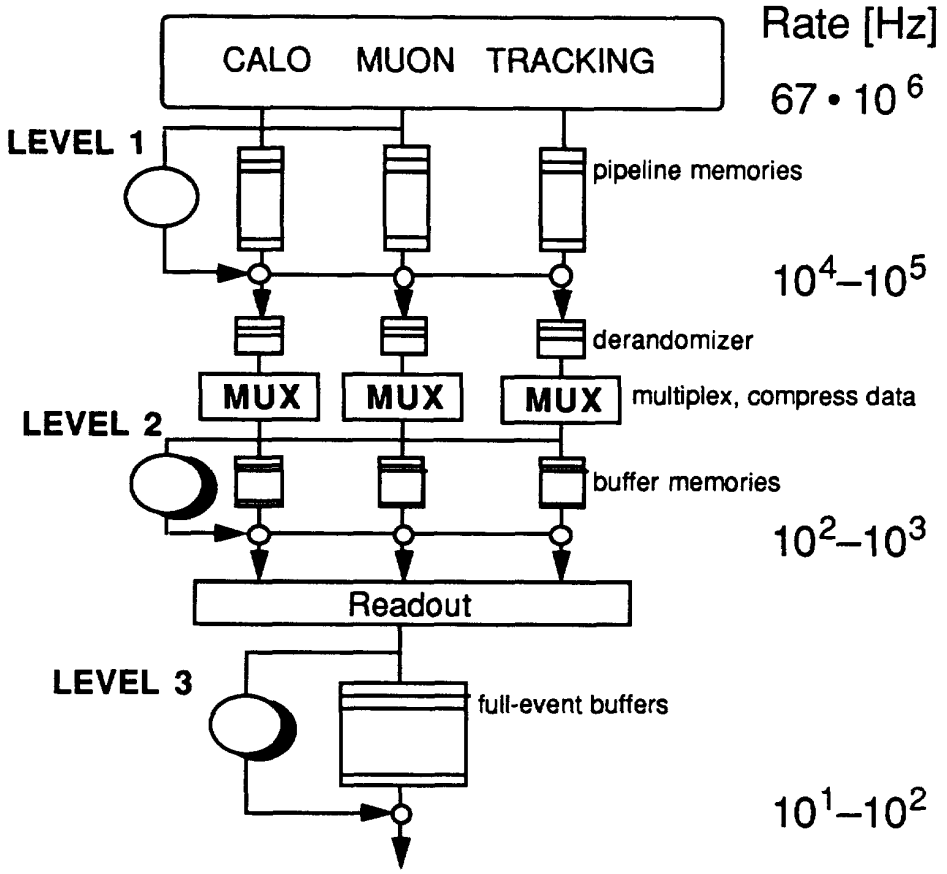

Figure 1: Multi-level trigger architecture

In the system which is illustrated in Fig. 1, data for events selected by the level-1 trigger will be transferred to buffer memories where they will be stored during level-2 trigger processing. For some detectors, data reduction and multiplexing will be done before the data are moved. Derandomizing memories, placed after the level-1 pipelines, will accept data in the 15 ns between bunch crossings, introducing no deadtime in normal operation. These data can then be compressed, multiplexed and transmitted to the buffer memories more slowly. The rate at which the derandomizing memories are emptied must exceed the level-1 trigger rate, but can be much less than the bunch-crossing rate.

The level-2 trigger system will use data from the buffer memories to refine the trigger decision made at level-1. Several techniques have been suggested for making data

stored in the level-2 buffers available to the level-2 processor system [4, 5]. For triggers which depend on local regions of the detector, such as high-p_T electron or muon triggers, the level-2 trigger can make use of regions of interest (ROI) defined by the level-1 trigger system [6]. This may require direct connections between the level-1 and level-2 trigger systems (not shown in Fig. 1) and is an area that we propose to study.

The level-2 trigger will reduce the event rate to the range $10^2 - 10^3$ Hz. These events will then be read out and subjected to further selection in the level-3 trigger system, which might consist of a farm of powerful processors.

Overview of the level–1 system

We envisage that the level-1 trigger system will contain a number of processors associated with different detectors or physics signatures, as shown in Fig. 2. The minimal LHC trigger system for a general-purpose experiment should have trigger processors associated with calorimetry and the muon detection system. These should provide triggers on high-p_T muons, electrons/photons (i.e. electromagnetic clusters), jets (possibly including taus), and large missing transverse energy. Some detectors, such as high-precision inner trackers, may not contribute to the level-1 trigger. Although the possibility of a level-1 inner tracking trigger is not excluded, it is less essential than the calorimeter and muon triggers.

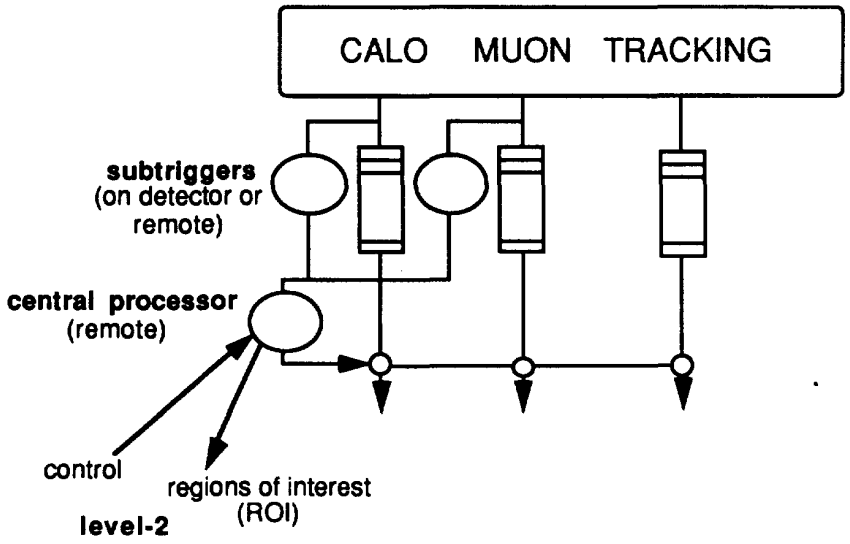

Figure 2: Level-1 architecture

A central level-1 processor is also required, for correlating the subtrigger results and forming combinational triggers. This central processor must interact with the level-2 system and with the front-end electronics.

The function of the level-1 central processor is illustrated in Fig. 3. It receives results from the subtriggers for a number of physics signatures and combines them to produce the overall "yes/no" level-1 trigger decision. This decision is then distributed to the front-end electronics. Interaction with the level-2 system will include the possibility to inhibit further level-1 triggers if the level-2 buffers become full, and the transmission of regions of interest and other useful data to the level-2 system.

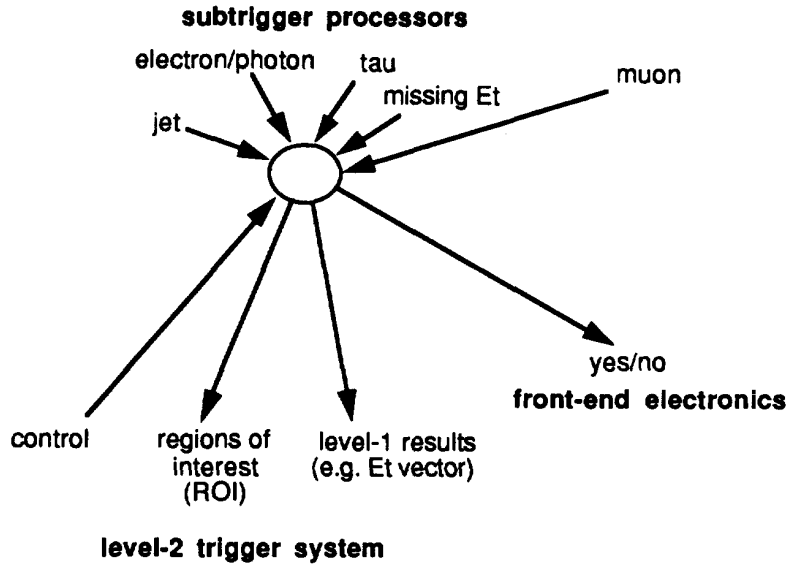

Figure 3: Central level-1 processor

Issues to be studied

We plan to study the problems of implementing a level-1 trigger system containing a number of different subtrigger processors and a central processor. We will also investigate the interaction between the level-1 system and the rest of the experiment.

A critical problem is that of feeding data into the level-1 processing system at the required rate. This is very severe in the case of the calorimeter trigger where we envisage several thousand "trigger channels" delivering digital (~ 8 bit) energy values every 15 ns; in the trigger scheme described below, the data rate is ≈ 5000 Gbits/s. The need for two-dimensional cluster-finding algorithms for high-p_T electron/photon triggers, which requires a fan-out of data between processing elements, makes this problem particularly acute.

The bunch-crossing period of 15 ns corresponds to an "event" rate into the level-1 system of 67 MHz, where each event may contain up to ≈ 20 inelastic proton–proton collisions. Considering the large physical size of the detectors and the consequent signal propagation delays, a level-1 trigger latency far in excess of the bunch-crossing interval is inevitable. Furthermore, complex electronic systems will be needed in order

to achieve the required background rejection power, and the processing time will certainly be long compared to 15 ns. In consequence, the trigger processing will have to be pipelined in 15 ns steps, with a large number of events being processed concurrently. The short bunch-crossing period of 15 ns has another important implication. In the time it takes muons to reach the farthest parts of the apparatus in typical LHC detectors, three bunch crossings will have taken place; time-of-flight corrections will be essential!

We plan to investigate the problem of keeping the trigger system synchronized with the necessary accuracy. Pipelined processing will be performed throughout the level-1 system with a clock speed of 67 MHz or more. Synchronization will be very challenging given that the processing system may be distributed over a large area. For example, the muon trigger system will probably be mounted on the detectors, and, therefore, spread over an area of several thousand square metres. We see a need to include synchronization mechanisms in the system design.

Synchronization between the inputs to each subtrigger processor from different regions of the detector must be established and maintained. Corrections will have to be made for different time-of-flight and signal propagation delays. Similarly, inputs to the level-1 central processor from different subtriggers will have to be brought into synchronization. Here, the results from subtriggers having short latencies will have to be delayed until the results of slower subtriggers are available, and the clock phases will have to be equalized. Finally, the "yes/no" trigger decision will have to be conveyed to the front-end electronics in synchronization with the readout pipeline memories.

The problem of giving a unique identification of the bunch crossing which contained the interaction of interest will be a difficult one to solve. Some of the detectors, especially calorimeters, may have a response time which is long compared to the 15 ns bunch-crossing period. Special processing will have to be included in the trigger system to do this. An alternative would be for the level-1 trigger to identify a time frame containing several crossings, one of which is the one of interest. However, this would put a large burden on the level-2 system, which would then have to handle much more data and perform the bunch-crossing identification analysis. We aim to achieve unique bunch-crossing identification in the level-1 trigger.

While the level-1 central processor will be fairly compact, possibly located somewhere near the detector, it will draw information from a number of trigger processors (subtriggers) associated with different subdetectors. Some of the processing, for example searching for tracks in muon detectors or clusters in electromagnetic calorimeters, could be done locally. The best location for the trigger electronics is a

complicated question which we plan to study. As discussed below, it affects the level-1 latency, an important parameter that may strongly influence the cost of detector readout systems.

We have made a very rough estimate of the level-1 latency by considering cable delays and estimated trigger processing times. The total level-1 latency is the interval from the bunch-crossing time to the time when the trigger decision gets back to the front-end electronics. The calculation for each subtrigger must include the time-of-flight to the detector, the response time of the detector, delays for signal shaping, digitization, digital processing, and propagation delays along electrical or optical cables. The calculation of the total latency is determined by the decision time of the slowest subtrigger, the processing time (in the level-1 central processor) to form compound triggers, cable delays back to the front-end electronics, which in many cases will be mounted on the detectors, and the delay through the electronics which fans the trigger decision out to the front-end electronics (up to $10^6 - 10^7$ channels). Assuming that all the electronics is mounted very close to the detector and that all the subdetectors contributing to the level-1 trigger produce prompt signals, the total level-1 latency is likely to be in the range $1 - 2$ μs. We plan to study this important issue in more detail. Overestimation of the latency would make the detector readout pipelines much more expensive than necessary, but underestimation could be disastrous.

Proposed R&D

Physics Simulation

Extensive physics simulation studies to estimate trigger rates and evaluate level-1 trigger algorithms are in progress within our collaboration [7, 8]. This work is coordinated with activities for level-2 triggers so that the combined rejection power of the level-1/level-2 system can be evaluated consistently. So far we have considered triggers for high-p_T muons, electrons, photons and jets, and for large missing E_T.

We expect to continue and extend the present simulation studies using more refined detector models. We will continue our evaluation of different trigger algorithms and investigate the effect of changing parameters such as the detector granularity used in the trigger. In addition, we will consider the possibility of a level-1 trigger for high-p_T tau production.

We will also evaluate the performance of compound triggers such as electron–muon coincidences by combining work being done by different groups within our collaboration. This work will determine the requirements of the central trigger.

These studies, which are likely to require significant computer resources, will mainly be performed outside CERN.

Our programme of physics simulation studies for the first year of the project is as follows:

- Evaluation of physics requirements and background rates for a full range of level-1 trigger options.
- Evaluation of different algorithms for muon, electron, photon, jet and missing transverse energy triggers.
- Evaluation of compound triggers based on several signatures in coincidence (e.g. electron–muon coincidence).
- Optimization of trigger algorithms with respect to parameters such as the detector granularity used in the trigger.
- Evaluation of the occupancy of detectors used in the trigger.

The above physics simulation activity will allow work in other project areas to continue on a sound basis.

Systems Aspects

Systems aspects are concerned with the interaction between different parts of the level-1 system, and between the level-1 system and its environment. The latter includes communications with the front-end electronics, the level-2 trigger and the DAQ system. There are a number of general issues that affect the level-1 trigger system:

- Establishing and maintaining synchronization throughout the level-1 system.
- Data transmission and connectivity.
- Location of the trigger electronics.
- Latency of the level-1 trigger.
- Interaction between level-1 and the rest of the trigger and data acquisition system.
- Making a unique determination of the bunch crossing responsible for the trigger.

We see synchronization as being a very challenging issue which must be considered at many points in the system. The inputs to each subtrigger from different regions of the detector must be brought into synchronization, as must the inputs to the central trigger processor from the different subtriggers. Each subtrigger will consist of a synchronous pipelined processor, parts of which may be distributed over a large area of the detector. The trigger decision will have to be distributed to the front-end electronics in synchronization with the readout pipelines. Facilities must be included to establish and monitor the synchronization throughout the system, and to recover it when it is lost.

Data transfer within the level-1 trigger system will also be very challenging given the massive data rates involved, especially for the calorimeter trigger as described below. We plan to investigate the possibilities offered by advanced technology. The use of very fast data links will be investigated for connections between electronic modules, while advanced interconnection technology such as Multi-Chip Modules (MCMs) and bump bonding of Application Specific Integrated Circuits (ASICs) will be studied for short-range communications. Data compression techniques may also play a role in reducing the bandwidth required within the level-1 system.

Simulation models of the level-1 system will be used to evaluate some of these issues. We envisage that different subtriggers will be simulated in detail, for example using VHDL and VERILOG programs. The interaction of the different level-1 subtriggers, and the interaction between the level-1 and level-2 systems, will also be investigated. Simulation at the system level will be performed, for example using the MODSIM program.

Many of these systems aspects will be studied in the context of the central trigger processor and subtrigger processors, as described in the following sections. The calorimeter trigger requires a massive input data rate which will have to be provided on high-speed connections. The muon trigger, with on-detector electronics spread over a huge area, will be very challenging from the point of view of synchronization.

We expect to identify components which are common to different subtriggers and will aim for a consistent treatment wherever possible. Possible examples are data links between the subtriggers and the level-2 system, and interfaces to the data acquisition and slow control systems. Similarly, calibration and monitoring should be included in the system design.

Our programme of work on systems aspects of the level-1 system is as follows:

- In the first year of the project, an outline specification and design will be developed for the timing and synchronization system to meet the requirements of an LHC level-1 trigger.
- Critical areas requiring further study, detailed design and prototyping will be identified for study in the second year.

These and other systems aspects will also be addressed in the context of the central level-1 processor and of the subtrigger systems, as discussed in the following sections.

Central Trigger Processor

The central trigger has to correlate information from a number of different subtriggers associated with different subdetectors or with specific physics signatures. Existing

experiments have adopted a variety of algorithms, some of which use topological information from the subtriggers in addition to multiplicity information. We plan to investigate the requirements for LHC experiments by physics simulation. The implementation of the central trigger processor will also be investigated.

We see a need for programmable delays between each subtrigger and the input to the central processor. This facility, already included in experiments at HERA where the bunch-crossing period is 96 ns, synchronizes the data from the different subtriggers. The central trigger processor is also likely to play a role in the distribution of clock and control signals (such as level-1 "yes/no") to the subtriggers and the front-end readout electronics.

In RD–12, the DRDC has funded the development of the basic hardware and software components of a multichannel optical-fibre timing distribution system for LHC detector front-end electronics, and an investigation of the feasibility of simultaneously exploiting the timing system for the transmission of level-1 trigger acceptance and addressable control information. This work has yielded very promising results during the first 12 months, and is now entering the second phase in which a complete medium-power prototype system will be developed.

Diode-pumped solid-state (DPSS) lasers are now becoming available which can generate high CW optical powers close to the preferred zero-dispersion wavelength of 1300 nm. As an extension of the preliminary work with laser diode sources, we propose to develop an externally-modulated high-power DPSS Nd:YAG laser source capable of transmitting the timing and control signals to a realistic number of front-end electronics modules. We shall also study the integration of the prototype system with the central trigger processor and front-end electronics.

Another important aspect of the central level-1 trigger processor is its interaction with the level-2 system. Clearly, the level-2 system must be informed when each new level-1 trigger is generated. However, the interaction is complicated by the need to inhibit level-1 triggers if the capacity of the derandomizing memories or level-2 buffer memories is exceeded (see Fig. 1). The mechanism for inhibiting further level-1 triggers must allow for the fact that the electronics will not be able to respond within the 15 ns between bunch-crossings. The communication of region-of-interest information is discussed below in the section on interaction between level-1 and level-2 systems.

Our proposed programme of work on the central trigger processor is as follows:

- Produce a specification and outline design within the first year of the project. This will require close collaboration and interaction with the physics simulation

groups, and will carefully balance electronics cost and complexity against real physics requirements.

- In the second year of the project, more detailed design and prototyping work will be performed for the critical areas identified in the first year.

Calorimeter Trigger

The calorimeter trigger system for a general-purpose LHC experiment should include electromagnetic cluster-finding (electron/photon), jet finding and missing-E_T triggers. The first of these is the most critical as it will probably dominate the level-1 trigger rate. We propose to make a detailed investigation of how a trigger processor for an LHC experiment could be implemented, building on work performed by members of our collaboration over the last two years.

The existing work [9, 10] is for a completely digital level-1 trigger processor. All the signatures described above are implemented in our conceptual designs and the electron/photon trigger algorithm includes the option to require isolation of the electromagnetic cluster.

We are already constructing a small prototype processor for the electron/photon trigger, based on a synchronous bit-parallel processor design, and it is planned to carry out beam tests with the prototype calorimeter of RD–3 towards the end of 1992. A more detailed description of this ongoing project is given in Appendix 1. The cluster algorithm, which views a 4×4 area of trigger cells in the electromagnetic calorimeter, is illustrated in Fig. 4. It has been implemented on an ASIC which is currently under test at the full 67 MHz clock speed. This algorithm offers a combination of excellent efficiency (no loss when showers are shared between adjacent cells) and up to an order-of-magnitude additional background rejection from the optional isolation requirement. While we remain open to alternative cluster algorithms, the algorithm details do not strongly influence the trigger system design.

The above processor design needs to be improved several ways. The number of interconnections in the prototype system described above is a big problem. The number of cluster positions processed by each ASIC is limited to one because of the large number of pins needed to input the 16×8-bit words of data required for the cluster algorithm, and present limits on the number of pins available on an ASIC package (the data are transmitted to the ASICs in bit-parallel form at 67 MHz). The interconnection problem becomes even more severe when inter-module and inter-crate data sharing is considered. While the extrapolation of such a system to a full LHC calorimeter trigger is possible, it would require a large number of crates of electronics with many connections between them. This interconnection problem must be overcome in future

designs by using more advanced technology and/or by using alternative architectures as discussed below.

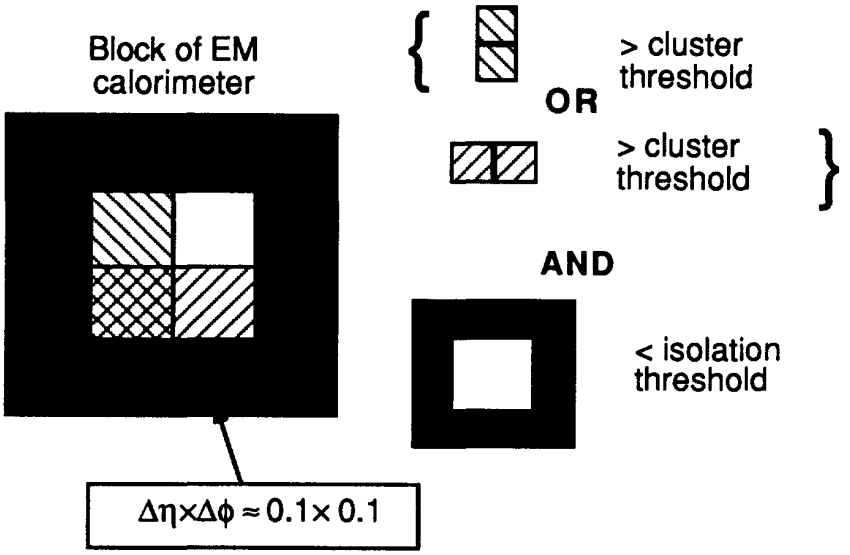

Figure 4: Electromagnetic cluster algorithm

Three possible alternative architectures are being studied:

- An improved single processor with synchronous bit-parallel data transmission and processing. This would have to use more advanced interconnection technology to overcome the problems described above and in Appendix 1.
- A single processor with asynchronous transmission of compressed zero-suppressed data, followed by synchronous bit-parallel data processing. More details of this option are given in Appendix 2.
- A farm of processors with bit-serial data transmission and processing, allowing a more compact design. A more detailed description of this option is included in Appendix 3.

The system aspects of all three possibilities are being actively pursued, each having some advantages. We will use detailed system design studies to select an architecture within the first year of the project.

In addition to overcoming the interconnection problem and making a more compact processor design, several significant improvements are required. We consider the addition of logic to uniquely identify the bunch-crossing responsible for the trigger to be essential. This might be achieved by digital (or analogue) signal processing at the level of individual trigger channels. We also plan to improve the effectiveness of the isolation requirement by including hadron calorimeter information in the cluster-finding algorithm.

We intend to construct a second small prototype processor. This will be used to assess the feasibility of key parts of the system. We plan to make beam tests of this second prototype trigger processor together with prototype calorimeters for LHC. This will require the design, development and construction of additional electronics. Analogue summation will be used to combine the highly granular calorimeter cells into "trigger cells" corresponding to a granularity of $\Delta\eta\times\Delta\phi \approx 0.1\times0.1$. The summed signals will then be digitized using 8-bit flash ADCs and the values will be transformed to E_T using look-up tables. Logic to uniquely identify the bunch crossing responsible for high-E_T data may also be included. As discussed above, various schemes are under consideration for transmitting the resulting data to the trigger processor. Tests with the FERMI modules of RD–16 are also envisaged when these become available.

As already stated, the prototype trigger processor which is currently under construction will be tested with the liquid argon calorimeter of RD–3 towards the end of 1992. These tests will use an analogue summing system constructed from commercial linear fan-in modules and a flash ADC system made from a modified drift-chamber readout board from the Zeus experiment. It is unlikely to be possible to adapt this digitization system to the alternative architectures described above.

The MPI–Munich group is involved in the construction of a prototype liquid Argon calorimeter [11]. A first prototype, which will use relatively slow electronics designed for the H1 experiment at HERA (bunch-crossing period 96 ns), will be available in 1993. In a second phase of the project (1994–5), a prototype calorimeter will be constructed with the properties required for an LHC experiment. We plan to design, develop and build analogue summing and digitization electronics to allow these calorimeters to be connected to the prototype trigger processor.

The connection between the trigger ADCs and the processor is a key issue and we consider it to be very important to test this part of the system in as realistic an environment as possible.

In summary, the following programme of work is then foreseen:

- System simulation of the full calorimeter trigger.
- Investigation of the application of advanced digital communication methods and technologies, including data compression techniques.
- The design and fabrication of ASICs for a small prototype processor which should include the additional features described above.
- The development of a prototype data transfer system to connect the digitization system to the calorimeter trigger processor. It will also be necessary to construct a

system to generate test data for transmission to the trigger processor, possibly making use of facilities provided by RD–12.

- Beam tests with prototype calorimeters. This will require the construction of a digitization system and/or tests with the FERMI modules of RD–16.

Muon trigger

Muon trigger systems based on Resistive Plate Chambers (RPCs) are already under study in the RD–5 project and in use in experiment WA–92. The RPCs consist of strips which give digital hit information. They are attractive as muon trigger detectors for LHC because of their fast (much less than 15 ns) response time and their low cost. The trigger system described below could be used with RPCs or with other detectors having good time resolution.

Level-1 muon triggers for LHC experiments can be made using two planes of strip detectors separated by a distance of, typically, 1 m. High-momentum muons will give hits in both layers approximately along a straight line which passes through the interaction region. Low momentum muons will be strongly deflected by magnetic bending, for example in magnetized iron, and by multiple Coulomb scattering. Therefore, for each possible hit position in the inner detector layer, only a limited number of hit positions in the outer layer correspond to high-momentum muon tracks. Fast coincidence logic can be used to identify valid coincidences, where the width of the "road" is related to the momentum cut. A typical granularity of the trigger, segmented in η and ϕ, would require ~ 10^5 strips.

The design and construction of such a trigger for LHC is complicated by the need to identify muons within the 15 ns bunch crossing interval. This puts stringent demands on the design of the system in terms of synchronization, which can only be achieved by means of careful signal routing between the detector layers and the trigger electronics.

The level-1 muon trigger system will be distributed over a very large detector area of the order of thousands of square metres. For this reason it must be logically divided into sectors (octants in the present design) and each sector must be further divided into subsectors. The trigger acts locally, identifying tracks in subsectors.

Each subsector is made of two RPC superlayers separated by about a metre. Each superlayer consists of two offset layers of RPCs 20 – 30 m long and a few metres wide. The coincidences between the inner and outer superlayers are made by coincidence matrices, each identifying tracks in a small area of the apparatus. The trigger information coming from the coincidence matrices is sent to the subsector trigger logic circuit. This local information must subsequently be combined with level-1 trigger information coming from different parts of the apparatus. This will be done in a global

muon trigger processor and the central level-1 processor. Region of interest information is also made available for use in the level-2 trigger.

We plan to investigate the design of such a level-1 muon trigger system through the following R&D programme:

- Design and implementation of a trigger subsystem and its associated electronics using RPC detector layers and existing coincidence matrix circuits. This will make use of commercial PLA chips or an ASIC developed for the H1 forward muon trigger [12]. The subsystem will be tested at the RD–5 test-beam facility.
- Design and implementation of a system for signal distribution with synchronization at the few-nanosecond level distributed over the length of the level-1 muon trigger subsector.
- Design and implementation of a prototype local muon trigger logic circuit having the performance required for LHC. This will be related to the design of the global muon trigger processor and the central level-1 processor.

Interaction between level–1 and level–2 systems

The level-1 and level-2 trigger systems at LHC experiments will have to be more tightly coupled than in existing experiments. Fast control signals will be required, for example to inhibit further level-1 triggers if the level-2 system becomes full. In addition, we see many benefits in using the level-1 trigger system to identify regions of interest (ROIs) for level-2 processing. This allows the level-2 trigger to act locally as described below.

"Feature extraction" is the first stage of level-2 processing and one which takes considerable time. Data from a limited volume of a detector are analysed and attempts are made to specify the information in a few words. For example, an electron candidate might be characterized by its E_T and $\eta - \phi$ position, plus a few binary words indicating its quality as an electron candidate, its isolation characteristics, etc.

Feature extraction operates on a limited volume of data and is therefore a local operation; the data input to the processor need only be connected locally. Because of this local nature of feature extraction, considerable benefits in data transfer and processing rates accrue from using the level-1 trigger system to provide information on regions of interest. The level-1 trigger system can flag ROIs, if necessary using additional thresholds which are lower than the ones used for the level-1 trigger decision. Only a small fraction (~ few percent) of the data from the detector need then be transferred into the level-2 trigger system. The cost of using ROI information in the level-2 trigger is a more complex information exchange between trigger levels as indicated in Fig. 5.

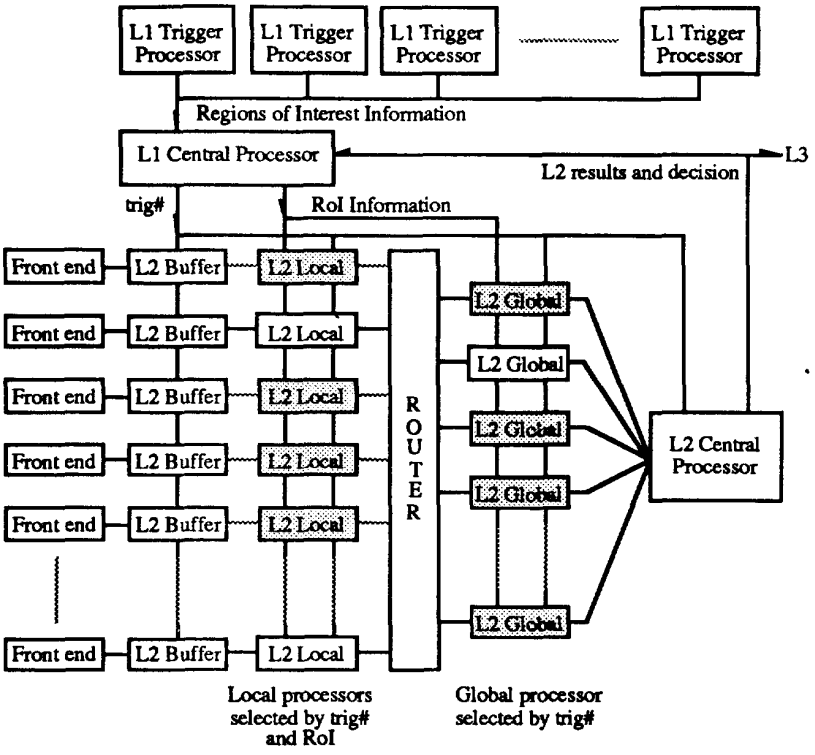

Figure 5: Interaction between level-1 and level-2 systems.

The use of ROI information greatly increases the number of options for the implementation of level-2 triggers. If no ROI information is used, the level-2 trigger must be based on dedicated custom hardware (all data must be processed for every level-1 trigger at a trigger rate of up to ~ 100 kHz). However, if ROI information is used, the level-2 trigger may use either special-purpose or general-purpose processors (or a combination of both) because local data transfers take place at a much lower rate (~ 1 kHz) and longer access and processing times (i.e. longer latency) can be accommodated.

The interaction between the level-1 and level-2 systems will require physics simulation to evaluate which data are required. Detailed design studies will then have to be made of the system linking the two trigger levels. We consider the construction of a prototype level-1/level-2 interface system as essential. Evaluation of the prototype will be performed on a test-bench facility. Initially, this will be VME-based and consist of the following parts:

- Modules which will emulate the level-1 central trigger processor by generating simulated ROI information.
- Active elements (e.g. DSP 56000 boards) which will collect relevant ROI data to be passed on to level-2 processors.

- Modules which will receive the ROI data and emulate level-2 processors.

This test facility will allow the development of software and the evaluation of prototype architectures for the level-1/level-2 interface in terms of speed, data integrity and other parameters which are important for level-2 processing.

In the longer term, it will be necessary to extend the test-bench system to a more realistic configuration for an LHC experiment. A new generation of hardware (e.g. newer DSPs) will probably have to be implemented and the software will have to be redesigned and tested. The transition to a new and higher-performance bus system such as FutureBus+ will probably be necessary.

In summary, we wish to investigate methods for using the detailed information generated at level-1 to guide level-2 processing. This will include:

- The study of which information needs to be provided following a level-1 trigger.
- The design of circuits and architectures to supply data from level-1 to level-2 in the form required. This may be affected by the design of the front-end electronics for the different subdetectors.
- Testing of the design by building a prototype system containing all the elements involved. The test-bench systems described above will be essential for this work.

Conclusions

The development of trigger systems for LHC will be extremely challenging, particularly at level-1 where the full 67 MHz rate must be processed. The proposed project aims to identify and solve some of the problems of constructing the first-level trigger systems for LHC experiments.

We will use physics simulation to evaluate the performance requirements of the level-1 trigger system. Calculations will be made of trigger rates and efficiencies for a wide range of physics signatures, and the effectiveness of trigger algorithms will be evaluated. Physics simulation will allow us to estimate the required detector granularity and occupancy, which have important implications for the trigger design.

A number of systems issues will be addressed, three of the most important being the problem of establishing and maintaining synchronization throughout the level-1 system, the problem of getting the large required amount of data into the level-1 processor system at the 67 MHz rate, and the problem of making a unique bunch-crossing identification at level-1. We also plan to study the question of level-1 latency which has important implications for the design of the readout electronics for all detectors, and which will be effected by the choice of location of the level-1 electronic systems.

We see a need for a central level-1 trigger processor which will interact with level-1 subtrigger systems, the front-end electronics and the level-2 system. We will investigate the requirements for this part of the system and carry out detailed design studies. We will then build and evaluate prototypes for some critical components.

We will continue and extend a programme of R&D on level-1 calorimeter triggers which was started outside the context of the DRDC. We expect to make extensive use of ASICs to obtain the required logic density. The problem of feeding the calorimeter data into the trigger processor system is a major challenge which is under study. We plan to make detailed design studies for the whole level-1 calorimeter trigger system. The feasibility of this system will then be studied by building a small prototype trigger processor containing as many of the critical components as possible. This will be evaluated in beam tests with prototype calorimeters.

We also plan to investigate level-1 muon trigger systems. A muon trigger subsystem will be designed and a prototype will be constructed using RPC chambers as the muon detector. This will require synchronization at the level of a few nanoseconds over the large area of a muon detector subsector. The interaction between the local muon trigger logic and the rest of the level-1 system will be studied. The prototype trigger system will be evaluated at the RD–5 test-beam facility.

Finally, we plan to study the requirements for transmitting region-of-interest information from level-1 to the level-2 system. We will design this interface and build a small prototype system to evaluate its performance.

Computing needs

Physics simulation studies will require significant computer time which will mainly be provided outside of CERN. We estimate a total annual requirement of 5000 hours (CERN IBM 168 accounting units), 1500 hours to be provided at CERN.

Sharing of responsibilities

The areas of responsibility of the participating groups are summarized in Table 1. Initially, each group will provide the funding required for its parts of the project. However, we are exploring the possibility of collaborative R&D with industry and non-HEP research institutions and may need to approach the DRDC for additional funds at a later stage of the project.

The proposed contributions of (financial and engineering manpower) resources are given in Table 2. All the figures are quoted for the two-year period of the proposed project. It is foreseen that an additional engineer will be recruited by CERN and join the project in 1993.

	Physics Simulat-ion	Systems Aspects	Timing & Control	Central Trigger	Calori-meter Trigger	Muon Trigger	Level-1 / Level-2
Birmingham	●	●			●		
CERN		●	●				
Heidelberg	●			●	●		●
Munich-MPI 1)	●				●		
QMW, London	●				●		
RAL		●		●	●		●
RHBNC, London	●						●
Rome	●	●				●	
Stockholm	●	●			●		●

1) Subject to approval by the institute.

Table 1: Areas of responsibility of groups.

	Financial Contribution (kSF)	Engineering Manpower (man years)
Birmingham	100	2
CERN	100 1)	2
Heidelberg	180 1)	2
Munich-MPI	50 1)	2
QMW, London	20	0
RAL		2 2)
RHBNC, London	35	1
Rome	120 1)	2
Stockholm	40 1) 3)	4

1) Funds requested but not yet approved. 2) Assuming current level of support is maintained.
3) Request for significantly larger funds under discussion.

Table 2: Financial and manpower contributions of groups for a two-year period.

Schedule

The first year of the project will mainly be used to perform design studies for the various parts of the level-1 system. However, in the case of the calorimeter and muon trigger systems, where work has already been in progress for some time, we hope make some initial tests of prototype trigger systems. Specific objectives for the first year are the following:

- Outline specification and design for the timing and synchronization system to meet the requirements of an LHC level-1 trigger.
- Outline specification and design of a central level-1 trigger processor.
- System design for a digital calorimeter trigger processor, including the data transfer system to connect the digitization system to the processor.
- Beam tests of the digital calorimeter trigger processor which is currently under construction. These tests will be performed at CERN using signals from prototype calorimeters.
- Design and implementation of a muon trigger subsystem and its associated electronics using RPC detectors. This will include a system for signal distribution and synchronization at the few-nanosecond level.
- Evaluation of the requirements of the level-2 trigger for region-of-interest information. Studies will be made of the architecture and system design to supply data from the level-1 system to level-2.

In the second year of the project, we will make more detailed design studies and construct prototypes for critical parts of the level-1 system.

References

[1] "Study of muon triggers and momentum reconstruction in a strong magnetic field for a muon detector at LHC", RD–5.
"Embedded architectures for second-level triggering in LHC experiments (EAST)", RD–11.
"Readout system test benches", RD–12.
"A scalable data-taking system at a test beam for LHC", RD–13.
"Opto-electronic signal transfer for LHC experiments", RD–23.

[2] S. Cittolin, N. Ellis, L. Mapelli, "Signal processing, triggering and data acquisition", *Proc. ECFA Large Hadron Collider workshop, Aachen, Germany*, 1990, p.504. CERN report 90–10 (vol I), eds. G. Jarlskog and D. Rein.

[3] "Embedded architectures for second-level triggering in LHC experiments (EAST)", RD–11.

[4] "A digital front-end and readout microsystem for calorimetry at LHC", RD–16.

[5] "Integrated high-rate transition radiation detector and tracking chamber for the LHC", RD–6.

[6] D. Crosetto, N. Ellis, G. Mornacchi, J. Strong, "A local/global architecture for level-two calorimeter triggers". *Proc. ECFA Large Hadron Collider workshop, Aachen, Germany*, 1990, p.145. CERN report 90–10 (vol III), eds. G. Jarlskog and D. Rein.

[7] A. Nisati, "Muon rates at the LHC", *Proc. ECFA Large Hadron Collider workshop, Aachen, Germany*, 1990, p.442. CERN report 90–10 (vol III), eds. G. Jarlskog and D. Rein.

[8] Transparencies presented by A. Watson at Eagle trigger/DAQ meetings.
S. Hellman et al, "Trigger rates at the LHC", *Proc. ECFA Large Hadron Collider workshop, Aachen, Germany*, 1990, p.72. CERN report 90–10 (vol III), eds. G. Jarlskog and D. Rein.

[9] N. Ellis et al, "Studies in fast calorimeter trigger design", status report to UK Particle Physics Grants Sub-Committee, October 1991.

[10] C. Bohm et al, "A tentative design of a bit-serial first-level calorimeter trigger for an LHC detector", technical note in preparation.

[11] "A thin gap liquid Argon calorimeter for LHC experiments", to be submitted to the DRDC.

[12] Track-finder IC number RAL–115, RAL micro-electronics group, SERC, UK.

[13] J. Yuan, C. Svensson, "Pushing the limits of standard CMOS", IEEE Spectrum, February 1991.

Appendix 1:

Ongoing work with a bit-parallel calorimeter processor design

The project *Studies of first-level calorimeter trigger systems* was funded by the UK Science and Engineering Research Council for two years from October 1990 and has now been extended for another two years; the participating groups are Birmingham, QMW and RAL. The aim was to investigate the feasibility of a digital level-1 calorimeter trigger for the LHC. Significant design progress has been made, and ASICs have been fabricated to perform an electron/photon trigger algorithm at the full 67 MHz rate of LHC as described below. We plan to continue and extend the work as part of the larger project that we are proposing to the DRDC.

Because of the limited resources available to us, most of our work has been restricted to the study of the electron/photon trigger, which we consider to be the most demanding part of the system. This has to use relatively fine granularity information from the electromagnetic calorimeters and must therefore handle a large number of input channels. There is a large physics background to the electron/photon trigger from high-p_T jets which must be overcome by using a relatively complicated algorithm. By demanding isolation of electromagnetic clusters, this background can be reduced by up to an order of magnitude while maintaining good efficiency for electrons or photons from W and Z decays.

We are building a small prototype electron/photon trigger processor, using ASICs to implement the cluster algorithm. This will initially be tested by using RAMs to provide the input data at the full 67 MHz rate of LHC. A small system of analogue summing circuits and fast ADCs will be provided to test the trigger on the liquid argon calorimeter prototype of RD–13 during a beam test at CERN towards the end of 1992.

Electron/photon trigger algorithm

The trigger algorithm which we have adopted was chosen after simulation studies of high transverse-energy jets, which are the dominant background to the electron/photon trigger. The algorithm is based on a 4×4 overlapping sliding window in the electromagnetic (em) calorimeter as shown in Fig. A1.1. Each square in the figure corresponds to the smallest unit used in the trigger, the "trigger cell". Lateral and longitudinal summation over em calorimeter cells is performed before the trigger processor, possibly using on-detector electronics, to provide signals for the reduced granularity used in the trigger. A trigger granularity in pseudorapidity–azimuth space of $\Delta\eta\times\Delta\phi \approx 0.1\times0.1$ is envisaged, giving ≈ 4000 channels for a calorimeter with rapidity coverage $|\eta| < 3$.

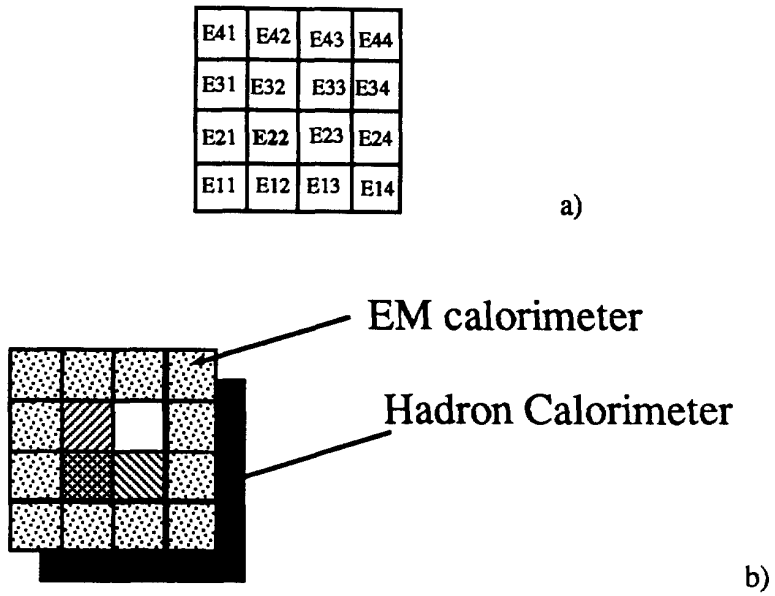

Figure A1.1: (a) Trigger window and (b) Cluster algorithm.

The algorithm requires a high E_T cluster contained in a pair of em cells and, optionally, isolation of the cluster. The cells containing the cluster may be a horizontal (E22+E23) or a vertical pair (E22+E32). This gives good efficiency even if the electron shower is shared between two trigger cells. There is some inefficiency if the shower is shared between three or more trigger cells. However, given the low probability for this (the trigger cells are relatively large), we believe that the inefficiency is acceptable.

The optional isolation condition requires that the transverse energy sum over the 12 em cells surrounding the cluster be small (see Fig. A1.1). All four cells in the centre of the window are excluded from the isolation area so that full efficiency is retained, from the point of view of isolation, even if the electron shower is shared between four trigger cells. We are aware that significantly improved discrimination can be obtained, between the electron/photon signature and the high-p_T jet background, by including the hadronic calorimeter in the isolation sum. We would like to implement this in the next phase of the project.

The electron/photon trigger provides several (currently two) programmable cluster thresholds, with corresponding independently programmable isolation thresholds. These are needed in order to allow inclusive triggers with and without isolation, and for multi-lepton triggers, etc.

We have implemented the electron/photon trigger algorithm described above as an ASIC, the internal operation of which has already been tested at the full 67 MHz input rate of LHC.

An important conclusion of our project so far is that it is relatively straightforward to implement cluster algorithms of this kind using ASICs. The main difficulty is in supplying the input data, which requires a large number of connections as discussed below. Therefore, changes in the details of the cluster algorithm do not critically affect the system design.

Implementation of a trigger for the LHC

A block diagram of the prototype trigger system is shown in Fig. A1.2. Signals from the calorimeter are summed using analogue electronics to form trigger cells. The summed signals are then digitized using fast ADCs. Simulation studies suggest that 8-bit ADCs give sufficient precision and dynamic range for the level-1 trigger. Look-up tables are included in the design to remove pedestals, apply calibration constants and convert to transverse energy units, although these will not be implemented for the initial beam tests in 1992. The 8-bit digital data are transmitted to the pipelined trigger processor, which performs the algorithms described above, using 8 parallel lines at 67 MHz.

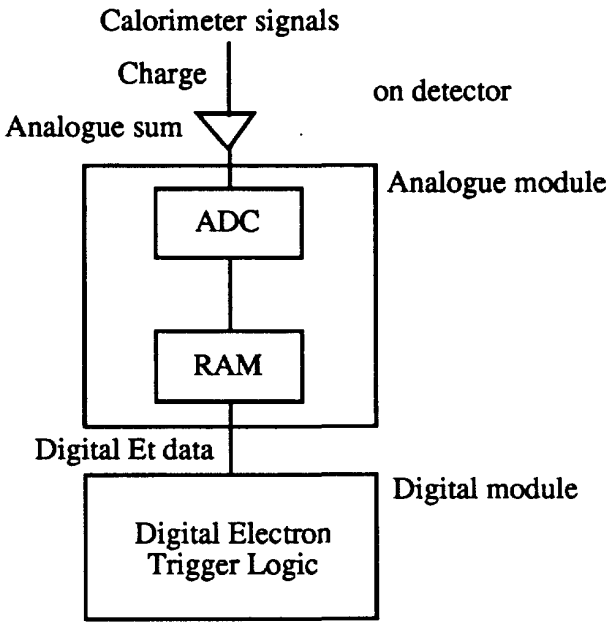

Figure A1.2: trigger system

The modularity of the electron/photon part of the system is shown in Fig. A1.3. Each ASIC performs the cluster algorithm for one trigger cell ("reference cell"), but needs data from all 16 cells in the window. Thus, 16×8 = 128 pins are needed for the input data alone. The large number of input connections limits the amount of processing that can be performed on each ASIC. Indeed, in the present design only about a third of the available logic gates are used. We are currently using a 0.8 micron gate array from Fujitsu with 20,000 gates in a 179-pin package.

The ASICs are mounted on multi-layer PCBs, each board viewing a 6×6 area of the calorimeter as shown in Fig. A1.3. A board contains 9 ASICs, each of which processes one reference cell; in this arrangement, each trigger ADC channel is connected to four PCB modules. The data from the digitization system are transmitted to the processor boards using a custom backplane with high-density connectors. These input data require 36×8 = 288 connections which occupy a significant fraction of the available connector space. The amount of logic that can be accommodated on each board is presently limited by the number of backplane connections and by the area of board occupied by the large ASIC packages. We are currently studying the possibility of multiplexing the 8-bit input data onto 4-bit interconnections at twice the rate. Initial simulation results indicated that this would be possible with the gate array technology we are currently using.

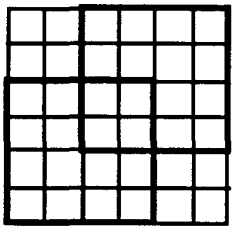

Overlapping, sliding windows in two dimensions:
- One chip for each em trigger channel in the calorimeter.
- Each ASIC uses data from 16 em channels in a 4 ×4 area of the calorimeter.

Figure A1.3: Modularity

One can make a rough estimate of the size of such a trigger system extrapolated to an LHC experiment. About 450 modules, each containing 9 ASICs, would be required to instrument a 4000-channel em calorimeter trigger. Each crate might contain ten such modules, the remaining space being used for ADCs, etc. Thus, the calorimeter trigger would occupy ≈ 45 crates. We consider it very important to reduce the size of the calorimeter trigger processor, since it will strongly affect the cost and the reliability of the level-1 system. However, we note that very large trigger processor systems are being considered for use in the SDC experiment at the SSC.

From the above discussion, it is apparent that the size of the trigger system is limited by the large amount of data that has to be input to the processing electronics. A more compact trigger system can only be achieved if this problem is overcome, either by more advanced interconnection technology or by improved architectures or both. The density of logic in the ASICs can be substantially increased using available technology, and is not yet a limiting factor. As discussed in the text and in Appendices 2 and 3, we are optimistic that a significant reduction in the size of the calorimeter trigger processor is possible.

Appendix 2:

Data compression and zero suppression in the bit-parallel calorimeter trigger design

We are investigating the possibility of reducing the number of connections between the digitization system and the processor system by using serial transmission of compressed, zero-suppressed data. Physics simulation studies show that the occupancy of both the electromagnetic and the hadronic calorimeters will be under 10% for an E_T = 1 GeV threshold and using a granularity of $\Delta\eta\times\Delta\phi = 0.1\times0.1$. Furthermore, the E_T values which remain after zero suppression are predominantly close to the threshold. Initial studies suggest that data compression algorithms can reduce the average word length from 8 to ~ 3 bits.

An example of such a data transmission scheme is shown in Fig. A2. The 8-bit E_T data from the digitization system are set to zero if the bunch-crossing identification logic identifies the energy as being due to a bunch crossing different from the present one. The data are then tagged using the bunch-crossing number modulo 16, and zero suppression is applied using a threshold of E_T ~ 1 GeV. The remaining 8-bit data are then compressed, for example using Huffman coding. A derandomizing buffer is included to average out fluctuations in the rate of arrival of the data. The E_T values are then transmitted to the processor system on one or more serial lines.

At the receiver end, the compressed data are expanded to 8 bits and placed in a FIFO buffer. Tag matching is used to re-establish synchronization prior to bit-parallel data processing. This tag-matching scheme is similar to the one used in the MEC2 project of the CERN micro-electronics group.

Preliminary simulation studies suggest that satisfactory performance can be achieved using a 32-bit deep derandomizing buffer and serial data transmission at 134 Mbits/s on two parallel lines. This introduces a maximum latency of eight bunch crossings or 120 ns. The fraction of data which are lost because the buffer overflows is very small. Using pessimistic assumptions about the occupancy and the average word length after data compression, we estimate a loss of high-E_T data values at the level of 0.1%.

More detailed studies will be made of data transmission schemes using compression and zero suppression algorithms. This will require physics simulation studies, computer modelling of the transmission system and micro-electronic design studies. Studies will also be made of the implementation of a calorimeter trigger system incorporating these data transmission schemes. We hope that it will be possible to overcome the interconnection problems encountered in the bit-parallel scheme described

in Appendix 1 in this way, and that the processor system can be reduced to a manageable size.

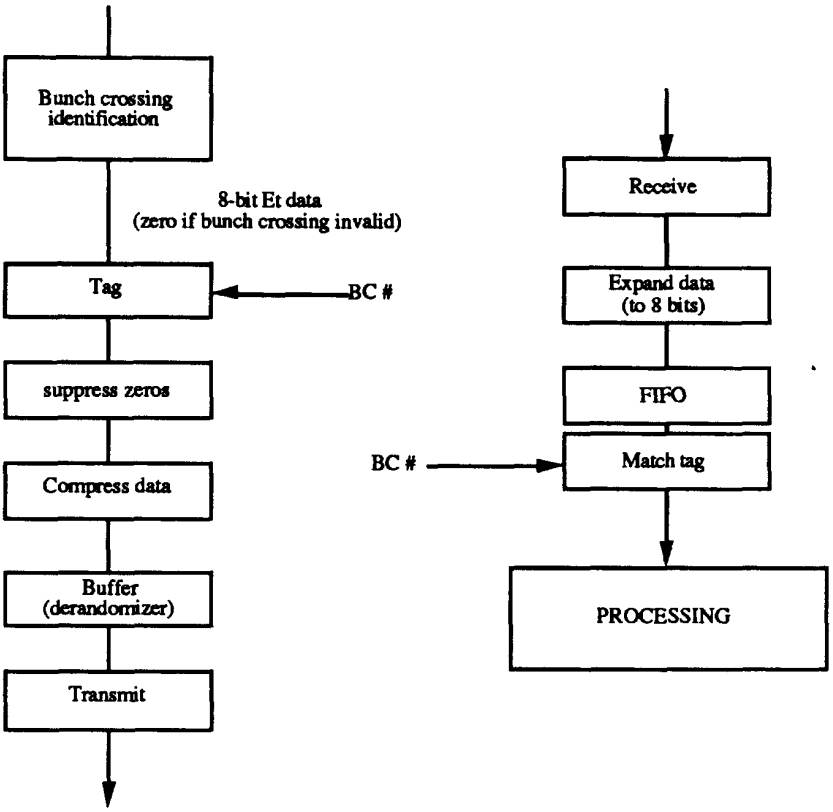

Figure A2: Transmission of compressed, zero-suppressed data.

Appendix 3:

A design study for a bit-serial calorimeter trigger processor

The bit-parallel conceptual design described in Appendix 1 would lead to a very large processor system if extrapolated to a full-size system for an LHC experiment. We estimate that ≈ 45 crates of electronics would be required, limited by inter-crate and inter-module interconnection density. This interconnection problem in the bit-parallel design described above can be greatly reduced by adopting an alternative bit-serial design.

In the bit-serial design, the data from the digitization system are transmitted to the processor system serially. The processor also processes the data serially (e.g. using serial addition and comparison operations), which results is a compact micro-electronic design. These bit-serial operations are more efficient than their bit-parallel counterparts, since they are not hampered by problems caused by carry propagation. It may thus be possible to operate such an implementation at a significantly higher clock-rate. The bit-serial solution also relieves the communication problems within the processor which were caused by pin count restrictions. It may therefore be possible to design processing ASICs that can perform all the required trigger calculations for a 6×6 trigger cell area in the em calorimeter.

The E_T value for each trigger channel is received serially at a rate of one or more bits every 15 ns. If each data word is M bits long and N bits are transmitted every 15 ns, $L = \frac{M}{N}$ bunch crossings will be required to transmit the data. This latency, inherent to the bit-serial design, can be compensated by introducing a farm of bit-serial trigger processors. A switching system is used to route data to the different trigger processors, each of which processes one event in L. Figure A3 shows an example of a system with $M = 12$ (e.g. 8-bit E_T data, 3 flags for bunch-crossing identification, etc, and 1 framing bit) and $N = 4$. We note that very high-speeds are now possible with standard CMOS [13].

While the bit serial approach is more demanding from the micro-electronics point of view, it has a number of advantages. The reduced number of ASICs should lead to increased reliability. Extending the farm by adding extra processors introduces redundancy which can be used to further increase the system reliability.

We plan to make a full design study for a bit-serial implementation of the level-1 calorimeter trigger. A tentative but fairly detailed design specification, which already exists, will be checked by simulation to evaluate consistency and performance. The system design currently being considered is based on optical readout of trigger cells using high-speed serial lines. The fibres are organized in groups of 36 (6×6 trigger

cells) which are sent to processing modules where the trigger algorithm is performed. Information from surrounding trigger cells is included to form expanded groups (9×9 trigger cells), using a passive optical fanout to share the data between processing modules. This provides sufficient information to evaluate isolation criteria around each of the 36 trigger cells.

All processing is performed in dedicated ASICs. This includes jet cluster-finding and the calculation of missing transverse energy, as well as the electromagnetic cluster-finding described above. The locations of clusters are stored and may be retrieved as regions of interest for level-2 processing if the event is accepted by the level-1 trigger. This also allows the possibility of upgrading the system by installing modified ASICs with improved algorithms.

A technical document on the bit-serial design is available [10].

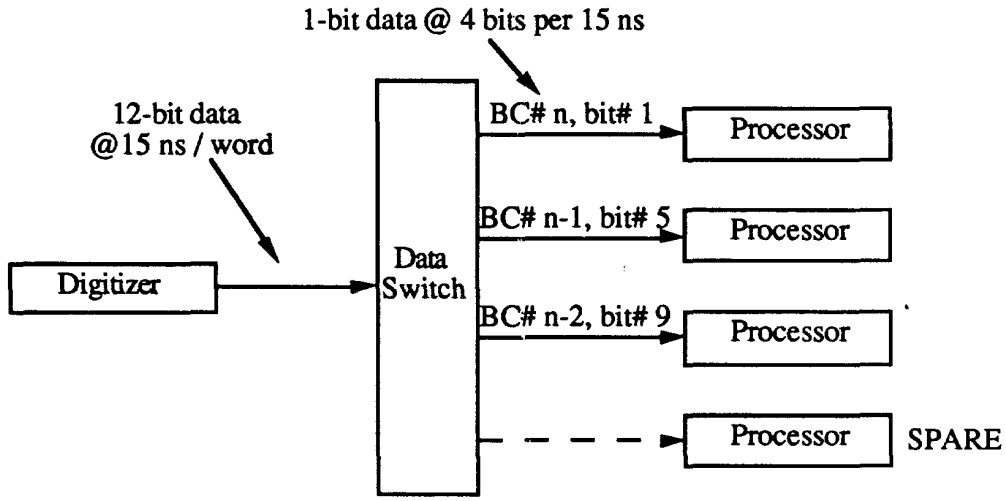

Figure A3: Farm of bit-serial processors.