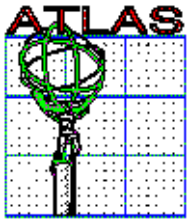

Level-1 Trigger

Technical Design Report

Issue: 1
Revision: 1
Reference: ATLAS TDR–12
Created: 24 June 1998
Last modified: 20 August 1998
Prepared By: ATLAS Level-1 Trigger Group

ATLAS Collaboration

Armenia
Yerevan Physics Institute, Yerevan

Australia
Research Centre for High Energy Physics, Melbourne University, Melbourne
University of Sydney, Sydney

Austria
Institut für Experimentalphysik der Leopold-Franzens-Universität Innsbruck, Innsbruck

Azerbaijan Republic
Institute of Physics, Azerbaijan Academy of Science, Baku

Republic of Belarus
Institute of Physics of the Academy of Science of Belarus, Minsk
National Centre of Particle and High Energy Physics, Minsk

Brazil
Universidade Federal do Rio de Janeiro, COPPE / EE / IF, Rio de Janeiro

Canada
University of Alberta, Edmonton
Department of Physics, University of British Columbia, Vancouver
University of Carleton / C.R.P.P., Carleton
Group of Particle Physics, University of Montreal, Montreal
Department of Physics, University of Toronto, Toronto
TRIUMF, Vancouver
University of Victoria, Victoria

CERN
European Laboratory for Particle Physics (CERN), Geneva

Czech Republic
Academy of Sciences of the Czech Republic, Institute of Physics and Institute of Computer Science, Prague
Charles University, Faculty of Mathematics and Physics, Prague
Czech Technical University in Prague, Faculty of Nuclear Sciences and Physical Engineering, Faculty of Mechanical Engineering, Prague

Denmark
Niels Bohr Institute, University of Copenhagen, Copenhagen

Finland
Helsinki Institute of Physics, Helsinki

France
Laboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP), IN2P3-CNRS, Annecy-le-Vieux
Université Blaise Pascal, IN2P3-CNRS, Clermont-Ferrand
Institut des Sciences Nucléaires de Grenoble, IN2P3-CNRS-Université Joseph Fourier, Grenoble
Centre de Physique des Particules de Marseille, IN2P3-CNRS, Marseille
Laboratoire de l'Accélérateur Linéaire, IN2P3-CNRS, Orsay
LPNHE, Universités de Paris VI et VII, IN2P3-CNRS, Paris
CEA, DSM / DAPNIA, Centre d'Etudes de Saclay, Gif-sur-Yvette

Republic of Georgia
Institute of Physics of the Georgian Academy of Sciences and Tbilisi State University, Tbilisi

Germany
Physikalisches Institut, Universität Bonn, Bonn
Institut für Physik, Universität Dortmund, Dortmund
Fakultät für Physik, Albert-Ludwigs-Universität, Freiburg
Institut für Hochenergiephysik der Universität Heidelberg, Heidelberg
Institut für Physik, Johannes-Gutenberg Universität Mainz, Mainz
Lehrstuhl für Informatik V, Universität Mannheim, Mannheim
Sektion Physik, Ludwig-Maximilian-Universität München, München
Max-Planck-Institut für Physik, München
Fachbereich Physik, Universität Siegen, Siegen
Fachbereich Physik, Bergische Universität, Wuppertal

Greece
Athens National Technical University, Athens
Athens University, Athens
High Energy Physics Department and Department of Mechanical Engineering, Aristotle University of Thessaloniki, Thessaloniki

Israel
Department of Physics, Technion, Haifa
Raymond and Beverly Sackler Faculty of Exact Sciences, School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv
Department of Particle Physics, The Weizmann Institute of Science, Rehovot

Italy
Dipartimento di Fisica dell' Università della Calabria e I.N.F.N., Cosenza
Laboratori Nazionali di Frascati dell' I.N.F.N., Frascati
Dipartimento di Fisica dell' Università di Genova e I.N.F.N., Genova
Dipartimento di Fisica dell' Università di Lecce e I.N.F.N., Lecce
Dipartimento di Fisica dell' Università di Milano e I.N.F.N., Milano
Dipartimento di Scienze Fisiche, Università di Napoli 'Federico II' e I.N.F.N., Napoli
Dipartimento di Fisica Nucleare e Teorica dell' Università di Pavia e I.N.F.N., Pavia
Dipartimento di Fisica dell' Università di Pisa e I.N.F.N., Pisa
Dipartimento di Fisica dell' Università di Roma 'La Sapienza' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Roma 'Tor Vergata' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Roma 'Roma Tre' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Udine, Gruppo collegato di Udine I.N.F.N. Trieste, Udine

Japan
Department of Information Science, Fukui University, Fukui
Hiroshima Institute of Technology, Hiroshima
Department of Physics, Hiroshima University, Higashi-Hiroshima
KEK, High Energy Accelerator Research Organisation, Tsukuba
Department of Physics, Faculty of Science, Kobe University, Kobe
Department of Physics, Kyoto University, Kyoto
Kyoto University of Education, Kyoto-shi
Department of Electrical Engineering, Nagasaki Institute of Applied Science, Nagasaki
Naruto University of Education, Naruto-shi
Department of Physics, Faculty of Science, Shinshu University, Matsumoto
International Center for Elementary Particle Physics, University of Tokyo, Tokyo
Physics Department, Tokyo Metropolitan University, Tokyo
Department of Applied Physics, Tokyo University of Agriculture and Technology, Tokyo

Morocco
Faculté des Sciences Aïn Chock, Université Hassan II, Casablanca, and Université Mohamed V, Rabat

Netherlands
FOM - Institute SAF NIKHEF and University of Amsterdam / NIKHEF, Amsterdam
University of Nijmegen / NIKHEF, Nijmegen

Norway
University of Bergen, Bergen
University of Oslo, Oslo

Poland
Henryk Niewodniczanski Institute of Nuclear Physics, Cracow
Faculty of Physics and Nuclear Techniques of the University of Mining and Metallurgy, Cracow

Portugal
Laboratorio de Instrumentação e Física Experimental de Partículas (University of Lisboa, University of Coimbra, University Católica-Figueira da Foz and University Nova de Lisboa), Lisbon

Romania
Institute of Atomic Physics, Bucharest

Russia
Institute for Theoretical and Experimental Physics (ITEP), Moscow
P.N. Lebedev Institute of Physics, Moscow
Moscow Engineering and Physics Institute (MEPhI), Moscow
Moscow State University, Institute of Nuclear Physics, Moscow
Budker Institute of Nuclear Physics (BINP), Novosibirsk
Institute for High Energy Physics (IHEP), Protvino
Petersburg Nuclear Physics Institute (PNPI), Gatchina, St. Petersburg

JINR
Joint Institute for Nuclear Research, Dubna

Slovak Republic
Bratislava University, Bratislava, and Institute of Experimental Physics of the Slovak Academy of Sciences, Kosice

Slovenia
Jozef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana

Spain
Institut de Física d'Altes Energies (IFAE), Universidad Autónoma de Barcelona, Bellaterra, Barcelona
Physics Department, Universidad Autónoma de Madrid, Madrid
Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia - CSIC, Valencia

Sweden
Fysiska institutionen, Lunds universitet, Lund
Royal Institute of Technology (KTH), Stockholm
University of Stockholm, Stockholm
Uppsala University, Department of Radiation Sciences, Uppsala

Switzerland
Laboratory for High Energy Physics, University of Bern, Bern
Section de Physique, Université de Genève, Geneva

Turkey
Department of Physics, Ankara University, Ankara
Department of Physics, Bogaziçi University, Istanbul

United Kingdom
School of Physics and Astronomy, The University of Birmingham, Birmingham

Cavendish Laboratory, Cambridge University, Cambridge
Department of Physics and Astronomy, University of Edinburgh, Edinburgh
Department of Physics and Astronomy, University of Glasgow, Glasgow
Department of Physics, Lancaster University, Lancaster
Department of Physics, Oliver Lodge Laboratory, University of Liverpool, Liverpool
Department of Physics, Queen Mary and Westfield College, University of London, London
Department of Physics, Royal Holloway and Bedford New College, University of London, Egham
Department of Physics and Astronomy, University College London, London
Department of Physics and Astronomy, University of Manchester, Manchester
Department of Physics, Oxford University, Oxford
Rutherford Appleton Laboratory, Chilton, Didcot
Department of Physics, University of Sheffield, Sheffield

United States of America
State University of New York at Albany, New York
Argonne National Laboratory, Argonne, Illinois
University of Arizona, Tucson, Arizona
Department of Physics, The University of Texas at Arlington, Arlington, Texas
Lawrence Berkeley Laboratory and University of California, Berkeley, California
Department of Physics, Boston University, Boston, Massachusetts
Brandeis University, Department of Physics, Waltham, Massachusetts
Brookhaven National Laboratory (BNL), Upton, New York
University of Chicago, Enrico Fermi Institute, Chicago, Illinois
Nevis Laboratory, Columbia University, Irvington, New York
Department of Physics, Duke University, Durham, North Carolina
Department of Physics, Hampton University, Virginia
Department of Physics, Harvard University, Cambridge, Massachusetts
Indiana University, Bloomington, Indiana
University of California, Irvine, California
Massachusetts Institute of Technology, Department of Physics, Cambridge, Massachusetts
University of Michigan, Department of Physics, Ann Arbor, Michigan
Michigan State University, Department of Physics and Astronomy, East Lansing, Michigan
University of New Mexico, New Mexico Center for Particle Physics, Albuquerque
Physics Department, Northern Illinois University, DeKalb, Illinois
Department of Physics and Astronomy, University of Oklahoma
Department of Physics, University of Pennsylvania, Philadelphia, Pennsylvania
University of Pittsburgh, Pittsburgh, Pennsylvania
Department of Physics and Astronomy, University of Rochester, Rochester, New York
Institute for Particle Physics, University of California, Santa Cruz, California
Department of Physics, Southern Methodist University, Dallas, Texas
State University of New York at Stony Brook, Stony Brook, New York
Tufts University, Medford, Massachusetts
High Energy Physics, University of Illinois, Urbana, Illinois
Department of Physics, Department of Mechanical Engineering, University of Washington, Seattle, Washington
Department of Physics, University of Wisconsin, Madison, Wisconsin

Acknowledgements

The Editors would like to thank Mario Ruggier for preparing the FrameMaker template upon which this document is based. The Editors also warmly thank Richard Cook, Michele Jouhet, Susan Leech-O'Neale and Janice Navarria for copy-editing the document, as well as the CERN Print-shop staff for the helpful, friendly and efficient service.

The level-1 trigger groups thank the rest of the ATLAS Collaboration for their help and advice in the preparation of this TDR, and for their cooperation in beam test evaluations of the trigger. In particular we thank the other ATLAS sub-systems for their assistance in reviewing the design presented here. We also wish to acknowledge the help of B.G. Taylor and the RD12 project group with whom we work closely on the Timing, Trigger and Control system.

The participating institutes would like to express their gratitude to their support staff for the invaluable help in designing and prototyping the many parts of the level-1 trigger. Finally, we wish to thank the following individuals who participated in the demonstrator prototyping work upon which the level-1 trigger design is based: A. Agapow, R. Carney, J. Leake, G.C. Lewis, A. J. Maddox, A. Mass, B. Niemann, D. Rees.

Revision Record

Issue	Revision	Date	Comments
1	0	24 June 1998	Original printing
1	1	20 August 1998	Editorial changes

Table Of Contents

1 Introduction

1.1 Purpose and scope of the TDR

The purpose of this document is to describe the design of the ATLAS level-1 (LVL1) trigger, as well as the plans for its construction, installation and commissioning, and operation. The design is based on an extensive programme of R&D and demonstrator prototyping work performed within the ATLAS, RD12 [1-1], and RD27 [1-2] projects since 1991. Most of the critical elements of the design have already been successfully demonstrated. Where this is not the case, the TDR describes the remaining demonstrator prototyping work required to validate the design. The design builds and improves upon the one presented in the ATLAS Technical Proposal [1-3].

This TDR documents the LVL1 trigger design at the system level, describing functional components at the level of electronic modules (printed circuit boards, multichip modules, application-specific integrated circuits) and their interconnections. The TDR includes a summary of trigger performance results that justify the choices that were made in the design. More details on the trigger performance of the LVL1 trigger, and of the Trigger, DAQ and Event Filter system as a whole, can be found in an accompanying document [1-4]. Another document [1-5] presents a technical progress report and workplan for the rest of the Trigger / DAQ / Event Filter / DCS system.

In order to keep the TDR to a reasonable size, not all of the details are given in this document. Where appropriate, reference is made to ATLAS internal notes and other documents.

1.2 Organization of the TDR

The level-1 trigger TDR is organized as follows. Chapter 2 gives a general description of the LVL1 trigger system, giving an overview of its architecture in the wider context of the ATLAS Trigger / DAQ system as a whole. It explains some of the general requirements on the LVL1 trigger, and discusses some of the more important implementation issues. Chapter 2 can be considered as an ‘executive summary’.

Chapters 3–16 describe the main hardware elements of the LVL1 trigger system: the calorimeter trigger (Chapters 3–8), the muon trigger (Chapters 9–14), the central trigger processor (Chapter 15) and the timing, trigger and control distribution system (Chapter 16). These are followed by a number of chapters that address general issues related to the LVL1 trigger: software (Chapter 17), a summary of the overall trigger latency (Chapter 18), the strategies for setting up the timing of the experiment (Chapter 19) and for handling deadtime (Chapter 20).

The remaining chapters address the issues of installation, access and maintenance (Chapter 21), safety (Chapter 22) and project organization and management (Chapter 23). Chapter 23 also summarizes the schedule and cost.

1.3 References

1-1 *RD12 Status Report*, CERN / LHCC / 97–29, April 1997.

1-2 *RD27 Status Report*, CERN / LHCC / 97–57, October 1997.

1-3 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, December 1994.

1-4 *ATLAS Trigger Performance Status Report*, CERN / LHCC / 98–15, June 1998.

1-5 *ATLAS DAQ, EF, LVL2 and DCS Technical Progress Report*, CERN / LHCC / 98–16, June 1998.

2 General description of the level-1 trigger system

2.1 ATLAS trigger and data-acquisition system overview

The ATLAS trigger and data-acquisition system is based on three levels of online event selection [2-1]. Each trigger level refines the decisions made at the previous level and, where necessary, applies additional selection criteria. Starting from an initial bunch-crossing rate of 40 MHz (interaction rate ~10^9 Hz at a luminosity of 10^{34} $cm^{-2}s^{-1}$), the rate of selected events must be reduced to ~100 Hz for permanent storage. While this requires an overall rejection factor of 10^7 against 'minimum-bias' processes, excellent efficiency must be retained for the rare new physics, such as Higgs boson decays, that is sought in ATLAS.

Figure 2-1 shows a simplified functional view of the Trigger/DAQ system. In the following, a brief description is given of some of the key aspects of the event-selection process.

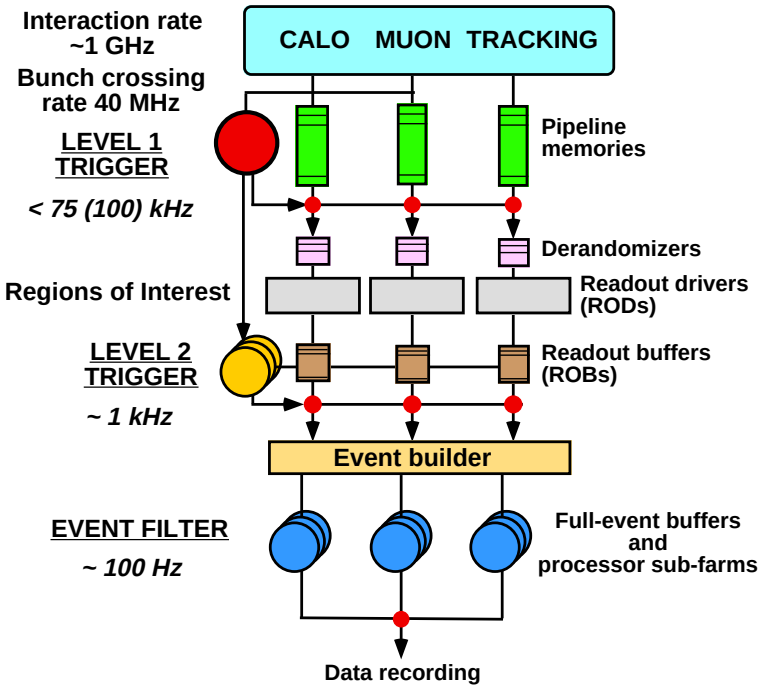

Figure 2-1 Block diagram of the Trigger/DAQ system.

The level-1 (LVL1) trigger described in this TDR makes an initial selection based on reduced-granularity information from a subset of detectors. High transverse-momentum (high-p_T) muons are identified using only the so-called Trigger chambers, resistive-plate chambers (RPCs) in the barrel, and thin-gap chambers (TGCs) in the endcaps [2-2]. The calorimeter selections are based on reduced-granularity information from all the ATLAS calorimeters (electromagnetic and hadronic; barrel, endcap and forward) [2-3], [2-4]. Objects searched for by the calorimeter trigger are high-p_T electrons and photons, jets, and taus decaying into hadrons, as well as large missing and total transverse energy. In the case of the electron/photon and hadron/tau triggers, isolation can be required. Information is available for a number of sets of p_T thresholds (generally 6–8 sets of thresholds per object type).

The missing and total scalar transverse energies used in the trigger are calculated by summing over trigger towers. However, a trigger on the scalar sum of jet transverse energies is also available.

The LVL1 trigger decision is based on combinations of objects required in coincidence or veto. As discussed in an accompanying document [2-5], most of the physics analyses that have been considered by ATLAS can be made using, at the trigger level, fairly simple selection criteria of a rather inclusive nature. However, the trigger implementation is flexible and it can be programmed to select events using more complicated signatures. Examples of LVL1 trigger 'menus' for 'high' (10^{34} cm^{-2}s^{-1}) and 'low' (10^{33} cm^{-2}s^{-1}) luminosity are shown in Tables 2-1 and 2-2. While the simple menus shown cover almost all of the mainstream discovery physics that ATLAS plans to study, in practice it is expected that real trigger menus will be significantly more complicated.

Table 2-1 Example of LVL1 trigger menu ($L = 10^{34}$ cm^{-2}s^{-1}).

Trigger	Rate (kHz)
Single muon, $p_T > 20$ GeV	4
Pair of muons, $p_T > 6$ GeV	1
Single isolated EM cluster, $E_T > 30$ GeV	22
Pair of isolated EM clusters, $E_T > 20$ GeV	5
Single jet, $E_T > 290$ GeV	0.2
Three jets, $E_T > 130$ GeV	0.2
Four jets, $E_T > 90$ GeV	0.2
Jet, $E_T > 100$ GeV AND missing $E_T > 100$ GeV	0.5
Tau, $E_T > 60$ GeV AND missing $E_T > 60$ GeV	1
Muon, $p_T > 10$ GeV AND isolated EM cluster, $E_T > 15$ GeV	0.4
Other triggers	5
Total	**~40**

The maximum rate at which the ATLAS front-end systems can accept LVL1 triggers is limited to 75 kHz (upgradable to 100 kHz). As indicated in Tables 2-1 and 2-2 (and discussed in Ref. [2-5]), the rates estimated in trigger performance studies, using trigger menus that meet the needs of the ATLAS physics programme, are about a factor of two below this limit. Given that there are large intrinsic uncertainties in the calculations, this safety factor is not over-generous. However, if necessary, rates could be significantly reduced without major consequences for the physics programme, for example by increasing the thresholds on some of the inclusive (single-object) triggers when operating at the highest luminosities, and by relying more heavily on multi-object triggers.

An essential requirement on the LVL1 trigger is that it should uniquely identify the bunch-crossing of interest. Given the short (25 ns) bunch-crossing interval, this is a non-trivial consideration. In the case of the muon trigger, the physical size of the muon spectrometer implies times-of-flight comparable to the bunch-crossing period. For the calorimeter trigger, a

Table 2-2 Example of LVL1 trigger menu ($L = 10^{33}$ cm^{-2}s^{-1}).

Trigger	Rate (kHz)
Single muon, $p_T > 6$ GeV	23
Single isolated EM cluster, $E_T > 20$ GeV	11
Pair of isolated EM clusters, $E_T > 15$ GeV	2
Single jet, $E_T > 180$ GeV	0.2
Three jets, $E_T > 75$ GeV	0.2
Four jets, $E_T > 55$ GeV	0.2
Jet, $E_T > 50$ GeV AND missing $E_T > 50$ GeV	0.4
Tau, $E_T > 20$ GeV AND missing $E_T > 30$ GeV	2
Other triggers	5
Total	~40

serious challenge is that the pulse shape of the calorimeter signals extends over many bunch crossings.

It is important to keep the latency (time taken to form and distribute the trigger decision) to a minimum. During this time information for all detector channels has to be retained in ‘pipeline’ memories. These memories are generally contained in custom integrated circuits, placed on or close to the detector, usually in inaccessible regions and in a high-radiation environment. The total number of detector channels, excluding the pixel detectors, exceeds 10^7. For reasons of cost and reliability, it is desirable to keep the pipeline lengths as short as possible. The LVL1 latency, measured from the time of the proton–proton collision until the trigger decision is available to the front-end electronics, is required to be less than 2.5 μs. In order to achieve this, the LVL1 trigger is implemented as a system of purpose-built hardware processors. The target latency for the LVL1 trigger is 2.0 μs (leaving 500 ns contingency).

Events selected by LVL1 are read out from the front-end electronics systems of the detectors into readout buffers (ROBs); present estimates foresee about 1700 ROBs in total. A large number of front-end electronics channels are multiplexed into each ROB. Intermediate buffers, labelled ‘derandomizers’ in Figure 2-1, average out the high instantaneous data rate at the output of the pipeline memories to match the available input bandwidth of the readout drivers (RODs).

All of the data for the selected bunch crossing from all of the detectors are held in the ROBs either until the event is rejected by the level-2 (LVL2) trigger (in which case the data are discarded) or, in case the event is accepted by LVL2, until the data have been successfully transferred by the DAQ system to storage associated with the Event Filter (which makes the third level of event selection). The process of moving data from the ROBs to the Event Filter (EF) is called event building. Whereas before event building each event is composed of many fragments, with one fragment in each ROB, after event building the full event is stored in a single memory accessible by an EF processor.

The LVL2 trigger makes use of ‘region-of-interest’ (RoI) information provided by the LVL1 trigger. This includes information on the position (η and ϕ, where η is pseudorapidity and ϕ is azimuthal angle) and p_T range of candidate objects (high-p_T muons, electrons / photons, hadrons / taus, jets), and energy sums (missing-E_T vector and scalar E_T value, where E_T is

transverse energy). The RoI data are sent by LVL1 to LVL2, for all events selected by the LVL1 trigger, using a dedicated data path. Using the RoI information, the LVL2 trigger selectively accesses data from the ROBs, moving only the data that are required in order to make the LVL2 decision. The LVL2 trigger has access to all of the event data, if necessary with the full precision and granularity. However, typically only data from a small fraction of the detector, corresponding to limited regions centred on the objects indicated by the LVL1 trigger, are needed by the LVL2 trigger. Hence, usually only a few per cent of the full event data are required thanks to the RoI mechanism.

Some of the signatures used at LVL2 to reduce the rate are discussed in the following; the reader is referred to Ref. [2-5] for more details. It is expected that LVL2 will reduce the rate to ~1 kHz. However, in contrast to the 75 kHz (upgradable to 100 kHz) limit for LVL1 that comes from the design of the detector front-end electronics, this is not a hard number. Optimization of the sharing of the selection task between LVL2 and the EF is being studied in the next phase of the project. The latency of the LVL2 trigger is variable from event to event; it is expected to be ~1–10 ms.

In the case of muon triggers, rejection power at LVL2 comes from sharpening (and where necessary raising) the p_T threshold compared to LVL1, and from applying isolation requirements. Sharper p_T thresholds are possible by using the precision muon chambers (MDTs and CSCs) [2-2] and also the inner detector. The isolation requirements use calorimeter information, demanding little E_T in a region around the muon candidate.

For isolated electrons, rejection power at LVL2 comes from using the full-granularity calorimeter information and requiring a matching high-p_T charged track in the inner detector; the transition-radiation signature gives additional rejection power. For photons, less rejection power is possible than in the case of electrons since the inner detector cannot be used. (Given the relatively high probability for photon conversion in ATLAS, it is not planned to use a track veto for the photon trigger.) However, for the important physics channel $H \rightarrow \gamma\gamma$, the trigger can require a pair of photons, with a rejection factor for each γ compared to LVL1 due to the use of high-precision, high-granularity calorimeter information.

For the hadron / tau trigger, rejection at LVL2 is achieved using the full-granularity calorimeter information and the inner detector. A localized, isolated (hadronic) calorimeter cluster with a matching high-p_T track is required.

In the case of jets, much less rejection power is possible. Jets are the dominant high-p_T process at the LHC, and the threshold behaviour of the LVL1 trigger is reasonably sharp. Hence for jet triggers LVL2 must either increase the threshold or make additional requirements in order to significantly reduce the rate. The possibility of identifying b-quark jets at LVL2 using inner-detector information for tagging is under study.

Concerning the energy-sum triggers (missing E_T, total scalar E_T), only limited improvement is possible using the RoI mechanism. The energy-sum values from LVL1 are provided to LVL2 and refinements can be made to correct, e.g., for high-p_T muons (the LVL1 missing-E_T trigger uses only calorimeter information, so muons contribute to the observed missing E_T) or for saturated trigger-tower signals. The possibility of performing a full missing-E_T recalculation at LVL2 for a small subset of events remaining after other LVL2 selection criteria have already been applied is being investigated.

The LVL1 trigger makes available RoI information for all of the objects that contributed to the event being selected; these are called primary RoIs. Furthermore, in order to allow additional

requirements to be made at LVL2, the LVL1 trigger can provide RoI information for objects that did not contribute to the selection of the event. Such RoIs, typically for objects of relatively low p_T, are called secondary RoIs.

After LVL2, a last stage of selection is performed in the EF. Here the algorithms will be based on offline code. The EF must reduce the rate to a level suitable for permanent storage, currently assumed to be ~100 Hz for full events of size ~1 Mbyte.

The status of work on the LVL2 trigger, DAQ, Event Filter and Detector Control systems, together with plans for work to be completed before submission of the Technical Proposal for these subsystems, is documented in a Technical Progress Report [2-6]. Details of the trigger performance studies for LVL2, as well as for LVL1, are also documented [2-5].

An important consideration is to define interfaces with other systems sufficiently early to avoid wasting effort on incompatible designs. The interface between the ATLAS front-end systems and the Trigger/DAQ system is specified in a requirements document [2-7] that has been endorsed by all systems. This defines (at a functional level) the interaction between the LVL1 trigger and the detector systems, including issues of trigger latency, rate capability, and permitted levels of data loss and deadtime. It does not cover the special cases of the provision of trigger signals from the muon Trigger chambers and from the calorimeters to the LVL1 trigger; these are addressed in this document and also in the corresponding detector TDRs that have already been approved, and in User Requirements Documents (URDs) [2-8], [2-9]. The requirements document [2-7] also defines the interaction and sharing of responsibilities between the detector systems on the one hand, and the LVL2 trigger and DAQ on the other hand, at the level of the ROBs. The interface between the LVL1 and LVL2 triggers is discussed in this TDR and also in Ref. [2-6].

2.2 Overview of level-1 trigger system

Figure 2-2 is a context diagram for the LVL1 trigger system, showing the interfaces to other systems. Input data/signals are received from the muon trigger chambers and from the calorimeters. The trigger decision has to be made available to the front end and readout systems for all the ATLAS subsystems — as shown, this is done using the Timing, Trigger and Control (TTC) distribution system (see Chapter 16). Additional information that can be sent over the TTC system are the event identifier number, the bunch-crossing number, and a trigger word that summarizes why the event was selected by the LVL1 system.

The LVL1 trigger is a system of synchronous, pipelined processors running at 40 MHz or multiples thereof. The clock signals used to drive the processing are derived from the LHC machine clock. Another signal received from the machine is a bunch-zero signal, sent once per turn of the machine (period 88 μs). This allows one to make an absolute identification of bunches within the machine.

The LVL1 trigger system can accept externally-generated trigger signals. These include calibration trigger signals associated, for example, with test-pulse generation in the calorimeters. Various other facilities included are related to calibration procedures, including the possibility to fire test pulses at well-defined points in the LHC cycle, and the arbitration of the use by detector systems of gaps in the LHC bunch structure for calibration triggers.

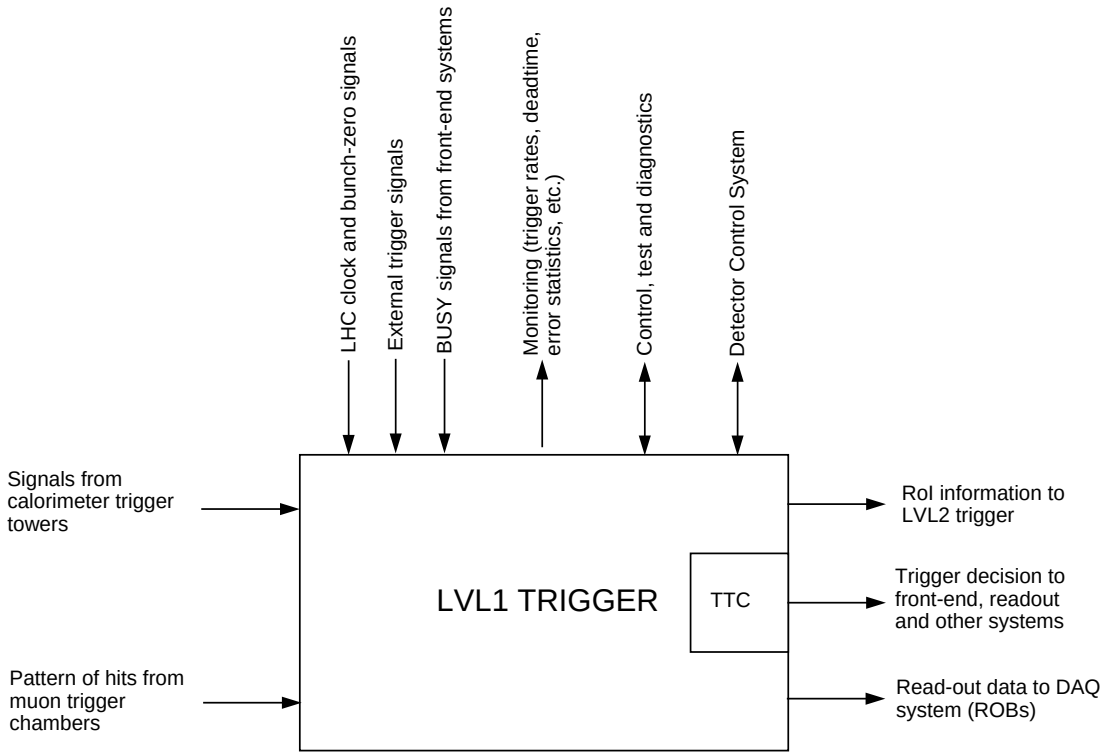

Figure 2-2 Context diagram for the LVL1 trigger system.

The LVL1 system receives inputs ('BUSY' signals) that can impose deadtime, for example if buffers become full in the readout systems. These are also used during initialization to prevent triggers being generated until the readout systems are ready.

The LVL2 trigger receives the accept / reject signal from LVL1. For accepted events, it receives in addition information on the candidate objects found by the LVL1 trigger which it uses to drive the RoI processing. More detailed information from LVL1, including input data, and intermediate and final results of the trigger processing, are stored in ROBs. These data are read out by the DAQ system for events that are selected by LVL2; the data may also be selectively accessed by the LVL2 trigger, as for detector data.

In addition, the LVL1 trigger is connected to the run control (an element of the DAQ system) and to the Detector Control System (DCS). The run control infrastructure is used for testing, monitoring and fault diagnosis, for running calibration programs, and for configuring the trigger system (e.g. loading all the parameters that control the operation of the trigger and specify the selection criteria). The run control system will be used to acquire data that are not associated with individual events, such as scaler information needed to monitor rates and deadtime. The functions of the DCS include control and monitoring of the electronics racks and crates (voltages, temperature, etc.).

As shown in Figure 2-3, the LVL1 trigger system is composed of a number of building blocks — the calorimeter trigger, the muon trigger, the Central Trigger Processor (CTP) and the Timing, Trigger and Control (TTC) system. These subsystems are functionally well-defined, and the interaction between them is limited. The interfaces are presently defined at the functional level, and their detailed physical implementation will be specified in the next phase of the project (this is considered to be straightforward). The compatibility of the different subsystems will be checked, prior to the start of construction, within the context of the ATLAS review process (i.e.

via Production Readiness Reviews). The functions of each of the subsystems are described in a bit more detail in the following, considering only the LVL1 trigger function *per se*.

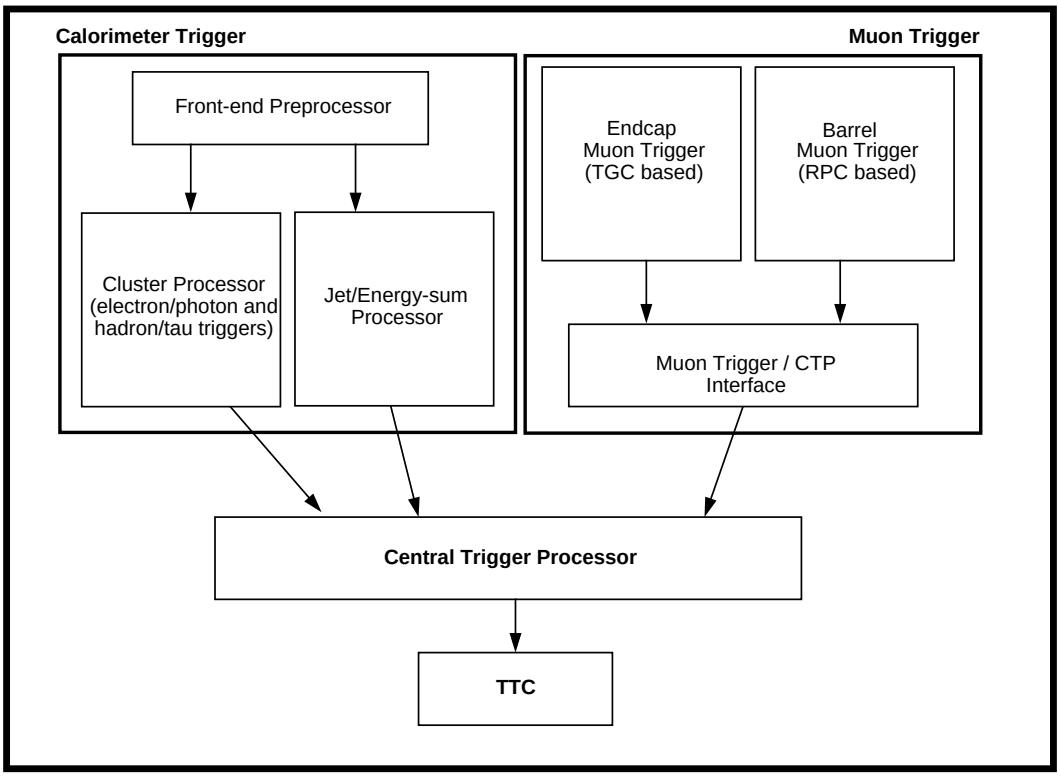

Figure 2-3 Block diagram of the LVL1 trigger system.

2.2.1 Calorimeter trigger

The inputs to the calorimeter trigger system are analogue signals from trigger towers that have a typical granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ (the granularity becomes coarser beyond $|\eta| \approx 2.5$). Separate tower signals are received from the electromagnetic and hadronic calorimeters, corresponding to a total of $\approx$ 7200 analogue input signals to the trigger system. The trigger towers are formed in the calorimeter front-end electronics by analogue summation over the corresponding calorimeter cells. Transmission to the calorimeter trigger system, located in USA15, is on individually-shielded twisted-pair cables.

In the calorimeter trigger system, the trigger-tower signals are digitized using a dedicated ADC system. Digital signal processing is applied to extract the E_T for calorimeter pulses and to assign it to the correct bunch crossing, since the shaped pulses from the calorimeters extend over several bunch-crossing periods. This part of the calorimeter trigger system is labelled 'front-end Preprocessor' in the figure.

The subsequent calorimeter-trigger processing is fully digital, and is divided into two parts. One part (labelled 'Cluster Processor' in the figure) performs a search for high-p_T electrons / photons and hadrons / taus using the full-granularity trigger-tower information from the Preprocessor. The other part (labelled 'Jet / Energy-sum Processor' in the figure) searches for high-E_T jets and calculates the missing-E_T and total scalar-E_T values.

For the electron / photon trigger, there are eight sets of thresholds that can be programmed independently; each set consists of a threshold on the E_T of the cluster, an isolation threshold on the surrounding E_T in the electromagnetic calorimeter, and a 'hadron-veto' threshold on the E_T in the associated hadron-calorimeter towers. Similarly, for the hadron / tau trigger, there are eight sets of thresholds that can be programmed independently; each set consists of a threshold on the E_T of the cluster, and isolation thresholds for the surrounding E_T in the electromagnetic and hadronic calorimeters. For the jet trigger there are eight thresholds that can be programmed independently and to which the E_T in 'jet windows' is compared. The jet-window size is also programmable.

Summation is performed over the trigger towers to calculate the missing-E_T vector and the total scalar E_T for the event. Thresholding is performed, with eight threshold values for the missing-E_T trigger (thresholds applied to modulus of missing-E_T vector), and four for the scalar-E_T trigger.

For each object type, information is sent to the CTP for each bunch crossing. In the case of the electron / photon, hadron / tau and jet triggers, this is the multiplicity of objects for each set of thresholds. For the energy-sum triggers, information is sent indicating which thresholds have been passed.

2.2.2 Muon trigger

The muon trigger receives as input the pattern of hit strips (and wire groups in the case of the TGC detectors) in the muon Trigger chambers. These data are produced by amplifier-shaper-discriminator circuits in the RPC and TGC front-end electronics. There are a total of more than 800 k input signals to the muon trigger system. The timing resolution is sufficiently good that the trigger can, with very high probability, identify the bunch crossing that contained the muon. The trigger searches for patterns of hits consistent with high-p_T muons originating from the interaction region. The logic provides six independently-programmable p_T thresholds. The output sent to the CTP for each bunch crossing is the multiplicity of muons for each of the six p_T thresholds (i.e. six multiplicity values).

As indicated in Figure 2-3, the muon trigger system is subdivided into a part specific to the RPC detectors, a part specific to the TGC detectors, and a part that combines information from the full system and prepares the input to the CTP. Some of the RPC- and TGC-specific logic is mounted on or near the detectors in the experimental cavern; the rest is located in the underground shielded counting room, USA15.

2.2.3 Central Trigger Processor

The role of the CTP is to combine the information for the different object types and to make the overall LVL1 accept / reject decision. Trigger menus can be programmed with up to 96 items, each item being a combination of requirements on the input data. The requirements can be very simple, demanding for example at least one muon for a given p_T threshold, or more complicated, for example requiring at least one electron above a specified p_T threshold, missing E_T above another specified threshold, and no high-p_T muons. The CTP can also combine the information on jet multiplicity versus threshold to estimate the jet-E_T sum, to which it can apply a threshold. The overall LVL1 decision will be to accept an event if any of the 96 menu items is satisfied, subject to prescale and deadtime considerations discussed below.

The CTP allows a prescale factor to be programmed individually for each of the menu items. This facility will be used to down-scale high-rate triggers, for example allowing data to be collected concurrently for jet production over a large E_T range; this is useful for QCD studies and also for monitoring the detector performance. The CTP is also responsible for deadtime control. Deadtime may be generated externally via a busy signal, for example if the LVL2 trigger is close to saturation. It will also be generated internally using algorithms that predict when overflow conditions are likely to occur in front-end systems. The simplest element of the deadtime logic is to prevent the occurrence of two triggers separated by fewer than five bunch crossings (125 ns). The CTP also contains a system of scalers used to monitor rates and deadtime.

2.2.4 Timing, Trigger and Control distribution system

The TTC system is responsible for distributing to the front-end systems a number of signals, including the LHC clock and the LVL1 trigger decision. The TTC backbone, which is described in this TDR, is based on the optical-broadcast system developed in the RD12 Collaboration [2-10]. The TTC system receives the LVL1 accept signal from the CTP and the LHC clock and bunch-zero signals from the machine. These and other signals are encoded and transmitted optically to the front-end systems, or to intermediate points at which a change is made to detector-specific protocols for TTC distribution.

2.3 Summary of requirements for the level-1 trigger

Detailed evaluations have been made of the requirements for the different parts of the LVL1 trigger system — the calorimeter trigger [2-9], the muon trigger [2-8], the CTP [2-11] and the TTC system [2-12]. A brief and mainly qualitative summary of some of the more important general requirements is given here. More details and quantitative requirements can be found in subsequent chapters of this TDR, and in the requirements documents for each of the subsystems.

2.3.1 General requirements

The primary function of the LVL1 trigger is to provide, for each bunch crossing, a signal specifying if the bunch crossing should be retained for further analysis, typically because it contains a potentially interesting physics signature. As discussed above, the decision is based on input signals received from the calorimeters and muon Trigger chambers. The trigger decision has to be distributed, along with other signals, to the front-end electronics of the detector systems. A secondary function of the LVL1 trigger is to provide RoI information to guide the LVL2 trigger.

2.3.2 Physics requirements

The physics performance of the trigger is discussed in Chapters 4 and 14 of this TDR for the LVL1 calorimeter and muon triggers, respectively, and in more detail in Ref. [2-5] which also addresses the LVL2 trigger. As stated above, the objects upon which the LVL1 trigger is based are high-p_T muons, electrons / photons, hadrons / taus and jets, and large missing and total

scalar E_T. Considerations include the p_T range over which the trigger must be able to operate, and the required acceptance and efficiency to find the objects, taking into account constraints on the acceptable trigger rate. Other issues are the ability to resolve nearby objects while not double-counting single objects; this is important for multi-object triggers such as the di-electron and di-muon triggers. Also important is the number of thresholds that can be concurrently used; typically different thresholds are used for single and multi-object triggers, and further thresholds are needed for prescaled triggers.

An indication of the required range of threshold settings can be obtained by considering the numerous studies that have been made to assess the physics capabilities of ATLAS. It turns out that the mainstream discovery physics is covered by fairly simple LVL1 trigger menus based on rather inclusive signatures, as shown in Tables 2-1 and 2-2. In practice, much more extensive menus will be used, including triggers for specialized physics studies (including Standard Model physics), and to collect samples for calibration and monitoring of the detectors and physics control samples (e.g. for background evaluation studies). However, for these additional menu items, often only limited statistics will be required, and prescale factors can be applied to limit the rate to acceptable levels.

The geometrical acceptance of the LVL1 trigger is driven by the design of the detector, where precision measurements in the calorimeters and the coverage of the inner detector are limited to pseudorapidity, $|\eta| < 2.5$. The muon, electron / photon and hadron / tau triggers are required to cover this range ($|\eta| < 2.4$ in the case of the muon trigger). For the jet trigger, the calorimeter trigger towers that are used extend up to $|\eta| < 3.2$, the edge of the endcap calorimeters. The missing and total scalar E_T triggers use all the calorimeters, giving a coverage of $|\eta| < 4.9$.

The LVL1 trigger has to take into consideration the need to provide RoI information to LVL2, with thresholds for flagging objects lower in p_T than those used actively in making the LVL1 trigger decision. The exact threshold requirements here will depend on the scope of the LVL2 trigger (compared to the Event Filter). Also important for LVL2 is the accuracy with which the position of the object within the detector can be specified by LVL1. More details on the requirements on the LVL1 trigger from LVL2 can be found in Ref. [2-13].

The trigger conditions (e.g. thresholds and multiplicity requirements) must be programmable to adapt to different luminosity conditions and changing physics requirements. There must be sufficient flexibility to cope with unforeseen background conditions or new physics.

Beyond the normal physics triggers, consideration has to be given to handling various kinds of special triggers needed for calibration and monitoring, and to understand the detector and background conditions. These include calibration triggers, such as test-pulse triggers associated with the different detector systems. They also include special triggers on cosmic rays and beam-halo particles (implemented using the LVL1 muon trigger system). In addition, provision has to be made for random triggers and for prescaling high-rate triggers as discussed above.

2.3.3 Measurement of trigger acceptance and efficiency

It will be important for physics studies to correct for acceptance and efficiency losses of the trigger. In contrast to the situation in e^+e^- machines such as LEP, where the interaction rates are relatively small and where all high-p_T events can therefore be accepted, the triggers at the LHC have to be extremely selective. In general, one cannot rely on a high level of redundancy in the trigger, with signal events being selected on the basis of multiple independent signatures.

(Some channels, such as $H \rightarrow l^+l^-l^+l^-$, would be selected by multiple independent signatures, but this is generally not the case.)

Despite the limitations imposed by the difficult environment of a hadronic machine, it is necessary to be able to measure the trigger efficiency. This can be achieved by determining the efficiency to trigger on objects as functions of p_T, η and ϕ. For multi-object triggers, these efficiencies can then be combined. Object efficiencies can be determined by using large samples of two-object events, where either object could have triggered. Such samples can, if necessary, be collected using prescaled single-object triggers with low p_T thresholds.

2.3.4 Rate and background rejection

The design of the front-end electronics for the detector systems imposes constraints on the rate of LVL1 triggers that can be accepted. The average rate must not exceed 75 kHz (upgradable to 100 kHz). Furthermore, the interval between successive LVL1 triggers must be at least five bunch crossings (125 ns), corresponding to 1% deadtime. As discussed later in this document in Chapter 20, more complicated deadtime logic is provided that allows for high- and low-priority triggers.

The design of the trigger must take into account the need to reduce the rate of triggers due to background processes to acceptable levels, while maintaining the required high efficiency for signal processes that are to be analysed. Adequate safety margins are required to allow for the large uncertainties in the rate calculations (typically of a factor of two or more).

The dominant background in the LVL1 trigger often comes from high-p_T jet production (except for the jet trigger where this is the signal). For example, jets may fake isolated electrons / photons if they fragment into leading π^0 particles. Similarly, background to the muon trigger comes from decays in flight of high-p_T charged pions and kaons. In the case of the muon trigger, the background radiation in the cavern also requires careful attention. In addition, possible false triggers due to cosmic rays and beam-halo particles have to be considered.

Also relevant to the discussion on rates is the threshold sharpness. Given that most of the processes that dominate the LVL1 rate, such as high-p_T jets and muons from b-quark decays, have steeply falling p_T distributions, it is desirable to make a sharp cut on p_T. Otherwise the rate may be dominated by objects that, with more detailed analysis using the full detector data, are found to have p_T below the nominal threshold.

2.3.5 Latency

A requirement that was already mentioned is that the latency should not exceed 2.5 μs. All the ATLAS front-end electronics systems are being designed and implemented with pipeline lengths sufficient to accommodate this latency. In order to keep some contingency, the target latency in designing the LVL1 trigger system is 2.0 μs. This contingency is needed to allow for unforeseen steps in the processing pipeline that may emerge in the detailed design, and for remaining uncertainties in the routing of cables carrying signals to or from the trigger. Note that care has been taken in specifying the latency requirements (see Ref. [2-7]) to allow for detector-specific treatment of TTC signals that may, in some cases, add extra latency but bring overall benefits.

2.3.6 Bunch-crossing identification

Another requirement that was already mentioned is the need for the trigger to uniquely identify the bunch crossing of interest. This is achieved for all the triggers that are considered; the probability to trigger on the wrong bunch crossing is expected to be negligibly small.

2.3.7 Timing calibration

An important consideration for the LVL1 trigger design (and also for all ATLAS front-end electronics systems) is the need to set up the timing under various operating conditions. The short bunch-crossing period, the large physical size of the apparatus, the huge number of electronic channels and units, and the inaccessibility of much of the electronics make this very challenging. Well thought-out strategies for setting up the timing need to be developed at an early stage so that the necessary tools can be included in the system designs.

It must be possible to set up the timing of the experiment for physics running and also for special runs, for example with cosmic rays or beam-halo particles. The timing also has to be set up for test and calibration runs with test pulses. These issues are discussed in more detail in Chapter 19.

2.3.8 Processing and transmission errors

False triggers may be caused by processing or transmission errors within the LVL1 system. For example, a single high-order bit being wrongly set in the E_{T} value for an electromagnetic-calorimeter trigger tower, due to a transmission error, would very likely result in a false electron / photon trigger. The rate of such false triggers is required to be much smaller than the rate of true triggers. Similarly, the loss of triggers due to processing or transmission errors is required to be very small.

An analysis of the effects of processing and transmission errors is required. Where necessary, error detection and protection are included in the design of the system. Also included in the design is the ability to disable parts of the system that have become faulty. This facility can be used to allow continued operation of the trigger, typically with only a small degradation in performance, until repairs can be made.

2.3.9 Operation of equipment in the experiment cavern

The special requirements for equipment that is installed in the ATLAS cavern are addressed. These include the need for radiation tolerance and for operation in a significant magnetic field, and constraints on the amount of heat that can be dissipated. They also include issues of reliability and maintainability for electronics that cannot be accessed easily (and not at all during running). Note that the only equipment in the cavern under the responsibility of the LVL1 trigger project is the on-detector part of the muon-trigger electronics. The calorimeter groups are responsible for sending to the USA15 shielded counting-room the analogue signals for the trigger towers.

2.3.10 Other requirements on the LVL1 trigger

There are numerous other requirements on the LVL1 trigger relating to commissioning and operating it *in situ*, including its connections to other parts of the Trigger / DAQ / DCS system. These are documented in the various requirements documents referred to above, and include the following:

- The system must read out the input data upon which the trigger decision was based, intermediate results, and the final trigger decision. This is required for use in monitoring the operation of the trigger and the quality of the input data to it, and is also used in trigger efficiency determination. The readout system has to respect the same requirements as for detector front-end systems (pipelined readout, etc.).
- A subset of the data that are read out has to be passed, via a separate path, to the LVL2 trigger. These data include information on all the candidate objects used in the RoI mechanism, as well as global quantities such as the missing-E_T vector and the total scalar-E_T value, and trigger words that specify why the event was selected by the LVL1 trigger.
- An extensive system (hardware and software) is required for control. This must be done in an integrated way across the different parts of the LVL1 trigger and more generally with the rest of the Trigger / DAQ system. A record of the trigger configuration, including details of selection criteria, calibration constants used, and other parameters, must be maintained. The time required to set up the trigger system or to change the trigger conditions must not be too long, so as to minimize loss of data.
- Provision has to be made for calibration of the system. Generally this requires interacting with the detector groups, for example for the use of test-pulse systems.
- Extensive capability must be built into the system for monitoring so that faults can be detected quickly. Online monitoring software can make extensive checks using the data that are read out from the trigger system. Scalers are required to monitor trigger rates and deadtime; provision has to be made to monitor some rates for individual bunches in the machine. The scaler data have to be read out separately from the event data. Finally, monitoring (using the DCS) is required of voltage levels, currents, temperatures, etc.
- The system must be designed with built-in test facilities to allow efficient and rapid detection and diagnosis of faults. This will be important both in the commissioning phase and during maintenance and repair operations.
- The failure rate and time required to perform repairs must be such that the resulting loss of data to the experiment is acceptably small. The system must be designed to be robust against localized failures. For example, it must be possible to prevent defective detector channels (noisy calorimeter cells or muon chamber strips) from seriously degrading the performance of the trigger.
- Maintenance requirements have been considered. Sufficient spares must be available taking into consideration the long lifetime of the experiment, and provision must be made to have adequate experts at CERN when the experiment is running.
- Standard safety requirements must be followed, for example concerning fire safety of cables.
- The ATLAS standard grounding rules should be followed where applicable.

2.3.11 Requirements on external systems

The requirements analyses that have been performed also identified numerous requirements on external systems, such as the calorimeter and muon detectors that provide the input signals, and also the DAQ and DCS systems that provide the infrastructure used to control, monitor, and read out the LVL1 system.

The requirements on the calorimeter and muon systems are discussed in detail in Chapters 3–8 and Chapters 9–14 of this TDR, respectively. These include a specification of the boundary of responsibility between the trigger and the detector systems, and a definition of the interface (including issues of detector granularity). Requirements are also specified on the properties of the input signals to the trigger system as discussed in the following paragraphs.

In the case of the calorimeters, requirements on the input signals cover issues of dynamic range, linearity, pulse shape, amplitude/transverse-energy calibration and timing, from the points of view of normalization, uniformity between cells, and stability. The requirements also address the issues of noise in and crosstalk between trigger towers, as well as the need to be able to remove individual 'noisy' cells from the trigger towers. It is required that the test-pulse systems of the calorimeter be available to generate test signals at the input to the calorimeter trigger for diagnostic and calibration purposes.

Requirements on the muon detector system include issues of optimizing the layout and granularity of the Trigger chambers taking into account the needs of the trigger, and of making sure that the absolute and relative positioning and alignment of the chambers is sufficiently precise that the performance of the trigger is not degraded. Other requirements relate to the timing uniformity of the input signals. It is also required that the test-pulse system of the Trigger chambers be available to generate test patterns at the input to the muon trigger.

2.4 Implementation

In the following chapters of this TDR, a detailed description is given of the LVL1 trigger architecture and the plans for implementing it, as well as results from an extensive R&D programme that has demonstrated the feasibility of the key elements of the design presented.

As discussed in Section 2.2 above, the LVL1 system is partitioned into the calorimeter trigger, the muon trigger, and the CTP and TTC systems (Figure 2-3). The following sections discuss briefly some of the implementation issues associated with these subsystems.

The LVL1 trigger system presents enormous challenges in terms of the required processing power and data movement capacities. Only a system of custom-electronics processors, with 'hard-wired' algorithms, can perform the necessary calculations within the allowed 2 μs latency, with new data arriving every 25 ns. The very high data rates within the trigger system, for example a total of about 2000 Gbit/s between the calorimeter trigger front-end Preprocessor and Cluster Processor subsystems, require the use of advanced, high-speed links.

The LVL1 system is implemented as a synchronous, pipelined, parallel processor driven by the 40 MHz LHC machine clock; the logic is clocked at 40 MHz or multiples thereof. At the start of the calculation many input channels are processed in parallel (e.g. data from all ~7200 calorimeter trigger towers); at the end the result is a single bit signalling if the event is to be accepted or not. Every clock tick one step in the calculation is performed and the results are

passed on to the next step. While 2 μs corresponds to a total of 80 25-ns clock ticks, an appreciable fraction of the available time is used for signal transmission on cables and optical fibres.

The processor designs make extensive use of application-specific integrated circuits (ASICs). These are employed where large numbers of units, high density and/or ultimate performance are required, justifying the development costs. Elsewhere, use is made of configurable logic (e.g. field-programmable gate arrays) and fast memories. For the on-detector part of the muon-trigger electronics, ASICs have the advantage that it is easier than with commercial off-the-shelf components to ensure that the required radiation tolerance will be achieved.

High-speed data transmission is required in many parts of the system. Within both the calorimeter and muon trigger systems large numbers of high-speed serial links are needed. For long links, such as those between the on-detector and off-detector parts of the muon trigger (length up to 80 m), optical links are used. For short links, such as those between different parts of the calorimeter trigger processor (length ~5 m), cheaper electrical links are used. For inter-module communication within electronics crates, very high-density, high-speed custom backplanes are required. For example, the calorimeter trigger Cluster Processor uses backplanes with >300 signal connections per board, each working at 160 Mbits/s.

A detailed system-level design exists for all parts of the LVL1 trigger; in many areas more detailed designs already exist. Extensive demonstrator prototyping work has been performed over a period of several years, including evaluations of prototypes for the muon, calorimeter and central trigger processors using signals from prototype detectors in test beams. In addition to proving the critical elements of the design, the demonstrator programme has allowed the LVL1 trigger community to acquire the necessary expertise and experience to build the final system.

2.4.1 Calorimeter trigger

The calorimeter trigger system is partitioned into three major parts as shown in Figure 2-3: the front-end Preprocessor, the Cluster Processor, and the Jet/Energy-sum Processor. Each of these subsystems consists of several crates of custom electronics — eight for the front-end Preprocessor, six for the Cluster Processor and five for the Jet/Energy-sum Processor. The full system is rack-based and is located in the underground counting room USA15 that is shielded against radiation.

2.4.1.1 Front-end preprocessor

The Preprocessor receives the analogue signals from the electromagnetic and hadronic calorimeters for ~7200 trigger towers, with a typical granularity of 0.1×0.1 in pseudorapidity–azimuth space. (Analogue summing is performed within the calorimeter front-end electronics to form the trigger tower signals from the finer granularity cell signals.) The input signals are digitized at 40 MHz rate using fast 10-bit analogue-to-digital converters (ADCs), with a least count corresponding to transverse energy $E_T \sim 0.25$ GeV. Digital signal processing is applied to the ADC data to perform bunch-crossing identification (BCID), apply a tower threshold and extract E_T for hit towers; the E_T value is set to zero unless the BCID assigns the pulse to the bunch-crossing under consideration.

The tower threshold, used to reduce the effects of electronic and pile-up noise, together with pedestal subtraction and the final calibration in transverse-energy units, are applied in a lookup table. The 10-bit data from the ADC and BCID logic are reduced to eight bits at the output of the lookup table, giving E_T in units of ~1 GeV. These data are transmitted to the Cluster Processor that searches for isolated high-p_T electrons and photons (isolated electromagnetic clusters), and hadrons and taus (isolated hadronic clusters). As discussed above, the total rate for transmitting these data is ~2000 Gbits/s. The system uses ~2000 serial data links operating in parallel; since the links are short, the signals are sent electrically.

The Preprocessor performs some additional processing for the Jet/Energy-sum Processor. Groups of 2×2 tower 8-bit E_T values are summed, separately for the electromagnetic and hadronic calorimeters. The resulting data are rounded down to nine bits (full-scale $E_T \sim 512$ GeV) before transmission using ~1000 serial data links similar to the ones used for transmission to the Cluster Processor.

2.4.1.2 Cluster Processor

The Cluster Processor acts on input data that can be viewed as an array of 50×64 E_T values in each of the electromagnetic and hadronic calorimeters. The algorithms for both the electron/photon trigger and the hadron/tau trigger search for isolated clusters using information from 4×4 trigger-tower windows (16 towers from each of the electromagnetic and hadronic calorimeters). A search is made for clusters for all possible window positions (so-called overlapping, sliding windows). As a consequence, each tower is used in the calculation for 16 different windows. This has very important implications for the design since data have to be fanned out between processing elements that sometimes are in different ASICs, boards or even crates.

The data fan-out task is addressed at several levels in different ways. Firstly, the mapping of the calorimeter onto the electronic modules of the Cluster Processor has be optimized to simplify as far as possible the sharing of data between modules. Some data have to be shared between crates, in which case the fan-out is done by duplicating the signals at the output of the Preprocessor. Within crates, data have to be shared between at most two electronic modules — this is done on a custom 160-Mbit/s backplane.

A fundamental ingredient to limiting data sharing between modules is to perform a lot of processing on each module. This is achieved using large 9U printed circuit boards with high density input/output and processing logic. High-density input using high-speed serial links is achieved by mounting commercial link receivers in die form on multi-chip modules (MCMs). High-density processing results from implementing in a single ASIC the logic required to process eight windows for both the electron/photon and tau/hadron triggers.

The results from the window processing are combined within the Cluster Processor subsystem to calculate multiplicity values for electrons/photons and tau/hadrons, with eight threshold sets per cluster type. The resulting 16 multiplicity values, rounded down to three bits per value, are passed to the CTP.

2.4.1.3 Jet/Energy-sum Processor

The Jet/Energy-sum Processor receives an array of $\sim 30 \times 32$ 9-bit E_T values in each of the electromagnetic and hadronic calorimeters from the Preprocessor. In contrast to electrons,

photons, and taus and hadrons, a jet is not a well-defined object. Different jet algorithms can be considered, both at the analysis level and at the trigger level. An important consideration is the window size used to search for jets. Here, the optimum choice will depend on many factors — the jet E_T of interest, the luminosity (level of pile-up within the window), and the need to resolve nearby jets in multi-jet events. Since there is no clear best choice, and to provide maximum flexibility to react to conditions at the LHC, the jet-processor design has a programmable window size. A choice can be made (for each set of thresholds) between a window size of 2×2, 3×3 and 4×4 jet elements, where each jet element is 2×2 trigger towers (i.e. each jet element has a size of $\sim 0.2 \times 0.2$ in pseudorapidity–azimuth space. In each case, the windows slide in steps of one jet element in each direction and hence are overlapped.

Many of the comments that were made in the context of the Cluster Processor also apply here. In particular, the fan-out of data is a major issue that is addressed in a similar way in both systems (high-density input/output electronics, optimized mapping of detector onto electronic modules). However, instead of using ASICs for the window processing, here FPGAs are used since the number of chips needed does not justify the cost of developing an ASIC.

The results from the window processing are combined within the Jet/Energy-sum Processor to calculate multiplicity values for jets with eight thresholds. The resulting eight multiplicity values, rounded down to three bits per value, are passed to the CTP.

The other function of the Jet/Energy-sum Processor is to calculate the missing E_T vector and total scalar E_T value for the full event. This is done by summing the E_T values over all of the jet elements and the forward calorimeters. In the case of the missing-E_T calculation, the vector energy components are calculated from the E_T values, using lookup tables to multiply by $\sin(\phi)$ and $\cos(\phi)$. After summation of E_x and E_y separately, a look-up table is used to compute the scalar missing-E_T value.

The missing-E_T value is compared with eight thresholds and the resulting eight-bit comparator pattern is transmitted to the CTP. The same is done for the scalar-E_T value, except that there only four thresholds are provided.

2.4.1.4 Demonstrator programme

The critical aspects of the design of the calorimeter trigger processor have been addressed in a demonstrator programme. This has shown the viability of implementing trigger algorithms in ASICs, and has demonstrated that a high density of high-speed serial links can be achieved by mounting commercial link components on multi-chip modules. The use of custom backplanes using single-ended signals at 160 Mbit/s per link for inter-module communication has also been validated.

Many specific evaluations have been done in the laboratory, for example detailed measurements of error rates for data transmission between modules. In addition, a ‘full-slice’ demonstrator prototype system has also been constructed, containing prototype ADC and front-end modules, cluster-processor modules, and a jet-processor module, as well as the CTP demonstrator module (see below). This has been successfully operated together with ATLAS prototype calorimeters at the test beam.

2.4.2 Muon trigger

As shown in Figure 2-3, the muon trigger is divided into three parts — one part associated with the RPC (barrel) detector, a second part associated with the TGC (endcap) detector and the third part that combines information from the two systems and forms the interface to the CTP.

For both the RPC and TGC based subsystems, part of the electronics is mounted on or close to the detectors. This is necessary for the logic that performs operations that are local to groups of chambers. However, as much as possible of the subsequent logic is placed in the USA15 counting room which is shielded against radiation. The on-detector logic will be radiation tolerant.

Although the algorithms used in the barrel and endcap are conceptually similar, there are significant differences, and different electronics is used in the two cases. Reasons for this include the following:

- The properties and sizes of the RPC and TGC detectors are different. In the case of the RPCs, the spread in time of signals at the output of the discriminator is dominated by the range of propagation delays along the strips, depending on the position of the muon in the chamber. As a consequence, there is a correlation in the timing of hits in different detector layers. In contrast, the timing spread for the TGCs is limited by the response of the detectors.
- The magnetic field map in the endcap is complex due to the interference between the fields generated by the barrel and endcap toroids. This means that there is not a clear separation between the ‘bending’ and ‘non-bending’ projections in the endcap system. In contrast to the barrel, information on bending in the two projections has to be combined before making a classification in p_T.

In both the barrel and the endcap the logic is based on coincidence-matrix ASICs that will be implemented using deep-submicron CMOS technology. These matrices search for patterns of hits in successive detector planes consistent with high-p_T muons originating from the interaction region. In each projection, one matrix implements the low-p_T trigger based on two planes of doublet chambers. A second matrix implements the high-p_T trigger using the output of the low-p_T matrix and an additional plane of chambers.

The coincidence matrices discussed above search for track candidates in projection. Further logic combines the information from the two projections. At this stage, track candidates in space are identified in small regions of the detector with a granularity in pseudorapidity–azimuth space smaller than or about 0.1×0.1, classified into six p_T ranges. These track candidates are collected together over sectors of the detector, retaining the two highest-p_T candidates per sector. In the barrel there are 32 sectors in azimuth per half barrel. Each endcap is divided into two pseudorapidity ranges, $1.05 < |\eta| < 1.92$ and $1.92 < |\eta| < 2.40$, with 48 and 24 sectors respectively.

The information from all the sectors is combined in the muon-trigger interface to the CTP (MUCTPI). This counts the number of muon candidates for each of the six p_T thresholds, and passes the multiplicity information on to the CTP. The MUCTPI is responsible for detecting cases where muons traverse more than one sector due to chamber overlaps, making sure that they are counted only once in the multiplicity calculation. Overlaps within sectors are handled by the logic specific to the barrel and endcap subsystems.

2.4.2.1 RPC-based subsystem

In the RPC-based subsystem, the coincidence matrices and the 'pad' logic that combines information from the two projections within regions, of size $\Delta\eta \times \Delta\phi \sim 0.2 \times 0.2$, are located on the detector. About 800 high-speed optical links are used to transmit the data from the pads to the sector logic that is located in the USA15 counting room.

The coincidence matrices operate with a clock speed of 320 MHz to match the intrinsic time resolution of the RPC detectors, so that the 'gate' width can adjusted to minimize the rate of fake triggers due to accidental coincidences between random hits, e.g. due to radiation in the cavern.

2.4.2.2 TGC-based subsystem

In the TGC-based subsystem, the coincidence matrices are preceded by 'patch panels' that synchronize the hits to the bunch-crossing clock using an adjustable gate width. The subsequent logic is clocked at 40 MHz. The on-detector system includes the low and high-p_T coincidence logic for each projection, but not the so-called r–ϕ coincidence logic that combines the two projections, which is located in the USA15 counting room.

The data from the on-detector logic are encoded before transmission on ~600 high-speed optical links to the r–ϕ coincidence circuits in the sector logic. They include the position of the track candidate within the region covered by the associated on-detector logic, and a measurement of the deflection compared to a straight-line extrapolation from the nominal interaction point. The information from the two projections is combined in order to determine the p_T range of the muon candidate, and the two highest-p_T candidates per sector are retained for transmission to the MUCTPI.

2.4.2.3 Muon–CTP interface

The MUCTPI consists of a single crate of electronics modules located close to the CTP in the USA15 underground counting room. It receives the input data electrically from the sector logic of the RPC- and TGC-based trigger subsystems. Given the small size of the system, it will be implemented using field programmable logic and fast memories. A prototype version of the MUCTPI is currently under design.

2.4.2.4 Demonstrator programme

Demonstrator prototypes have already been tested for both the RPC- and the TGC-based systems. These have proven the design concepts and the performance of the muon trigger. In the case of the RPC logic, demonstrator prototype ASICs have been incorporated in a 'full-slice' test system and operated at the test-beam connected to full-size prototype detectors.

2.4.3 Central trigger processor and TTC system

The CTP is responsible for combining information from the calorimeter and muon trigger systems. As discussed above, it receives the multiplicities of candidate electrons / photons, hadrons / taus, jets and muons, as well as information on the missing and total scalar E_T. It

forms the overall LVL1 trigger decision, accepting events on the basis of up to 96 menu items. Each menu item is a list of criteria that, if satisfied, will result in the event being selected, subject to deadtime and prescale conditions. In the CTP the input data are combined to form decisions for each menu item using lookup tables followed by combinatorial logic. This gives a very high degree of flexibility in specifying the criteria for each menu item.

The CTP takes care of handling the deadtime of the experiment, preventing the overflow of data buffers elsewhere in the experiment by inhibiting triggers. Scalers monitor this deadtime as well as giving information on the rates for each type of trigger. Each menu item can be prescaled by a programmable factor, allowing data to be collected with high-rate triggers concurrently with low-rate ones. In addition to triggers based on information from the calorimeter and muon processors, the CTP accepts external inputs (e.g. for calibration and test triggers) and can generate random triggers and prescaled bunch-crossing triggers.

The CTP, located in the USA15 counting room, will be implemented on one or possibly two 9U printed circuit boards. The use of ASICs is not justified in such a small system. The logic is therefore implemented using a combination of configurable logic (FPGAs and CPLDs) and fast memories. A reduced-size version of the CTP has been constructed as a demonstrator prototype. This has been successfully tested in the laboratory and, together with the calorimeter trigger demonstrator system, at the test beam.

The TTC system is responsible for distributing the LVL1 trigger decision (L1A), together with the 40 MHz clock and other signals, to the detector front-end and readout systems. It uses the system developed in the RD12 collaboration, in which members of the ATLAS LVL1 trigger community are actively participating. This system uses an optical-broadcast network connecting up to about 1000 destinations to each source. The clock, L1A and other signals are encoded and broadcast as a single optical signal to all destinations. The TTC system is separated into zones, with a few zones for each detector system. All of the TTC sources (one source per zone) are located in the USA15 counting room.

Each TTC destination will include an optical / electrical converter followed by a receiver ASIC (TTCrx) that decodes the signal, providing the clock, L1A and other signals electrically. Prototype TTCrx chips have already been produced with performance close to that required for ATLAS, and the next version, implemented in a radiation-hard technology, will be available soon.

2.5 References

2-1 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, December 1994.

2-2 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

2-3 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN / LHCC / 96–41, December 1996.

2-4 *ATLAS Tile Calorimeter Technical Design Report*, CERN / LHCC / 96–42, December 1996.

2-5 *ATLAS Trigger Performance Status Report*, CERN / LHCC / 98–15, June 1998.

2-6 *ATLAS DAQ, EF, LVL2 and DCS Technical Progress Report*, CERN / LHCC / 98–16, June 1998.

2-7 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

2-8 *LVL1 Muon Trigger User Requirements Document (Draft version 1.4)*, ATLAS working document, ATL–DA–ES–0002, March 1998.

2-9 *ATLAS Level-1 Calorimeter Trigger User Requirements Document*, version 1.1.0, ATLAS working document, ATL–DA–ES–0001, April 1998.

2-10 *RD12 Status Report*, CERN / LHCC / 97–29, April 1997.

2-11 *LVL1 Central Trigger Processor (User) Requirements Document (Draft version 1.1)*, ATLAS working document, ATL–DA–ES–0003, February 1998.

2-12 *Timing, Trigger and Control System (TTC) (User) Requirements Document (Draft version 0.2)*, ATLAS working document, ATL–DA–ES–0004, February 1998.

2-13 *ATLAS Level-2 Trigger User Requirements Document*, ATLAS note DAQ–NO–79, November 1997.

3 Calorimeter trigger requirements and environment

The Level-1 Calorimeter Trigger must have high selectivity for several different types of high-E_T objects — electrons and photons, jets, and hadrons from tau decays, as well as missing and total transverse energy. The trigger uses reduced-granularity data from the calorimeters amounting to ~7200 trigger towers, and processing is performed using pipelined digital custom electronics. The range of algorithms is therefore limited, but all algorithms and selection criteria are programmable at the level of parameters. Formal requirements for the trigger are described in the *Level-1 Calorimeter Trigger User Requirements Document* [3-1].

In addition to sending its results to the Central Trigger Processor, the calorimeter trigger must send region-of-interest (RoI) information for all potential trigger objects to the level-2 trigger, in order to limit the amount of data it must transfer to verify these features in more detail. Information from the calorimeter trigger is read out and recorded by the data acquisition system in order to be able to tell what caused the trigger, and to allow monitoring of the performance of the trigger system. A diagram showing all the calorimeter trigger's external interfaces is shown in Figure 3-1.

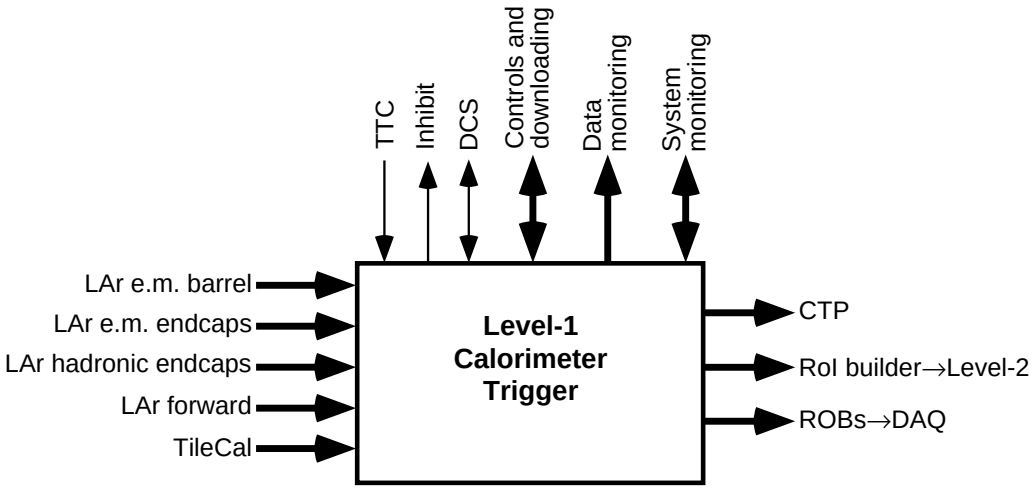

Figure 3-1 Inputs and outputs of the Level-1 Calorimeter Trigger system.

The calorimeter trigger receives one analogue signal per trigger tower from the ATLAS calorimeters. A trigger tower covers, in general, an area of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. For $|\eta| > 2.5$ the tower size changes to larger values, and for $|\eta| > 3.2$ the tower definition becomes more complicated.

The 'building' of tower signals is done by the calorimeter electronics, separately for electromagnetic and hadronic. They are summed over the full depth of each calorimeter. Due care must be taken to ensure similar pulse shapes, relative timing, and calibration of the calorimeter cells that form the trigger tower signals so that the sum represents a sufficiently good, linear energy value in its peak amplitude. It should be noted that the calorimeter trigger will work with values of transverse energy (E_T) only, with a dynamic range of at least 250 GeV. Performance at the low-energy end is limited by noise of typically ~ 0.5 GeV per trigger tower.

Pulses from all ATLAS calorimeters are several bunch-crossings wide. For every LHC bunch-crossing, the calorimeter trigger has to decide whether potentially interesting interactions occurred, and to do this the transverse energy deposited in each trigger tower for that bunch-

crossing must be extracted. To achieve this the trigger carries out a digital process of bunch-crossing identification (BCID) on all trigger towers. It must be stressed that, despite this nomenclature, BCID involves accurate energy measurement as well as assigning every pulse to a unique bunch-crossing.

Trigger tower signals from very large energy deposits saturate before reaching the calorimeter trigger's digitizers. Such signals must also be assigned to the correct bunch-crossing and must *always* produce a trigger and an RoI, since they represent signals above all trigger thresholds and could be due to new and exciting physics. To do this, the shape characteristics of saturated pulses must be accurately known and controlled in order for the BCID logic to function correctly.

In order to give an indication of the performance required of the calorimeter trigger, Table 3-1 lists maximum allowed rates for triggers on specific types of trigger objects with the transverse-energy thresholds shown, and minimal efficiencies, at two values of the luminosity. The maximum rates allowed must include ample safety margins for such things as uncertainties in cross-sections, machine backgrounds, etc.

Table 3-1 Calorimeter trigger rate and efficiency requirements. Energies shown are transverse.

Trigger signal	Level-1 maximum rate (kHz)	Level-1 minimum efficiency
Luminosity = 10^{34} cm^{-2}s^{-1}:		
≥1 isolated e / γ > 30 GeV	22	95%
≥2 isolated e / γ > 20 GeV	5	90%
≥1 jet > 290 GeV	0.2	95%
Large missing-E_T	1	90%
Luminosity = 10^{33} cm^{-2}s^{-1}:		
≥1 isolated e / γ > 20 GeV	8	95%
≥2 isolated e / γ > 15 GeV	2	90%
≥1 jet > 180 GeV	0.2	95%
Large missing-E_T	0.1	90%

The rapidity coverage of the electron / photon and hadron / tau triggers must match that of the ATLAS precision tracking, namely $|\eta| < 2.5$. For the jet trigger the coverage must extend to $|\eta| < 3.2$. Trigger information beyond this is needed only for the missing-E_T trigger.

So as to reduce the amount of data that the level-2 trigger must examine, the level-1 trigger system sends it so-called region-of-interest (RoI) information for each bunch-crossing that produces a trigger at level-1. This tells the level-2 trigger where in the detector to look for features of interest. RoI information must be produced by both the level-1 calorimeter and muon triggers. In the case of the calorimeter trigger, this must be done for all types of trigger object listed above. The Central Trigger Processor also sends information, so that the level-2 trigger knows which kinds of RoI can be objects that caused the level-1 trigger.

The RoIs that are sent to the level-2 trigger for local calorimeter triggers (i.e. electron / photon, jet, hadron / tau) consist of information on the coordinates of the feature, and on which threshold conditions it has passed. The coordinate information is at the granularity of the basic

elements of that type of trigger, and corresponds to the local maximum in transverse energy. The missing-E_T trigger sends its E_x and E_y components, and the total-E_T trigger sends its energy sum.

Data are read out to the data acquisition system from the calorimeter trigger in order to allow a full understanding of the trigger path for all events found to be interesting in the offline data analysis. To achieve this, the E_T values after BCID as well as the results sent to the Central Trigger Processor must be recorded for all triggered events. Combined with a knowledge of trigger parameters, this allows the full trigger logic to be reconstructed offline. For a subsample of events, additional data will be read out in order to monitor and check the correct operation of the trigger. The raw calorimeter data, extending over several bunch-crossings, combined with the BCID results allows verification of the synchronization and the BCID. The correct functioning of the trigger processor itself can be verified by comparing trigger results with the input data and simulating the operation of the trigger processor in software. In case of problems, intermediate trigger data will be read out to allow more detailed verification of its operation

Calibration of the calorimeter input signals must be done using a combination of real data and the calorimeter pulser system. Online checks for tower-to-tower variations can be done using data from the trigger system alone, but full calibration will involve detailed comparison of the readout data from the calorimeters with that for the trigger. The quality of the input data to the trigger must be monitored during normal operation by making online plots of such things as trigger-tower pulse height and timing.

The latency allowed for the calorimeter trigger obviously depends on the cable delays for receiving signals and for sending level-1 trigger results back to subdetector front-end buffers, as well as the latency of other logic, such as the Central Trigger Processor, in the trigger chain. The total level-1 trigger latency, which must be less than 2.0 μs, is in fact dominated by these external factors and so the calorimeter trigger latency must be very short. An estimate is given in Section 8.3 and the latency of the entire level-1 system is summarized in Chapter 18.

In the following five chapters we describe the current situation concerning the level-1 calorimeter trigger. First, in Chapter 4, we describe the simulation work that has led to our choice of trigger algorithms as well as our current understanding of their performance. Then we describe the organization of the signals from the calorimeters, and also the vital problem of bunch-crossing identification and how we intend to solve it, in Chapter 5. Chapter 6 gives a detailed description of the current design for the calorimeter trigger system. In order to give an idea of our experience and the lessons we have learned from the technologies involved, we then have a description of our test-beam and laboratory work over the past six years in Chapter 7. Finally, in Chapter 8 we discuss a number of important issues for the implementation of our proposed design for the trigger.

3.1 References

3-1 *Level-1 Calorimeter Trigger User Requirements Document*, version 1.1.0, ATLAS working document, ATL–DA–ES–0001, April 1998.

4 Calorimeter trigger algorithms and performance

4.1 Introduction

In this chapter we will present the algorithms chosen for triggering, and for delivering regions-of-interest to the level-2 trigger. The simulation results on which the choices were based will be discussed, and some results on the performance of the trigger will be shown. We begin with the electron / photon trigger, and then the hadron / tau trigger and jet triggers are described. Finally, the missing-E_T and total-E_T triggers are discussed together since much of the work was done in common. An accompanying document presents expanded versions of the material in this chapter [4-1].

4.2 Electron/photon trigger

The tasks of the electron / photon trigger are:

- to identify electron and photon candidates using the calorimeter data;
- to classify these according to E_T and isolation;
- to provide multiplicities of candidates passing each classification to the CTP;
- to provide the coordinates of candidates and their classification to the level-2 trigger (electron / photon regions of interest or RoIs).

The main requirements are then:

- good discrimination between isolated electromagnetic (e.m.) showers and QCD jets (the dominant high-E_T process and the main background for the e.m. cluster trigger);
- high efficiency for electrons and photons of $p_T > 10$ GeV, even in events with complex topologies (multiple electrons / photons plus jets);
- accurate location of candidates within the calorimeters, to minimize the size of RoI the level-2 trigger needs to read out;
- unambiguous classification of candidates found, and accurate count of their multiplicity.

The choice of the granularity of input data, and of the algorithm used, are determined by these requirements, and also by technical and financial constraints.

4.2.1 Granularity and algorithm

The inputs to the algorithm are a set of ‘trigger towers’, of granularity 0.1×0.1 in $\Delta\eta \times \Delta\phi$. These are formed by analogue summation of calorimeter cells [4-2]. There are separate sets of trigger towers from the (e.m.) and hadronic calorimeters. The electromagnetic cluster trigger uses data from the region $|\eta| < 2.5$, which is the limit of inner detector coverage and of high-granularity electromagnetic calorimetry. The algorithm is required to be efficient out to the ends of this region.

The trigger-tower granularity was determined by a balance between rejection of jet background and the cost and complexity of the trigger processor. In Figure 4-1 we show the variation in the rate of an inclusive electron trigger with the granularity of the trigger towers. A 40 GeV threshold was applied to a cluster of two trigger towers (adjacent in either η or ϕ), and no isolation requirements were imposed. As can be seen, the rate varies approximately linearly with the trigger tower dimension, i.e. as the square root of the cluster area. Conversely, the number of trigger towers, and hence the cost and complexity of the trigger system, varies quadratically with the tower dimension. In fact, Figure 4-1 overstates the gains to be had from a very fine trigger-tower granularity; for towers much finer than 0.1×0.1 the two-tower cluster used would give an unacceptably soft trigger threshold. This would then require that a larger cluster be used, which would give a higher rate. A trigger-tower granularity of 0.1×0.1 was therefore chosen as giving a good compromise between performance, and the cost and complexity of the trigger. This decision was made before the ATLAS Letter of Intent of 1992 [4-3], and is essentially 'frozen' into the design of the calorimeter granularity and readout.

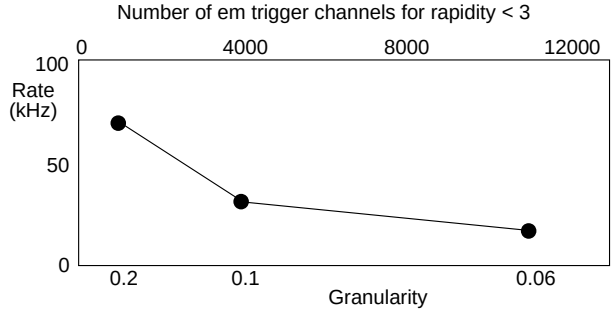

Figure 4-1 Inclusive electron/photon trigger rate vs trigger-tower granularity, for a threshold of 40 GeV and luminosity of 10^{34} $cm^{-2}s^{-1}$. A two-tower electromagnetic cluster was used and no isolation requirement made.

The proposed electron / photon trigger algorithm is illustrated in Figure 4-2. It is based on a window of 4×4 towers in the electromagnetic and hadronic calorimeters, and consists of four elements:

- A 2×2-tower electromagnetic cluster, used to identify candidate RoIs.
- A 2-tower electromagnetic cluster, used to measure the E_T of electromagnetic showers. There are four such clusters within the RoI cluster, and the most energetic of these is used.
- A ring of 12 electromagnetic towers surrounding the clusters, which is used for isolation tests in the electromagnetic calorimeters.
- The 16 hadronic towers behind the electromagnetic clusters and isolation ring, which are used for isolation tests in the hadronic calorimeters.

The window slides in steps of one trigger tower in both the η and ϕ directions, and so there is one window for each tower within the acceptance of the electron / photon trigger. The overlap of the windows has a major impact on the trigger processor design — since the windows overlap, elements of the processor (chips, boards or crates) processing neighbouring windows need to share many of the same inputs, and hence fan-out of trigger tower signals to multiple destinations is required. For the preferred algorithm, any element of the trigger processing $n \times m$ windows needs data from $(n + 3) \times (m + 3) \times 2$ trigger towers (the factor 2 is for e.m. and hadronic calorimeters). Of these, $(n - 3) \times (m - 3) \times 2$ are needed by this processing element only, and the rest are shared with one or more other processing elements. A larger window would result in a greater degree of overlap and signal sharing, and so the preference is to use an

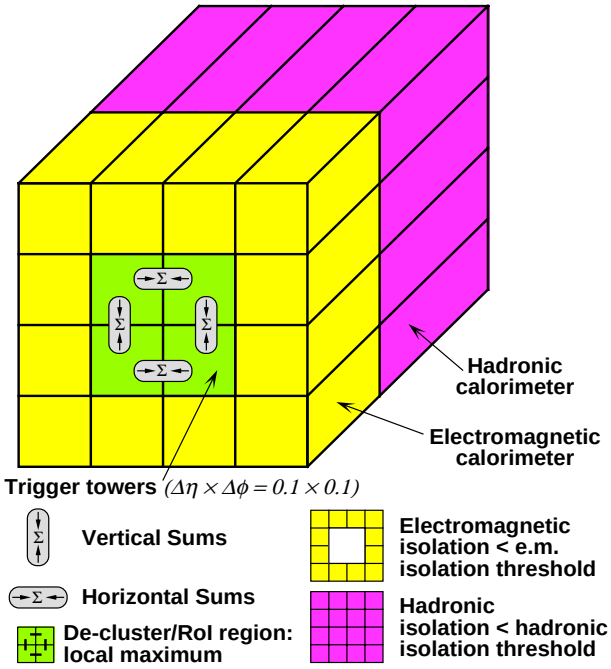

Figure 4-2 Electron/photon algorithm.

algorithm which can be executed in as small a window as possible without compromising performance.

The requirements for a trigger object to be found within the window are:

- the RoI cluster must be a local E_T maximum (see below);
- the most energetic of the four trigger clusters must pass the electromagnetic cluster threshold;
- the total E_T in the electromagnetic isolation region must be less than the e.m. isolation threshold;
- the total E_T in the hadronic isolation region must be less than the hadronic isolation threshold.

If all of these conditions are met, then the window is considered to contain an electron / photon candidate (no distinction can be made here between electrons and photons). Eight sets of trigger E_T thresholds (combinations of cluster, e.m. isolation and hadronic isolation) are foreseen, and the candidate is classified according to which sets it passes.

For each of the eight sets of thresholds, the multiplicity of candidates passing that selection is counted and passed to the CTP, as an input to its decision. Three bits are used to indicate the multiplicity for any selection, and so for each set of thresholds the multiplicity can range between 0 and 7 (multiplicities higher than 7 must be counted as 7). This restriction on the multiplicity passed to the CTP does not affect the number of RoIs which may be indicated to the level-2 trigger.

This algorithm was arrived at after studies of its performance and that of alternatives, which will now be described. In all of these studies, the full GEANT-based simulation of the ATLAS detector was used. For high-luminosity studies, the effects of pile-up were simulated for a

sequence of bunch-crossings, with the contributions from bunch-crossings other than the one of interest weighted according to the pulse shapes expected from the ATLAS calorimeters.

4.2.2 Declustering and regions of interest

An RoI is generated for each electron / photon candidate. The RoI information consists of the location of the RoI, in the form of the η, ϕ indices, plus bits indicating which sets of thresholds it has passed. The RoI coordinate is therefore indicated at the granularity of the trigger towers. Because each set of thresholds tests three variables (electromagnetic cluster E_T, electromagnetic isolation, and hadronic isolation) there is no unambiguous 'ranking' of the different sets, and so the classification word indicates all sets of thresholds passed by the RoI.

What are needed for the CTP and RoI information are, respectively, the multiplicity and location of objects passing the different trigger selections. However, because the algorithm windows (Figure 4-2) overlap, it is insufficient to simply count the number of windows within which the cluster and isolation criteria are met. For example, a high-E_T deposit in a single e.m. trigger tower (with no significant E_T deposits nearby) will pass the cluster and isolation criteria in four overlapping algorithm windows. In order to avoid multiple-counting of electron / photon candidates, and in order to provide an unambiguous coordinate for RoIs, some additional logic is required. This process is often referred to as 'declustering', in that it resolves clusters of hits which naturally arise in overlapping-window algorithms.

A declustering scheme based on sharing of results between neighbouring algorithm windows runs into problems where neighbouring windows are evaluated in different chips, modules, or even crates. One solution is to transfer the results from all algorithm windows to another processor, where they may be compared and ambiguities resolved. This, however, is only practical if a reduced granularity is used for this comparison (as was proposed earlier in [4-2]), since otherwise this 'compare and resolve' logic will be as complex as the one executing the original algorithm. This then limits the ability to separately count nearby candidates, and does not solve the problem of producing an unambiguous, high-precision RoI coordinate. The preferred solution is then to add an extra element to the processing within the algorithm window (in addition to the trigger cluster and isolation tests) which removes this possibility of multiple-counting. The method adopted is to use an algorithm which unambiguously defines an RoI coordinate, and then to count as trigger objects only those RoI candidates which are associated with valid trigger cluster and isolation sums. This same approach is also used for the hadron / tau and jet triggers.

The obvious criterion for finding an unambiguous RoI coordinate is to choose the 4×4 window whose central cluster, used to define the RoI E_T, best contains the shower, i.e. within which the RoI cluster E_T is a maximum when compared with neighbouring RoI clusters. The question then is: what cluster should be used to define the RoI? The preferred solution is to use a cluster of 2×2 electromagnetic trigger towers for this purpose (Figure 4-2). Requiring that such a cluster be an E_T maximum, i.e. selecting the RoI cluster which best contains an electromagnetic shower, should also minimize leakage of that shower into the electromagnetic isolation region around it (Figure 4-3). The 4×4-tower algorithm window is large enough to contain all of the neighbouring 2×2 clusters, and hence we can determine whether the RoI cluster for a given window is an E_T maximum without needing to enlarge the window (which would have adverse consequences for signal fan-out). Since an e.m. shower which is contained entirely in one or two towers will result in neighbouring RoI clusters of exactly equal E_T, the requirement that the RoI cluster must be more energetic than its overlapping neighbours is relaxed slightly

— instead we require that the RoI cluster be more energetic than its neighbours along two connected edges, and at least as energetic as its neighbours along the opposite two edges (Figure 4-4). This ensures that for *any* isolated shower a valid RoI will be found.

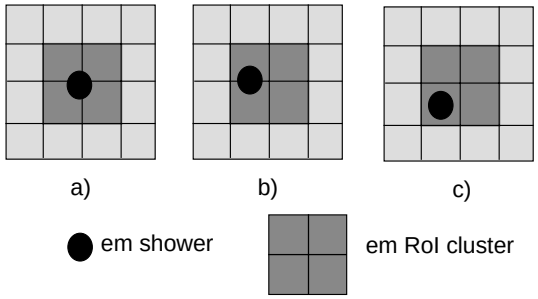

Figure 4-3 Relative position of RoI cluster and shower centroid, (a) for showers shared between four 0.1×0.1 towers, (b) for showers shared between two towers, and (c) for showers contained within a single tower. The types of the inequalities in the 'local maximum' test (Figure 4-4) bias well-contained showers into a particular corner of the RoI cluster.

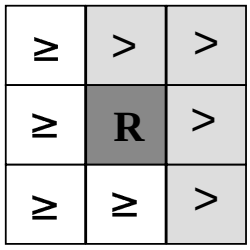

Figure 4-4 'Local E_T maximum' test for an RoI candidate. The 0.2×0.2 RoI cluster 'R' is accepted if it is more energetic than the neighbouring clusters marked '>' and at least as energetic as those marked '≥'. The use of the two types of inequalities in this way ensures that where a shower generates two equal RoI clusters, one and only one of these will always be selected.

Whenever we say that some object is required to be a 'local E_T maximum', it is this type of condition which is meant.

If one algorithm window contains an RoI cluster which is a local E_T maximum, then by definition the neighbouring windows (i.e. those whose trigger clusters potentially overlap) do not, and hence there can be no ambiguity about RoI coordinate. The electron / photon candidate multiplicity is simply the number of algorithm windows passing both RoI and trigger conditions.

An alternative definition of the RoI coordinate would have been to require that a single e.m. trigger tower were a local E_T maximum. This was used in the ATLAS Technical Proposal [4-2] to provide RoI coordinates. This might appear to give better RoI coordinate precision since it is based on a smaller object. However, the coordinate resolution is determined primarily by the size of the step by which the RoI cluster slides, rather than by the size of the RoI cluster, and so the two algorithms in fact give comparable RoI precision, as shown in Figure 4-5. (The offset is not significant.) The single-tower RoI is not suitable for declustering in the preferred algorithm. The reason is that, in contrast to a 2×2 cluster which uniquely identifies a *single* 4×4 algorithm window, a single e.m. tower is part of the trigger clusters in *four* such algorithm windows. Thus, with a single-tower RoI, the trigger logic must evaluate all four algorithm windows, which requires data from a larger environment (5×5), and hence a big increase in the signal fan-out. The single-tower RoI would instead be a natural solution for algorithms based on a 3×3 window, which is disfavoured due to weaker isolation performance.

4.2.3 Choice of trigger cluster size

There are two main requirements on the trigger cluster:

- it should be large enough to contain electromagnetic showers sufficiently well so as to give a sharp trigger threshold;

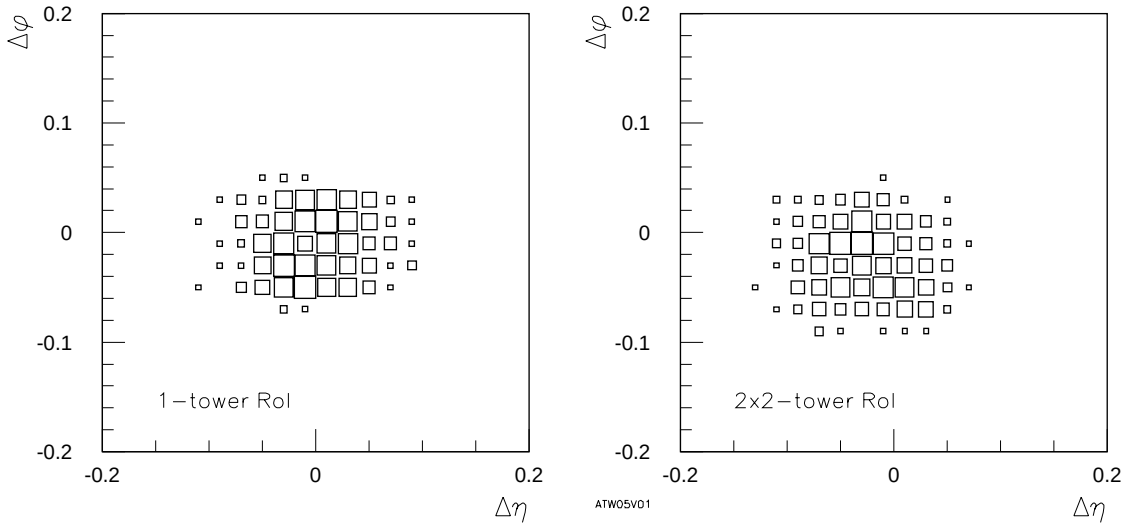

Figure 4-5 Difference ($\Delta\eta$ vs. $\Delta\phi$) between RoI centre and actual electron position for a single-tower and 2×2-tower RoI cluster.

- it should be small enough to provide good discrimination between electrons or photons and QCD jets.

The ultimate criterion is the trigger rate due to QCD jets for a threshold meeting the physics efficiency requirement. Three different trigger clusters have been compared, each of which is completely contained within the 2×2 tower RoI cluster. These are

- 2-tower clusters, where the two towers may be neighbours in either η or ϕ;
- 3-tower clusters, where the three towers form an 'L' shape (with any orientation);
- 4-tower clusters, which are the same as the RoI cluster itself.

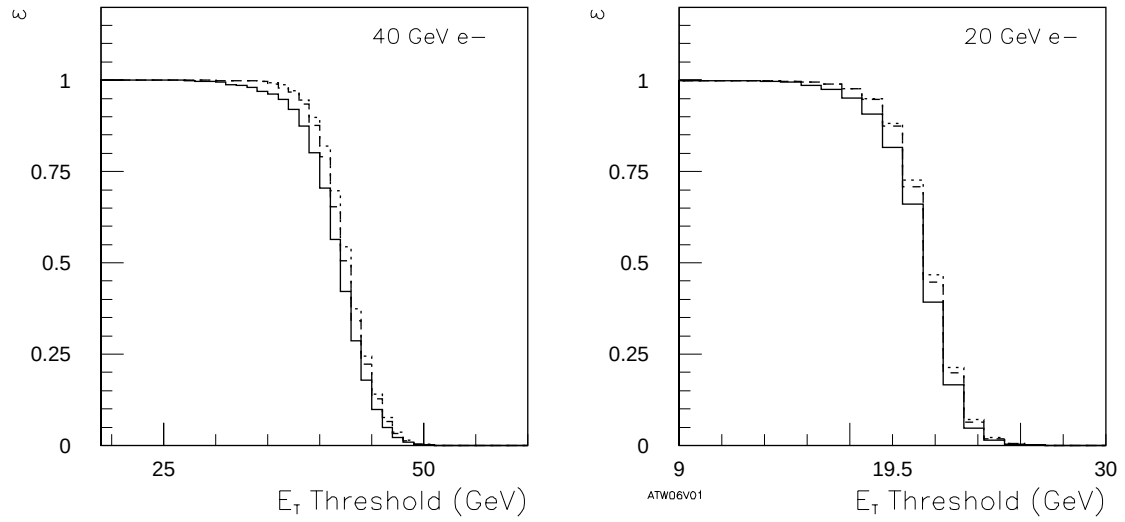

Figure 4-6 Cluster threshold efficiency curves for the three different cluster options, for electron E_T of 40 and 20 GeV. The solid line shows the performance of the two-tower cluster, the dashed line the three-tower cluster, and the dotted line the four-tower cluster.

Figure 4-6 compares the efficiencies of these clusters as a function of threshold for electrons of E_T = 20 and 40 GeV. There is very little difference in threshold sharpness between the three-tower and four-tower clusters, with the two-tower cluster giving a slightly softer threshold. Figure 4-7 compares the thresholds for electrons with bremsstrahlung, and for converted photons, for the different algorithms. It can be seen that for all three algorithms these distributions differ little from those of Figure 4-6 (which included the mixture of bremsstrahlung and non-bremsstrahlung events expected from the distribution of material in the ATLAS detector). Figure 4-8 shows the trigger cluster E_T as a function of rapidity, showing a similar degradation of performance for all algorithms in the calorimeter barrel/endcap transition, due to the extra material in this region.

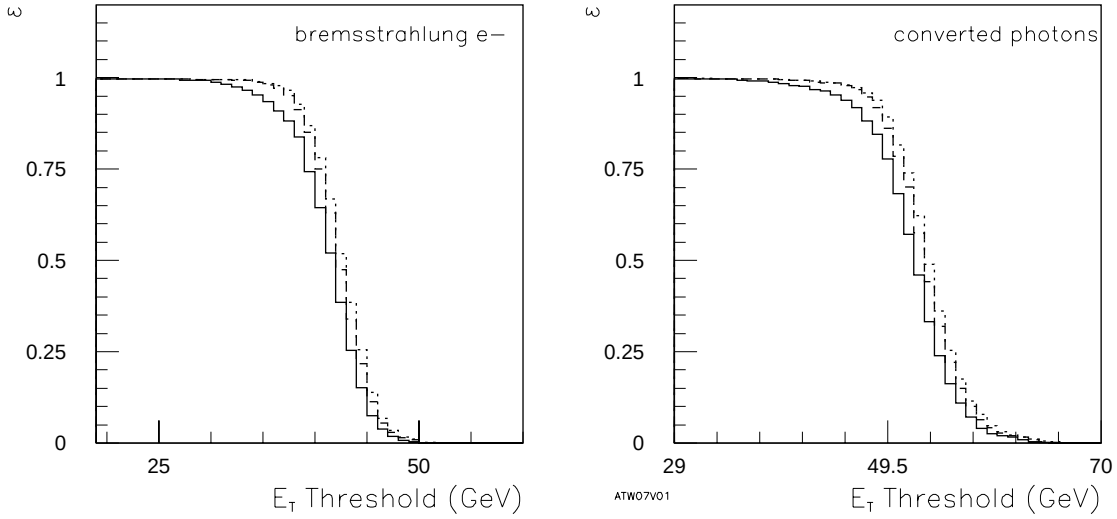

Figure 4-7 Cluster threshold efficiency curves for electrons with bremsstrahlung (40 GeV) and converted photons (50 GeV). The solid line shows the performance of the two-tower cluster, the dashed line the three-tower cluster, and the dotted line the four-tower cluster.

The larger three- and four-tower clusters also contain a greater fraction of jet E_T, resulting in a higher rate for the same threshold. Figure 4-9 shows the rate due to QCD jets vs. the efficiency for 20 GeV electrons for the three algorithms, at a luminosity of 10^{33} cm^{-2}s^{-1}. As can be seen, the smaller two-tower cluster gives a significantly lower rate. The differences are summarized in Table 4-1, requiring 95% efficiency for electrons of E_T = 20 GeV and 40 GeV. Even at low luminosity and relatively high cluster threshold, the smaller cluster would still be favoured.

At the higher luminosity of 10^{34} cm^{-2}s^{-1}, we again show the trigger rate vs. the efficiency, for 30 GeV electrons, in Figure 4-10. The smaller cluster is also favoured at high luminosity, as might be expected from the effect of pile-up. Figure 4-11 compares the threshold sharpness of the algorithms. Here a greater separation of the curves can be seen, due to the additional pile-up E_T, though the effect is again small.

The two-tower trigger cluster is preferred, since it offers superior performance in terms of rate vs. efficiency, and gives a trigger threshold behaviour that is acceptably sharp. The lower cluster-trigger rate also reduces the dependence on isolation, increasing the robustness of the trigger performance.

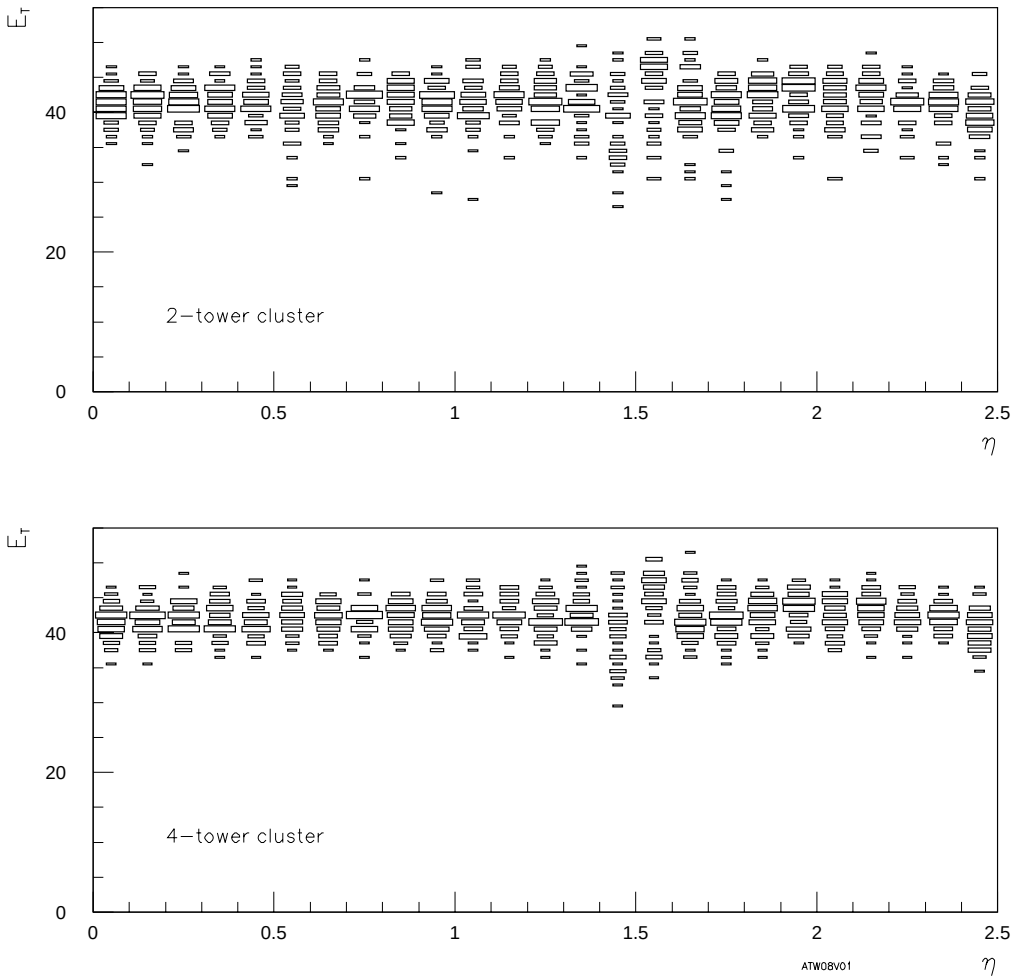

Figure 4-8 Rapidity-variation of cluster E_T for 40 GeV electrons, for two-tower and four-tower clusters. The area of degraded resolution is the transition between the barrel and endcap electromagnetic calorimeters. The calibration was optimized for a two-tower cluster.

Table 4-1 Comparison of thresholds and rates for three cluster algorithms at two values of electron E_T, for luminosity 10^{33} cm^{-2}s^{-1} and requiring 95% efficiency.

	E_T = 20 GeV			E_T = 40 GeV		
Cluster size	**2-tower**	**3-tower**	**4-tower**	**2-tower**	**3-tower**	**4-tower**
Threshold	16 GeV	17 GeV	17 GeV	35 GeV	36 GeV	37 GeV
Rate	24 kHz	29 kHz	31 kHz	2.6 kHz	3.2 kHz	3.7 kHz

4.2.4 Choice of isolation regions

The cluster threshold alone does not provide sufficient jet rejection to meet the ATLAS physics efficiency requirements with a sustainable trigger rate. Thus isolation is used to obtain greater rejection. The sizes and shapes of the isolation regions are restricted by the allowable sizes of the algorithm window itself. With a two-tower electromagnetic trigger cluster the smallest possible algorithm window is 3×3 towers (i.e. the environment needed for declustering by requiring that a single e.m. trigger tower be a 'local E_T maximum' as defined above). Equally, the increase

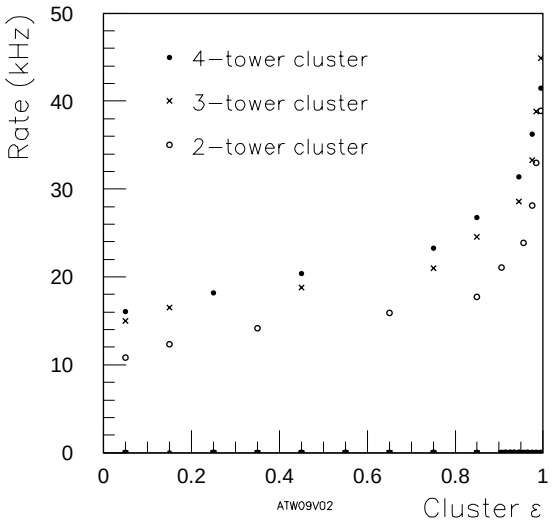

Figure 4-9 QCD background rate vs. efficiency for 20 GeV electrons, for luminosity 10^{33} cm^{-2}s^{-1}.

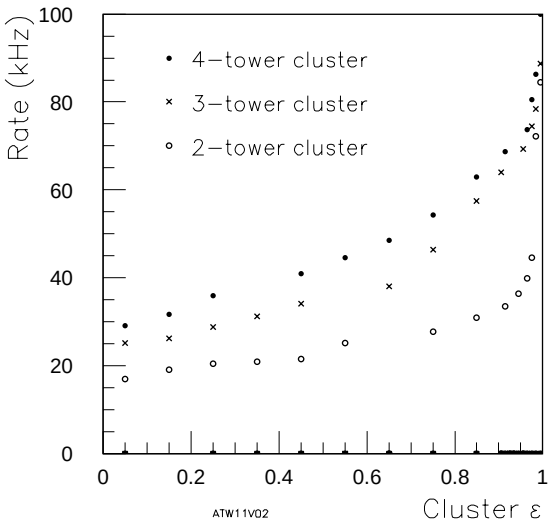

Figure 4-10 QCD background rate vs. efficiency for 40 GeV electrons, for luminosity 10^{34} cm^{-2}s^{-1}.

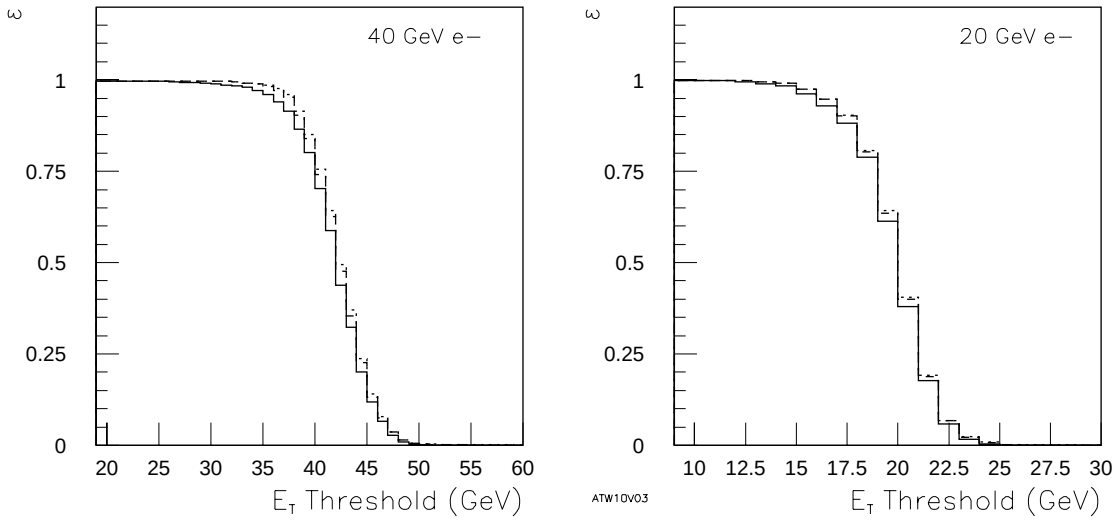

Figure 4-11 Threshold curves for electrons of E_T = 20 and 40 GeV, for luminosity 10^{34} cm^{-2}s^{-1}. The solid line shows the performance of the two-tower cluster, the dashed line the three-tower cluster, and the dotted line the four-tower cluster.

in signal fan-out (and hence cost, complexity and/or technical risk) with environment size makes an algorithm window larger than 4×4 trigger towers problematic. Hence the studies were restricted to these two possible algorithm windows.

Within a 4×4 algorithm window there are 'obvious' choices:

- electromagnetic isolation based on a ring of 12 towers surrounding the RoI and trigger clusters;
- hadronic isolation based on the sum of hadronic E_T in the entire 4×4 window.

Another possibility would be a smaller hadronic isolation region, consisting of the 2×2 hadronic towers directly behind the RoI. This might allow triggering on electrons closer to jets than the larger windows would. This was therefore also studied.

Within a 3×3 algorithm window there is an 'obvious' choice for hadronic isolation, namely the sum of all hadronic E_T in the window. Electromagnetic isolation is more problematic. One possibility would be to use the sum of the 3×3 e.m. towers minus the most energetic two-tower trigger cluster. However, since the two-tower cluster does not completely contain all showers (as can be seen from the threshold curves), such an algorithm would suffer from relatively long tails in the isolation E_T for true electrons and photons, limiting the usefulness of electromagnetic isolation. Instead, the sum of the 3×3 e.m. towers minus the most energetic 2×2 e.m. cluster was used, an algorithm which should be no more vulnerable to e.m. shower leakage than that used for the 4×4 window. Such an algorithm has also been studied by the CMS Collaboration [4-4].

Figure 4-12 shows the E_T distributions in these different isolation regions for electrons and for jets, both passing an appropriate cluster-E_T cut, without pile-up. In Figure 4-13 the same distributions are shown for events with minimum bias added to simulate a luminosity of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. In both figures the effects of the signal preprocessing (bunch-crossing identification and tower E_T thresholds) have been simulated. It can be seen that in both cases a clear separation between electrons and jets can be made, and that the isolation distributions, especially hadronic isolation, are relatively insensitive to pile-up.

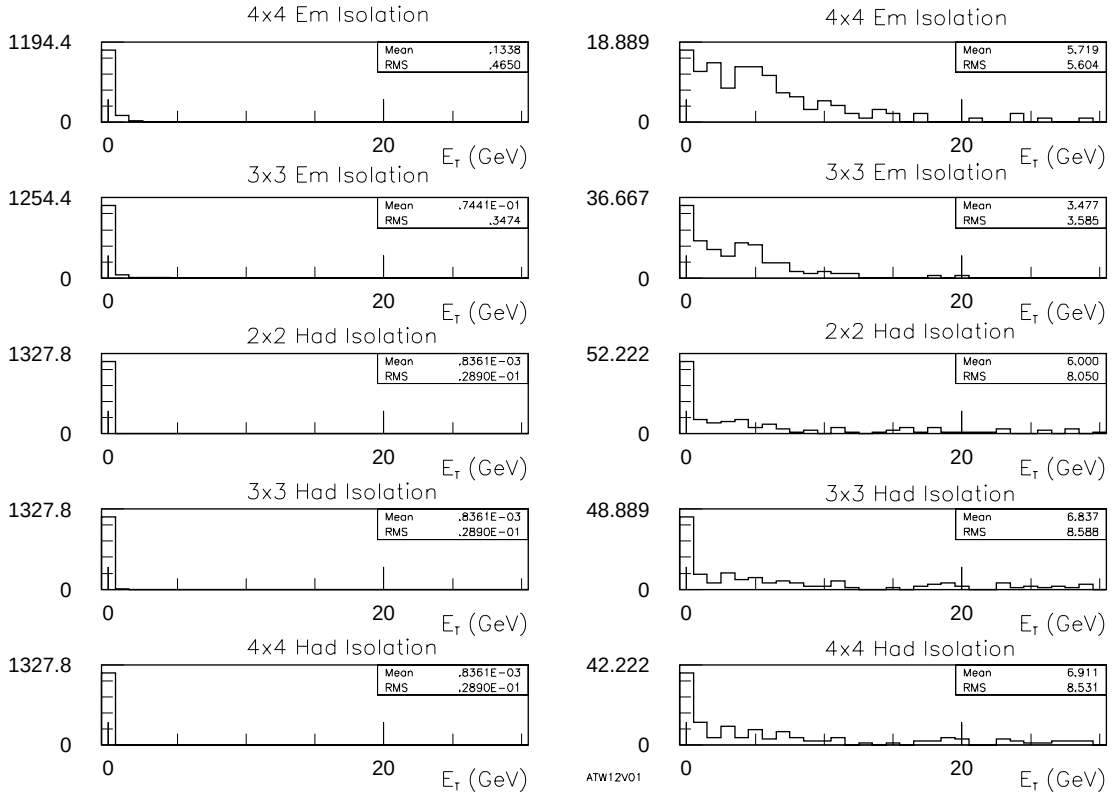

Figure 4-12 Isolation E_T distributions in different regions, for electrons (left) and jets (right). A cluster E_T > 17 GeV was required in each case.

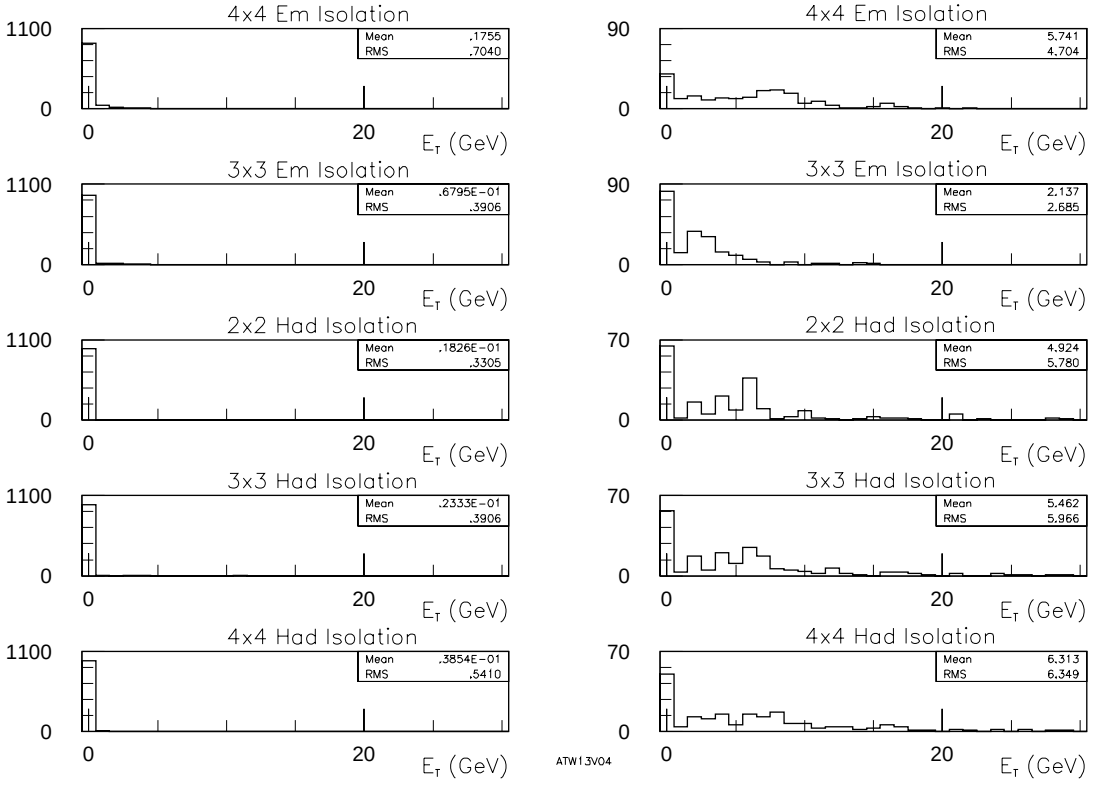

Figure 4-13 Isolation E_T distributions for electrons (left) and jets (right), as Figure 4-12 but with pile-up added to simulate design luminosity.

Table 4-2 summarizes the additional background rejection which may be obtained using isolation at low luminosity. The isolation cuts (which were in the region 1–3 GeV in all cases) were always chosen to be at least 95% efficient for the electron sample. Table 4-3 shows similar results for a luminosity of 10^{34} cm^{-2}s^{-1}. Electromagnetic and hadronic isolation are combined by applying separate thresholds to both, rather than by adding them and applying a threshold to the sum. This is done for reasons of flexibility, and also because of the different sensitivities of the electromagnetic and hadronic isolation sums to pile-up and to cluster E_T (see below).

Table 4-2 Additional rejection factors obtainable against jets passing three different cluster-E_T thresholds for the different isolation algorithms at low luminosity. Isolation is chosen such that electron trigger efficiency is at least 95% in each case.

Isolation algorithm	E_T > 16 GeV	E_T > 27 GeV	E_T > 37 GeV
E.M. 3 × 3 window	1.8	2.4	3.5
E.M. 4 × 4 window	2.7	4.1	4.8
Hadronic 2 × 2 window	1.9	2.4	3.3
Hadronic 3 × 3 window	2.1	2.6	3.7
Hadronic 4 × 4 window	2.3	3.0	4.2
E.M. 3 × 3 + hadronic 3 × 3	3.1	4.5	6.6
E.M. 4 × 4 + hadronic 2 × 2	3.9	5.4	7.3
E.M. 4 × 4 + hadronic 4 × 4	4.3	6.0	8.1

Table 4-3 Additional rejection factors obtainable against jets passing three different cluster-E_T thresholds for the different isolation algorithms at a luminosity of 10^{34} $cm^{-2}s^{-1}$. Isolation is chosen such that electron trigger efficiency is at least 95% in each case.

Isolation algorithm	E_T > 16 GeV	E_T > 27 GeV	E_T > 37 GeV
E.M. 3×3 window	1.7	2.2	2.6
E.M. 4×4 window	2.2	3.7	3.6
Hadronic 2×2 window	1.9	3.1	3.4
Hadronic 3×3 window	2.1	3.4	4.5
Hadronic 4×4 window	2.3	3.8	4.5
E.M. 3×3 + hadronic 3×3	3.0	5.5	7.8
E.M. 4×4 + hadronic 2×2	3.0	6.0	7.8
E.M. 4×4 + hadronic 4×4	3.5	6.4	8.6

From these data we can see that:

- larger isolation regions give better jet rejection;
- there are gains to be made from using both electromagnetic and hadronic isolation in combination;
- the rejection obtainable from isolation increases with increasing cluster E_T.

In particular, the greater rejection obtainable in a 4×4 environment if only one or other of electromagnetic or hadronic isolation is used is advantageous, since this provides extra flexibility should there be problems with one of the calorimeters, or offers the possibility of relaxing isolation (e.g. by removing the electromagnetic isolation requirement) at a lower cluster E_T.

In the above discussion, the isolation requirement was imposed by requiring that isolation E_T be less than a fixed (though programmable) threshold. Another possibility would be to require that the isolation E_T be less than some fraction of the cluster E_T. This would be more complex to implement (requiring a division or multiplication as well as a threshold), but would be a reasonable approach if the isolation E_T for genuine electrons and photons were dominated by shower leakage into the isolation sums. The e.m. isolation distributions for electrons of different E_T are shown in Figure 4-14. As would be expected, the hadronic isolation sums are quite insensitive to electron E_T. There is some dependence of electromagnetic isolation E_T on electron E_T, but this is a relatively small effect for electrons of $E_T < 50$ GeV. Other noise sources also contribute to the electromagnetic isolation sums, and so there is not a simple scaling with E_T. These plots are for low luminosity. At high luminosity, pile-up makes additional contributions to these sums and so the shower leakage is a smaller proportion of the total isolation E_T. It is therefore felt that applying fixed thresholds (which may be different in different ranges of trigger cluster E_T) is a more appropriate approach than requiring isolation to be less than some fraction of cluster E_T.

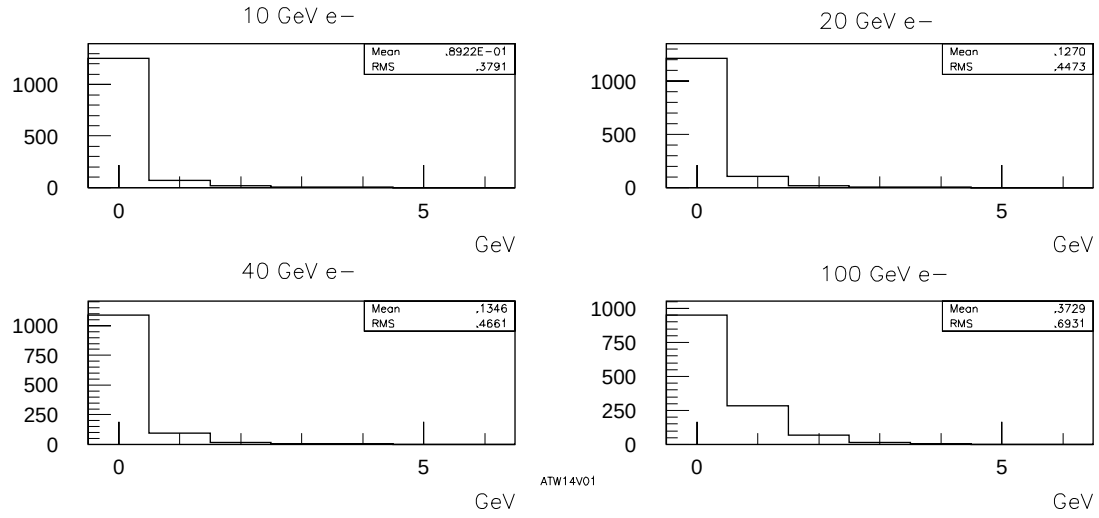

Figure 4-14 Electromagnetic isolation E_T distributions for electrons of different E_T.

4.2.5 Overall trigger performance

It is foreseen that electron / photon candidates may contribute to the level-1 trigger decision in three ways:

- as inclusive triggers, where at least one signal above a given threshold is sufficient to cause an event to be accepted;
- in electron / photon multiplicity triggers, e.g. dielectron / diphoton triggers;
- in combination with other trigger inputs, e.g. electron and missing-E_T or electron and muon.

Figures 4-15 and 4-16 show the estimated inclusive trigger rates as a function of electron p_T (for 95% electron efficiency), for luminosities of 10^{33} cm^{-2}s^{-1} and 10^{34} cm^{-2}s^{-1}. Each plot shows the rate without isolation, using only hadronic isolation, and using both electromagnetic and hadronic isolation. Figures 4-17 and 4-18 show similar distributions for a dielectron / diphoton trigger. In these, the isolation cuts were chosen to give 95% efficiency for triggering on the pair, rather than a single electron or photon. As can be seen, by requiring both electromagnetic and hadronic isolation the trigger efficiency and rate requirements can be comfortably met at both low and high luminosity.

As shown above, there is a dependence of isolation E_T on electron / photon E_T, and so one would not require the same isolation thresholds irrespective of cluster E_T. Also, since the trigger rate falls quite rapidly with increasing cluster E_T, there is no need to require equally stringent isolation for higher-E_T clusters, as the effect on trigger rate is negligible. It is therefore anticipated that the isolation conditions will be progressively loosened with increasing cluster E_T. An example of how this could work is shown in Table 4-4 for low luminosity, and Table 4-5 for high luminosity.

Level-2 trigger selections may use information about lower-E_T electrons or photons than those which contribute to the level-1 decision, as low as E_T = 10 GeV even at high luminosity. Thus a low-E_T ‘secondary RoI’ selection is needed in order to produce RoI flags for such objects. In order to maximize efficiency, these are assumed to be non-isolated. The number of such RoIs in

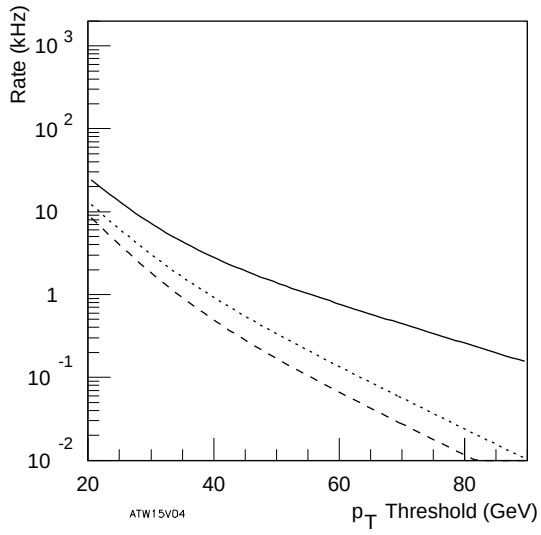

Figure 4-15 Inclusive electron trigger rate for luminosity 10^{33} $cm^{-2}s^{-1}$, without isolation (solid), requiring only hadronic isolation (dotted) and requiring both electromagnetic and hadronic isolation (dashed).

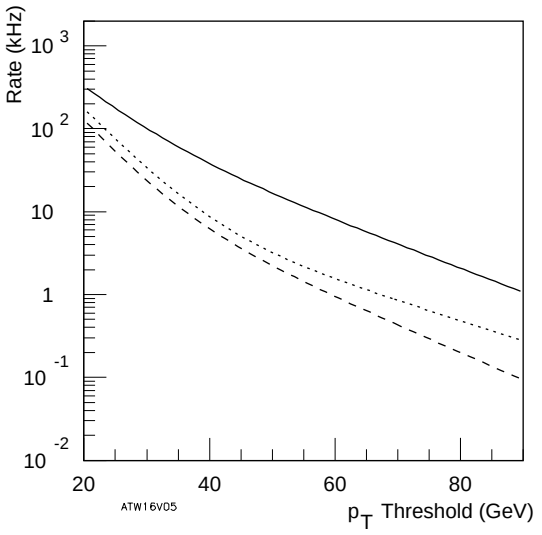

Figure 4-16 Inclusive electron trigger rate for luminosity 10^{34} $cm^{-2}s^{-1}$, without isolation (solid), requiring only hadronic isolation (dotted) and requiring both electromagnetic and hadronic isolation (dashed).

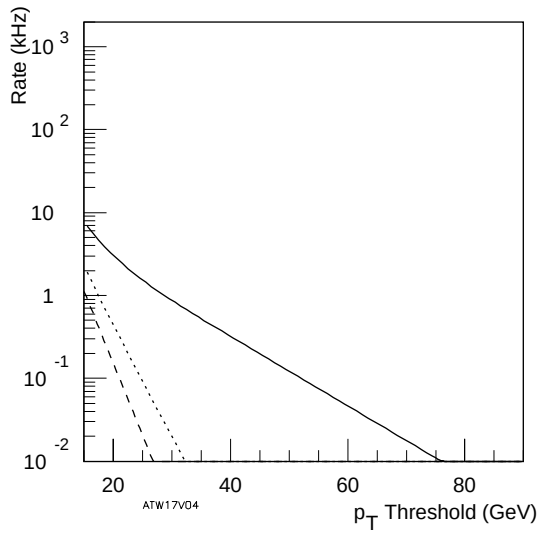

Figure 4-17 Electron/photon pair trigger rate for luminosity 10^{33} $cm^{-2}s^{-1}$, without isolation (solid), requiring only hadronic isolation (dotted) and requiring both electromagnetic and hadronic isolation (dashed).

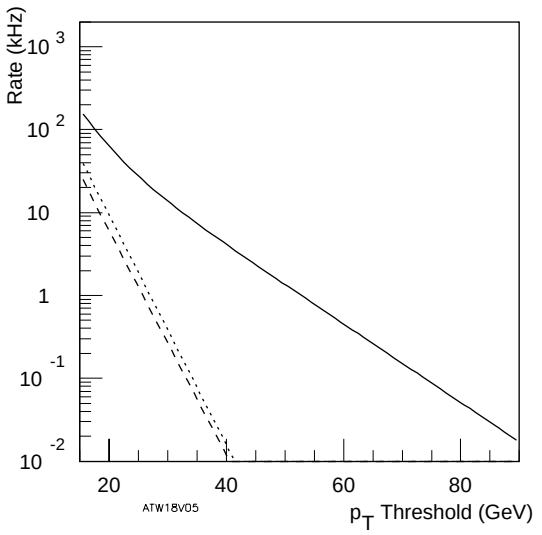

Figure 4-18 Electron/photon pair trigger rate for luminosity 10^{34} $cm^{-2}s^{-1}$, without isolation (solid), requiring only hadronic isolation (dotted) and requiring both electromagnetic and hadronic isolation (dashed).

a typical triggered event will have a major impact on the data transfer needed for the level-2 trigger. The average RoI multiplicities in events passing the trigger selections of Table 4-4 and Table 4-5 are shown in Figure 4-19. As can be seen, the average RoI multiplicities are reasonably low. If needed, one could also require isolation for such secondary RoIs, reducing the average multiplicity further.

Table 4-4 An example of how isolation criteria might be progressively loosened with increasing E_T for luminosity 10^{33} cm^{-2}s^{-1}. The total rate is less than the sum of the parts due to overlaps between the different selections.

Trigger selection	Threshold	Isolation	Rate
≥ 1 electron / photon	E_T > 17 GeV	E.M. + hadronic	11 kHz
≥1 electron / photon	E_T > 35 GeV	Hadronic	1.2 kHz
≥1 electron / photon	E_T > 60 GeV	None	0.6 kHz
≥ 2 electron / photons	E_T > 12 GeV	E.M. + hadronic	1.4 kHz
≥ 2 electron / photons	E_T > 20 GeV	Hadronic	0.1 kHz
≥ 2 electron / photons	E_T > 35 GeV	None	0.3 kHz
Total trigger rate			**13 kHz**

Table 4-5 An example of how isolation criteria might be progressively loosened with increasing E_T for luminosity 10^{34} cm^{-2}s^{-1}. The total rate is less than the sum of the parts due to overlaps between the different selections.

Trigger selection	Threshold	Isolation	Rate
≥ 1 electron / photon	E_T > 26 GeV	E.M. + hadronic	21.5 kHz
≥1 electron / photon	E_T > 45 GeV	Hadronic	2.6 kHz
≥1 electron / photon	E_T > 75 GeV	None	3.0 kHz
≥2 electron / photons	E_T > 15 GeV	E.M. + hadronic	5.2 kHz
≥2 electron / photons	E_T > 25 GeV	Hadronic	0.4 kHz
≥2 electron / photons	E_T > 45 GeV	None	1.5 kHz
Total trigger rate			**29.2 kHz**

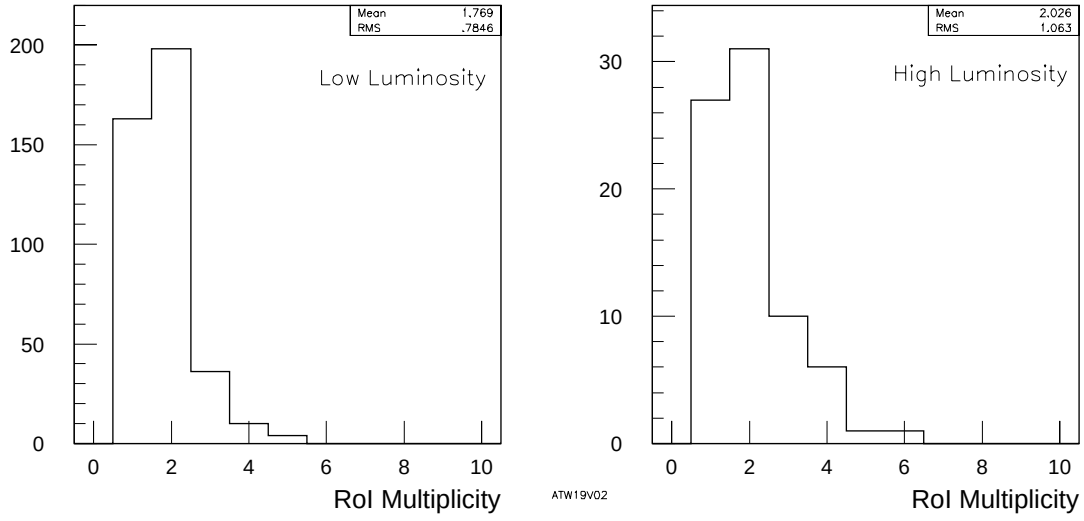

Figure 4-19 RoI multiplicity distributions in triggered events, at luminosity 10^{33} cm^{-2}s^{-1} and 10^{34} cm^{-2}s^{-1}. RoIs were required to contain a cluster with E_T > 7 GeV, and no isolation requirement was made.

4.2.6 Dynamic range required

The electron / photon trigger does not place particularly stringent requirements on the dynamic range of the trigger tower signals, which is proposed to be ~250 GeV in steps of ~1 GeV. Since even at design luminosity the rate for non-isolated clusters of $E_T > 80$ GeV is a negligible fraction of the total trigger rate, there is no need for the range to extend above 100 GeV for the electromagnetic trigger tower signals. With hadronic signals used only for isolation purposes, a very small range would be adequate for these. However, since the same trigger tower signals are used for the hadron / tau trigger, in practice a rather larger range is required. The electron / photon trigger is sensitive to the resolution of the trigger tower signals, primarily in the setting of isolation thresholds. Since the r.m.s. noise in the e.m. trigger towers is expected to be at least 500 MeV, a least-count in the region of 500–1000 MeV would be appropriate here.

4.3 Hadron/tau trigger

4.3.1 Motivation

There could be several purposes for a hadron/tau trigger, implemented in the level-1 calorimeter trigger:

1. In coincidence with a muon trigger or an electron trigger, it could improve the efficiency for triggering on $Z \rightarrow \tau^+ \tau^-$ decays or on low-mass $A \rightarrow \tau^+ \tau^-$ decays.
2. In coincidence with missing-E_T, it could provide a trigger on $W \rightarrow \tau \nu$ and $Z \rightarrow \tau^+ \tau^-$ hadronic decays (at low luminosity).
3. It could select high-E_T hadron decays for calibration of the hadron calorimeter, through the measured momentum from the tracking system.

Item 3 requires the trigger to be used stand-alone, where it would be prescaled and/or used at low luminosities, in order not to saturate the trigger rate. An isolation requirement in the calorimeter would be desirable for items 1 and 2, since the tau decays will give rise to narrow energy depositions. Isolation might be undesirable for item 3, since it would bias the shower shape.

While this trigger might be considered a relatively low priority for physics, it can be implemented relatively easily and cheaply at level-1, using the same inputs and much of the same logic used for the electron/photon trigger.

4.3.2 Algorithms considered

Several algorithms have been considered, subject to the constraint that the minimum granularity in η and ϕ be 0.1. All algorithms have a core energy which is a sum of energy in the electromagnetic and hadronic layers, covering a certain number of 0.1×0.1 trigger towers that is not necessarily the same in the two layers. The algorithms then require this core energy to exceed some threshold. Five different core definitions were considered (see Figure 4-20):

1. A 2×1 or 1×2 trigger-tower group (two neighbouring towers at the same η or ϕ) in the electromagnetic layer, together with the same towers (i.e. the towers directly behind) in the hadronic layer.
2. A 2×1 or 1×2 trigger-tower group in the electromagnetic layer together with a 2×2 tower group in the hadronic layer.
3. A 2×2 trigger-tower group in both electromagnetic and hadronic layers.
4. A 2×2 trigger-tower group in the electromagnetic layer together with a 4×4 tower group in the hadronic layer (centred on the same point).
5. A 4×4 trigger-tower group in both electromagnetic and hadronic layers.

In all these algorithms, the central 2×2 trigger-tower block, summing over electromagnetic and hadronic layers, is required to contain more energy than any of the other eight possible 2×2 tower blocks that can be defined in the same 4×4 window. In practice, the block is required to contain more energy than its neighbours along two edges (top and right), and at least as much energy as its neighbours along the other two edges (they can be equal since the energy is

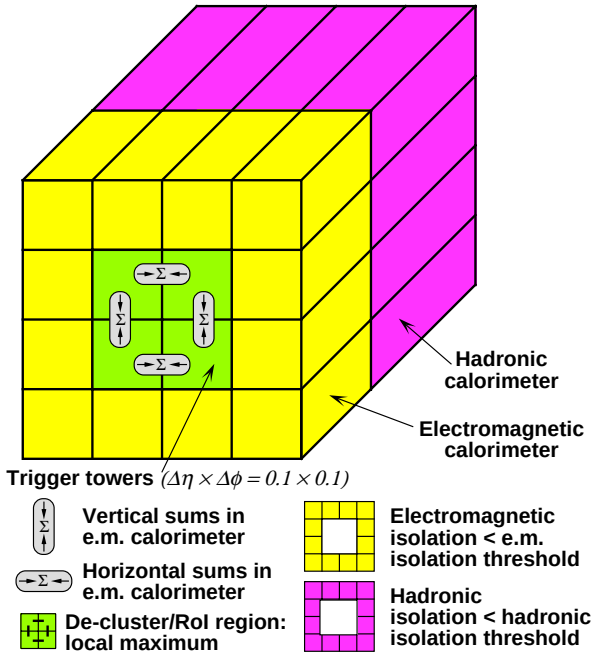

Figure 4-20 Diagram of the trigger towers used in the hadron/tau trigger, core definition 2.

digital). This 2×2 block (which slides by 0.1 in both directions) gives the RoI coordinate for a trigger. This condition ensures that towers are not used twice — i.e. a high energy hadron will give a single RoI. Note that for the 2×1 core definitions, the trigger condition is satisfied if any of the four possibilities (2×1 or 1×2) in the 2×2 block passes the cluster-energy threshold. An alternative 'declustering' algorithm was studied (see Section 4.2.2) but gave very similar results. The declustering scheme described above has already been chosen for the electron / photon trigger, and so is the natural choice for the hadron / tau trigger.

For the isolation definition, the most obvious choice is to use the 12 trigger towers surrounding the 2×2 core, as shown in Figure 4-20, summing the towers in the two layers separately. An alternative is to use a 14-tower isolation definition (in the case of a 2×1 core), including all towers in the 4×4 block except the 2×1 core. Finally, in the case of a 4×4 core definition, isolation was defined in an 8×8 window. This could be the same algorithm as the first, but using an element size of 0.2×0.2 rather than 0.1×0.1. If this latter algorithm were favoured, then the trigger might be more naturally implemented in the Jet / Energy-sum Processor.

4.3.3 Evaluation of candidate algorithms

The evaluation of the different algorithms was performed using two types of data-sets: signal events from $A \rightarrow \tau^+ \tau^-$ decays with m_A in the range 150 GeV to 450 GeV, and background events containing QCD jets. These events were run through the full GEANT simulation.

For the tau sample, only events where a tau was found close to the highest-E_T particle were selected. In order to evaluate the efficiency as a function of E_T, the summed E_T of the hadronic daughters of the tau was used rather than the E_T of the tau itself. One-prong tau decays were compared to three-prong decays, but no difference was observed. For the rest of the analysis, no distinction was made between the different hadronic tau decay modes.

4.3.3.1 Core algorithm

In order to make a fair comparison of the efficiencies for the tau selection, E_T thresholds were defined for the different core algorithms such that all selected similar numbers of jet events, i.e. would lead to the same trigger rate. Two sets of thresholds were defined, and are shown in Table 4-6.

Table 4-6 Thresholds (in GeV) and numbers of jet and tau events selected by the different core algorithms. Also indicated is the ratio of tau to jet events selected.

Algorithm	Threshold	No. jets	N_{tau}	tau/jets	Threshold	No. jets	N_{tau}	tau/jets
2×1	40 GeV	1829	1810	0.99 ± 0.03	60 GeV	588	1235	2.10 ± 0.11
$2 \times 1 + 2 \times 2$	43 GeV	1834	1802	0.98 ± 0.03	65 GeV	579	1213	2.09 ± 0.11
2×2	49 GeV	1788	1715	0.96 ± 0.03	72 GeV	594	1158	1.95 ± 0.10
$2 \times 2 + 4 \times 4$	52 GeV	1865	1700	0.91 ± 0.03	78 GeV	585	1104	1.89 ± 0.10
4×4	62 GeV	1804	1505	0.83 ± 0.03	88 GeV	576	1022	1.77 ± 0.09

The efficiency curves for the tau events vs. the tau hadronic E_T are shown for these algorithms and thresholds in Figures 4-21 for the lower thresholds in the table, and 4-22 for the higher thresholds. The tau closest to the highest-E_T particle in the event was used.

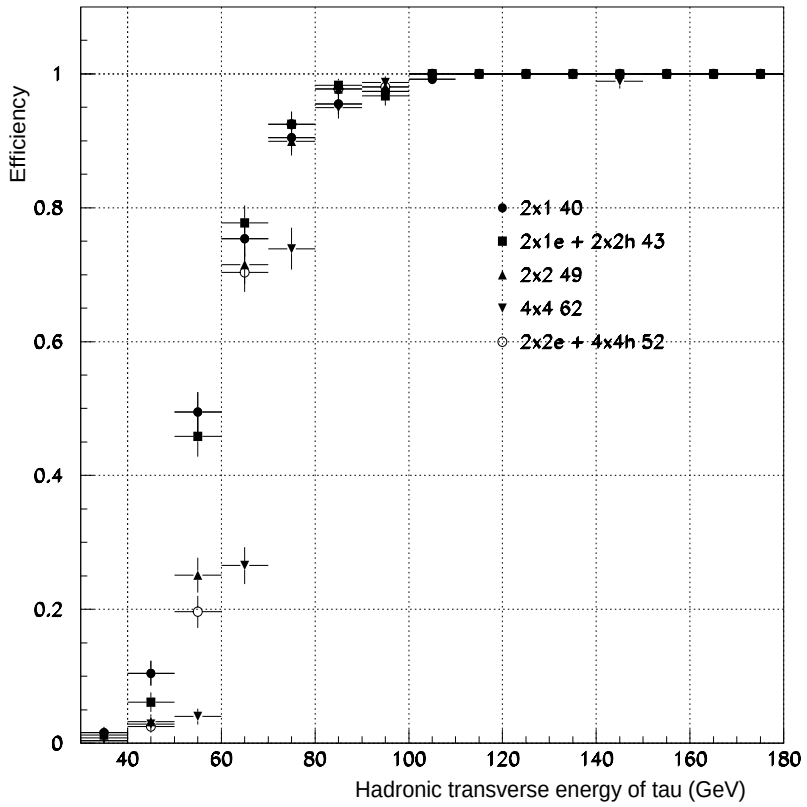

Figure 4-21 Efficiency vs. tau hadronic E_T in GeV for the five different core algorithms with the thresholds indicated.

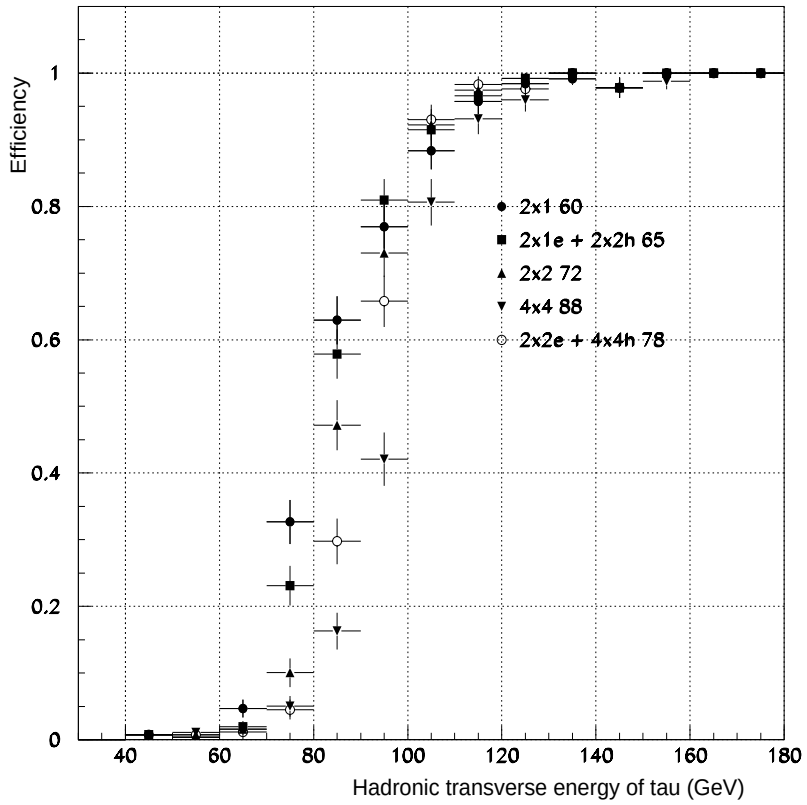

Figure 4-22 Efficiency vs. tau hadronic E_T in GeV for the five different core algorithms with the thresholds indicated.

Since the jet rates were equalized, the algorithm yielding the highest efficiencies should be the best one. From these results, algorithms 1 and 2 are judged to be superior to the others, with algorithm 1 giving higher efficiency in the low-E_T range and algorithm 2 giving higher efficiency in the high-E_T range. This sharper turn-on would make algorithm 2 preferable for calibration purposes, where it would be desirable to restrict to a region of high trigger efficiency to reduce trigger bias. (N.B. For the tau trigger in coincidence with a lepton trigger, an improved efficiency at low energy might perhaps be useful, since the level-2 trigger may be able to accept such events, by reducing the lepton background and applying more sophisticated isolation criteria.)

The total efficiency for the tau events depends on the E_T spectrum. For the E_T spectrum corresponding to the data samples analysed here, the acceptance for the tau events is compared in Table 4-6, including the ratio of tau events relative to jet events accepted. This ratio is a measure of the algorithm performance. At face value the errors look too large to draw firm conclusions, but there are large correlations between the numbers of events accepted by the different algorithms. For example, 98% (93%) of the tau (jet) events accepted by algorithm 1 (threshold = 40 GeV) are also accepted by algorithm 2 (threshold 43 GeV). This means that, when comparing the performance of the different algorithms, the errors should be reduced by a factor of about five. One may then conclude that, for this distribution of taus, algorithms 1 and 2 are superior to the others for both sets of thresholds, but the difference between them is not significant.

One may also study the ability to trigger on hadrons in an inclusive way. For this purpose, the jet events were used, but only events where the highest-E_T particle was hadronic were retained.

The same thresholds were used as above, in order to equalize the overall trigger rate. The efficiency spectra as a function of the hadron's E_T (using the highest-E_T particle in the event) are shown in Figure 4-23. The differences between the various algorithms are not very significant, but algorithms 1 and 2 still have relatively good performance.

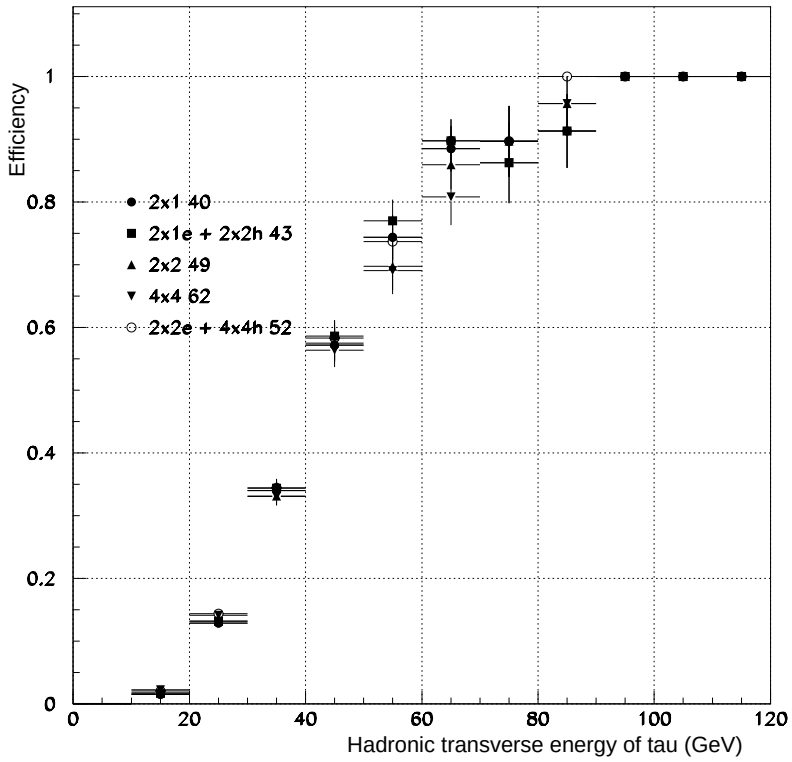

Figure 4-23 Efficiency vs. hadron E_T (GeV) for the five different core algorithms, with the thresholds indicated.

4.3.3.2 Isolation

The isolation in the tau events, summed over electromagnetic and hadronic layers, is compared to that in the jet events in Figure 4-24 for the first three core algorithms, for events passing the lower set of thresholds defined above. In each case, the default 12-tower isolation definition was used. A clear distinction is seen between the tau and jet events, but this distinction does not seem to depend on the core algorithm in a significant way. Also shown in Figure 4-24 is the result using a 14-tower isolation definition, which can only be used for the 2×1 core definition. The distinction between the tau and jet events is reduced with this definition compared to the 12-tower case.

Restricting ourselves to core algorithm 2 (2×1 e.m. + 2×2 hadronic), the isolation in tau events is compared in Figure 4-25 to that in jet events separately for the electromagnetic and hadronic layers. The electromagnetic isolation is seen to be much more powerful than the hadronic, and therefore should be treated separately. The correlation between the isolation E_T values seen in the two layers is shown in Figure 4-26 for the two classes of events. From this plot, it is not clear that the hadronic isolation is a very powerful tool.

The other isolation option studied, which would apply only for the 4×4 core (algorithm 5), is to define it in an 8×8 window. (Effectively, it is a 12-tower isolation with 0.2×0.2 towers.) Using the lower core threshold defined above, the isolation defined in this way is compared between

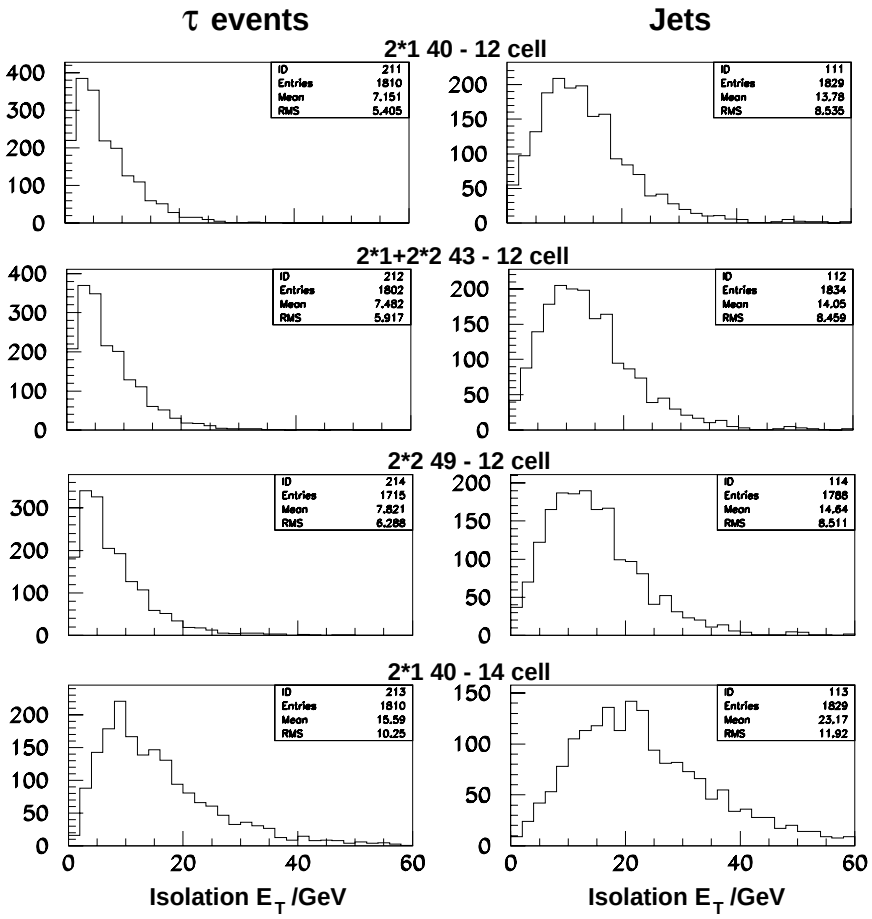

Figure 4-24 Isolation E_T for tau and jet events, summed over electromagnetic and hadronic layers, for core algorithms 1, 2 and 3. The bottom two plots show the same thing for algorithm 1, but using a 14-tower isolation definition.

the two data samples in the two layers in Figure 4-27. Compared to the default isolation definition, the electromagnetic separation is worse, and the hadronic is slightly better. Thus, use of the 4×4 core with this isolation definition would result in an inferior performance relative to algorithm 2 with the standard definition. (In fact, the 8×8 window isolation is largely uncorrelated with the default isolation, so combining the two would improve the performance, but the implementation of such an algorithm would be impractical at level-1.)

4.3.3.3 Absolute trigger rate

Selecting core algorithm 2, Figure 4-28 shows the absolute trigger rate that would result from using the hadron / tau trigger in a stand-alone way as a function of core-E_T threshold, assuming a luminosity of $10^{33}\,cm^{-2}s^{-1}$. The effect of pile-up has been neglected. The figure shows the rate with and without an electromagnetic isolation cut, where two possibilities are indicated for the dependence of such an isolation cut on the core threshold. The first possibility is no dependence — i.e. a fixed cut, while the second possibility is a direct proportionality with the core threshold. The optimal choice most probably lies somewhere in between these extremes.

The expected trigger rates for a few specific core thresholds are also indicated in Table 4-7.

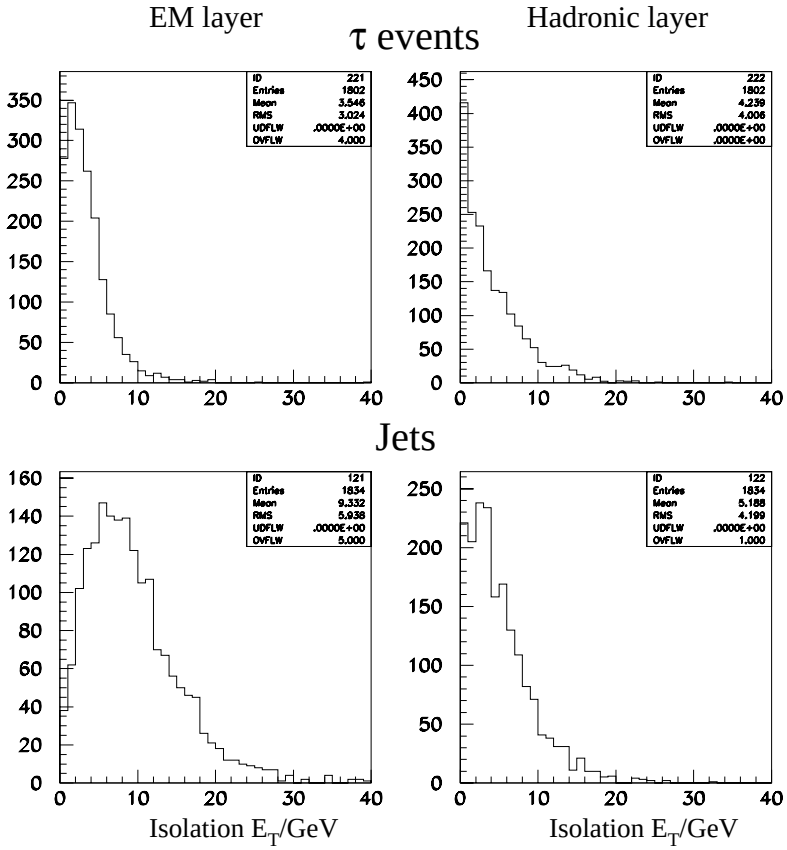

Figure 4-25 The isolation E_T (12-tower), shown separately for the electromagnetic and hadronic layers for tau and jet events, using core algorithm 2.

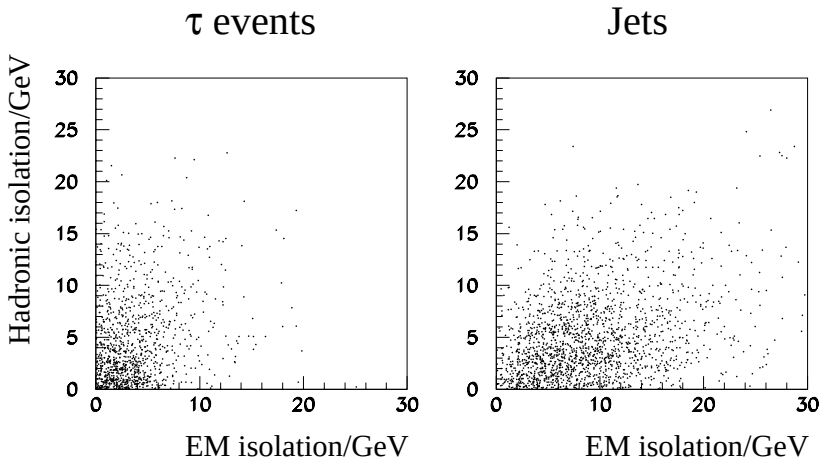

Figure 4-26 The hadronic isolation E_T vs. the electromagnetic isolation E_T (12-tower) for tau and jet events, using core algorithm 2.

4.3.4 Conclusion

Several possibilities have been studied for the hadron/tau calorimeter trigger at level-1. From studies of the tau efficiency curves at a fixed rate of jet events, two algorithms are preferred for

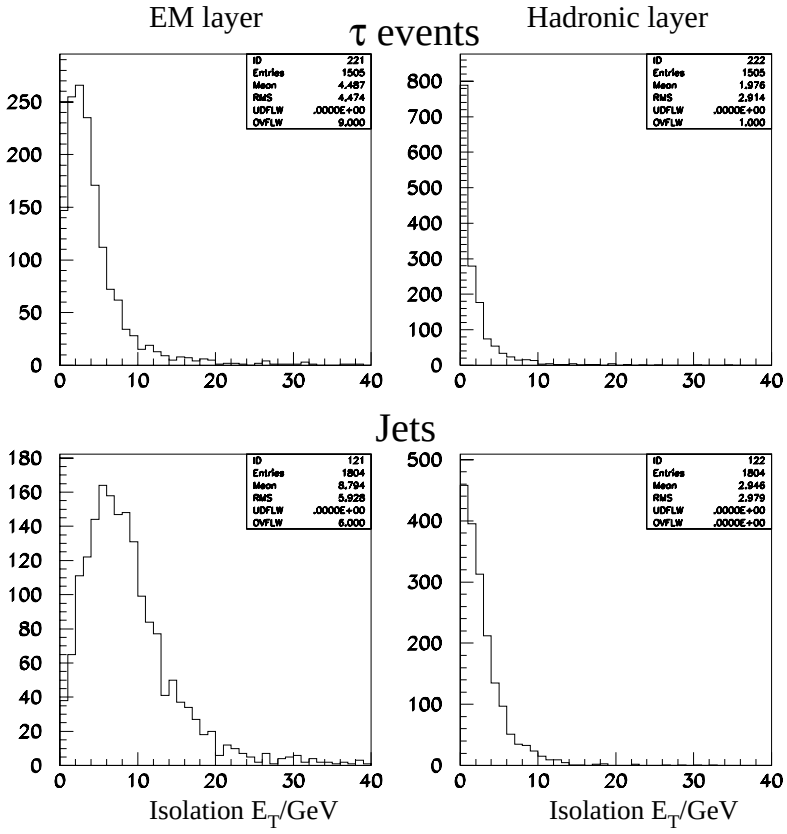

Figure 4-27 The isolation E_T defined in the 8×8 window, using a 4×4 core, separately for the electromagnetic and hadronic layers, for the tau and jet events.

Table 4-7 Trigger rates for a stand-alone tau trigger for a few different core-E_T thresholds and the isolation requirements indicated.

Core-E_T threshold	E.M. isolation threshold	Rate at 10^{33} cm^{-2}s^{-1}
20 GeV	7 GeV	16.3 ± 0.4 kHz
40 GeV	10 GeV	2.1 ± 0.2 kHz
60 GeV	12 GeV	0.6 ± 0.1 kHz

the core definition: algorithms 1 (2×1) and 2 (2×1 e.m. + 2×2 hadronic). Between these two, algorithm 2 has the faster turn-on in efficiency and so is preferred.

An isolation requirement may also be made, and the most suitable definition was found to be that which sums the E_T in the electromagnetic layer in the 12 trigger towers surrounding the 2×2 core (which acts as a declustering object). The corresponding isolation sum in the hadronic layer may also be used, but its discrimination power is not very large.

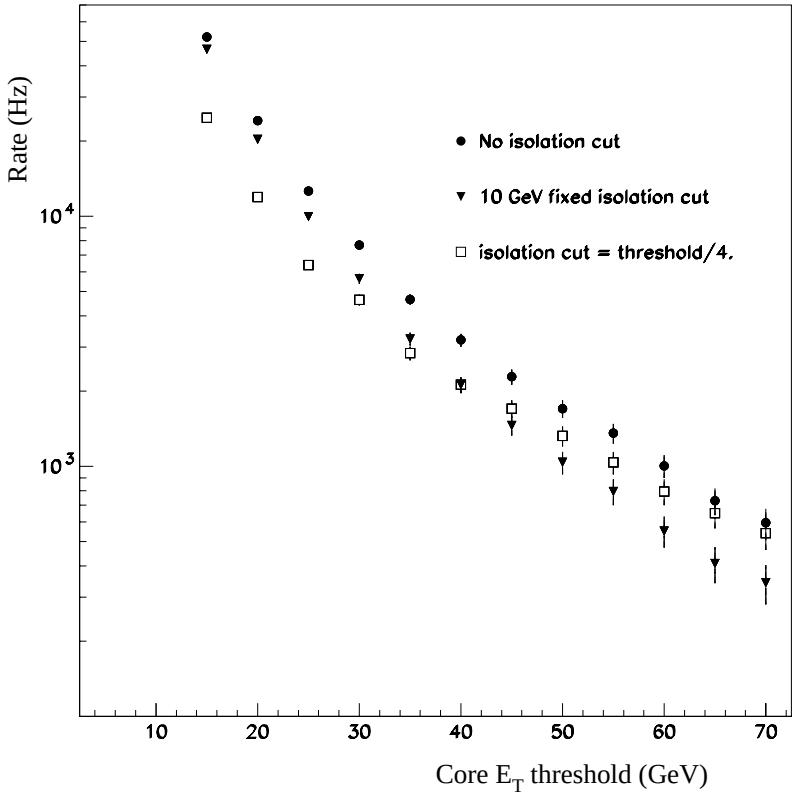

Figure 4-28 Trigger rate vs core-E_T threshold for an inclusive tau trigger, assuming a luminosity of 10^{33} $cm^{-2}s^{-1}$ without pile-up. The effect of using an electromagnetic isolation requirement is indicated.

4.4 Jet trigger

Jet production is expected to be the dominant hard process at the LHC. Unlike the electron / photon and hadron / tau triggers, the main requirement on the jet trigger is therefore not that it should discriminate between two different types of object, but rather that it should discriminate on the basis of the E_T and multiplicity of jets. Only when trying to flag the lowest-E_T jets (20–40 GeV) as secondary RoIs for level-2 is the question of background from other sources (noise and pile-up) expected to be a problem.

The tasks of the jet trigger are:

- to identify hadronic jets using calorimeter data;
- to classify these according to E_T;
- to provide multiplicities of jets passing each classification to the CTP;
- to provide the coordinates of candidates and their classification to the level-2 trigger (jet RoIs).

In order to achieve the physics goals, the jet trigger should:

- have as good an energy resolution as possible for both high-E_T (100–200 GeV) jets and low-E_T (< 50 GeV) jets, at both low and high luminosity;
- provide an accurate count of jet multiplicity above the required thresholds, even in complex multijet events;
- provide accurate coordinates for jets within the calorimeters, to minimize the size of the RoIs which level-2 needs to read out.

The choice of the granularity of input data, and the algorithms used, are determined by these requirements, and also by technical and financial constraints.

For most of the studies described below, a ‘fast’ simulation was used. This included a realistic model of the electronics effects, preprocessing, and trigger algorithms, but lacked the detailed simulation of the detector and material. This model allowed large, high-E_T datasets to be produced easily. Cross-checks were performed using the full GEANT-based simulation, particularly in the delicate area of low-E_T jet performance.

4.4.1 Granularity and algorithm

The proposed jet trigger algorithm is illustrated in Figure 4-29. It is based on a window of 4×4 ‘jet elements’, which have a granularity of 0.2×0.2 in $\Delta\eta \times \Delta\phi$ and are summed in depth between the electromagnetic and hadronic calorimeters. The algorithm has two components:

- A 2×2-element cluster (0.4×0.4), used to identify candidate jet RoIs.
- A trigger cluster, used to measure the jet E_T. This cluster can be 2×2, 3×3 or 4×4 jet elements (0.4×0.4, 0.6×0.6 or 0.8×0.8 in $\Delta\eta \times \Delta\phi$).

The window slides in steps of one element in both the η and ϕ directions, and so there is one window for each jet element within the acceptance of the jet trigger, $|\eta| < 3.2$, which is the limit of the endcap calorimeter coverage. As in the electron / photon and hadron / tau algorithms, the overlap of the windows requires fan-out of the input data between different processing units of

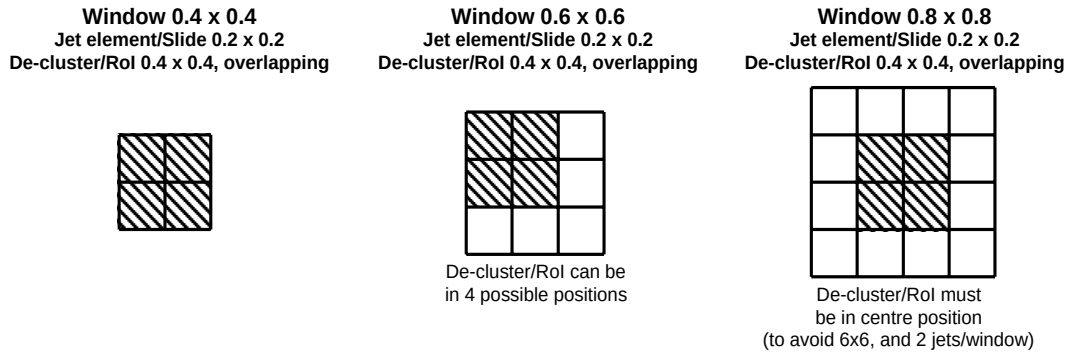

Figure 4-29 Jet trigger algorithms based on a step of 0.2×0.2. In the 0.6×0.6 algorithm, the RoI (shaded) may occupy any corner of the cluster.

the jet trigger system. The 4×4 algorithm window used here is the same as that used for the electromagnetic and hadronic triggers (except that there is now only one sample in depth), and so results in a similar level of signal fan-out within the jet system.

The two requirements for a trigger object to be found within the window are:

- the RoI cluster is a local E_T maximum;
- the trigger cluster E_T passes the jet cluster threshold.

The optimum size of the jet cluster will depend on both the jet E_T and the luminosity. The resolution for high-E_T jets at low luminosity will be dominated by the containment of the jet E_T within the cluster, favouring a larger cluster. Conversely, when flagging low-E_T jets, especially at high luminosity, the amount of electronic and pile-up noise within the jet cone will be the limiting factor in jet trigger performance. For this reason, a flexible system is foreseen, in which different jet cluster sizes may be used simultaneously at different E_T thresholds, allowing optimization of different jet selections for different purposes. The choice of which cluster size can be used with which threshold will be programmable.

Because of the overlap of algorithm windows there is, just as for the electron / photon and hadron / tau triggers, the possibility of multiple-counting of overlapping clusters due to a single jet. This is especially true for the 0.8×0.8 cluster, where the degree of overlap is particularly large. This problem is resolved in the same way as for the other local-object triggers: rather than count trigger clusters above threshold, we instead count RoI clusters which are identified as local E_T maxima (Section 4.2) and which are part of a trigger cluster passing one or more trigger thresholds. The local E_T maximum condition ensures that the RoI clusters are distinct (non-overlapping) objects, and so eliminates the problem of multiple counting.

An RoI is generated for each jet candidate passing the above conditions. The RoI information consists of the location of the RoI, in the form of the η, ϕ indices of the window, plus bits indicating which jet cluster thresholds it has passed. The RoI coordinate is therefore the centre of the RoI cluster (which is the centre of the algorithm window). Because it is possible for different trigger cluster sizes to be used for different trigger thresholds, the ‘threshold classification’ word for each RoI should indicate the complete set of thresholds passed by the RoI, rather than merely the most energetic.

As with the electron / photon and hadron / tau triggers, eight jet trigger thresholds are foreseen. The multiplicity of candidates passing each threshold is counted and passed to the CTP as a

three-bit word. Again, the restriction of the CTP data to three-bit multiplicities does not affect the number of RoIs which may be indicated to level-2.

4.4.2 Choice of algorithm options

The trigger architecture essentially forces the jet algorithm to be based upon a fixed-size sliding window. There are three basic elements to the actual algorithms studied:

- the size of the trigger cluster;
- the size of the step by which it slides;
- the algorithm used for declustering and RoI generation.

There is a strong coupling between the step size, which is equal to the size of the jet elements, and the trigger and RoI cluster options, since both must be an integer multiple of the jet element size. We compare below algorithms based on elements of 0.2×0.2 and 0.4×0.4. The latter are illustrated in Figure 4-30. For a jet element of 0.4×0.4, a trigger cluster of 0.8×0.8 was used, i.e. a 2×2 elements. Two decluster/RoI algorithms were considered. In one, the entire 0.8×0.8 cluster was required to be a local E_T maximum. In the other, one of the four jet elements making up the cluster was required to be a local E_T maximum. In the latter case, the RoI element could occupy any corner of the cluster.

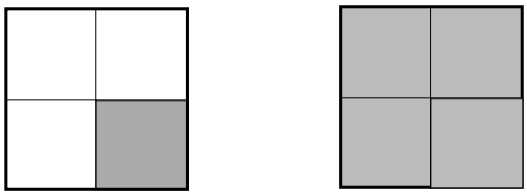

Figure 4-30 Alternative jet trigger algorithms based on a step of 0.4×0.4. In the left-hand figure, the RoI is a single jet element, which may occupy any corner of the cluster. In the right-hand figure the cluster itself is also the RoI, and so is required to be a local E_T maximum.

Only algorithms based on overlapping clusters are considered here. Non-overlapping clusters could also be used, and are in principle simpler to implement as they would not require fan-out of input data. However, such algorithms have a fundamental problem: jets falling near the edge of such a cluster have their energy split between two or more clusters. This has three adverse consequences:

- It causes tails in the jet resolution curve, and hence a softened jet efficiency threshold. This results in degraded inclusive jet trigger performance.
- Single high-E_T jets may be reconstructed as two or more lower-E_T clusters. This degrades the ability to classify events according to jet multiplicity.
- Since a single jet may be reconstructed as two or more adjacent clusters, additional processing is required to provide reliable RoI coordinates for level-2. This then eliminates the advantage of simplicity of implementation of such triggers.

For these reasons, it is felt that such algorithms do not meet the requirements of a jet trigger for ATLAS, and so they are not considered further here.

4.4.3 Inclusive trigger performance

One problem with studies of jet triggers is there is no unique definition of what constitutes a ‘jet’. Hence one must use a particular jet-finder as a ‘reference’ against which the trigger algorithms are compared. Fixed-cone algorithms are widely used within ATLAS, but relying on one of these as a reference carries the risk that it would bias studies of the optimum cluster size. For this reason both a fixed cone algorithm (with $\Delta R = 0.4$) and a k_T algorithm [4-5] were used as references for comparison with the trigger algorithms. The plots shown are mostly for the k_T algorithm, but the results did not significantly depend on which reference algorithm was used.

Figure 4-31 shows the threshold efficiency curves for 100 GeV E_T jets for different cluster sizes, at a luminosity of 10^{33} cm^{-2}s^{-1}. Such jets are of interest for the inclusive jet trigger at this luminosity. As can be seen, the threshold sharpness for jets of 0.6×0.6 and 0.8×0.8 is very similar, while the smaller 0.4×0.4 cluster produces a much softer threshold. Figure 4-32 shows the correlation between efficiency for these jets and the inclusive trigger rate for the same algorithms. From this we see that the larger clusters produce a lower rate when high efficiency is required. The same quantities are shown in Figures 4-33 and 4-34 for 200 GeV E_T jets at the design luminosity of 10^{34} cm^{-2}s^{-1}. Again, a larger cluster size is favoured.

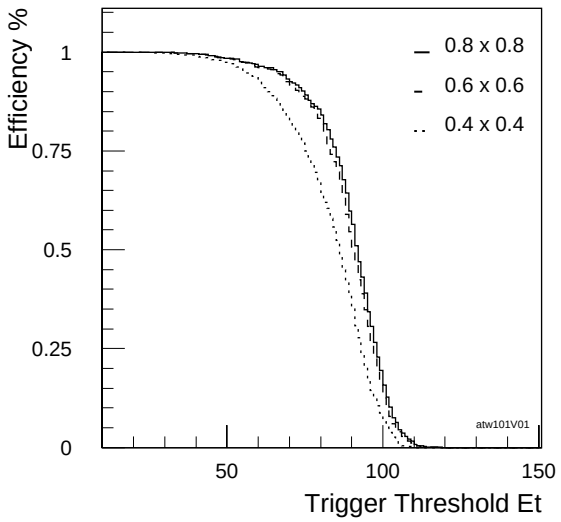

Figure 4-31 Jet trigger efficiency curves for 100 GeV E_T jets, for different cluster sizes, at luminosity 10^{33} cm^{-2}s^{-1}.

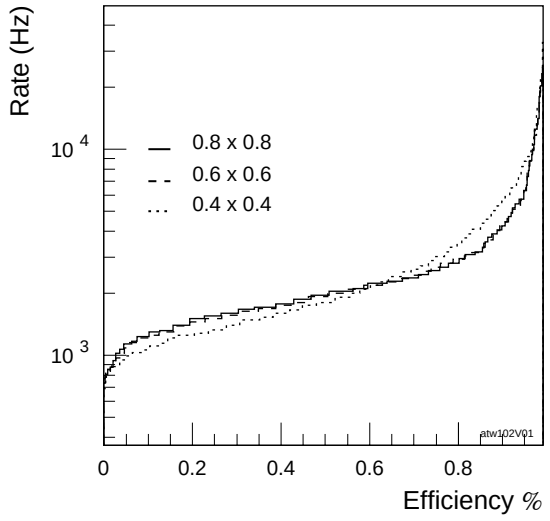

Figure 4-32 Trigger rate vs. efficiency for 100 GeV E_T jets, for different cluster sizes, at luminosity 10^{33} cm^{-2}s^{-1}.

Table 4-8 compares the inclusive jet rates for different cluster and step sizes, for thresholds chosen to be 95% efficient for jets of 100 GeV and 200 GeV E_T at luminosities of 10^{33} cm^{-2}s^{-1} and 10^{34} cm^{-2}s^{-1}, respectively. As can be seen, there is a strong preference for a larger jet cluster size for such trigger selections. There is also a small gain from using a step of 0.4×0.4. This is understood to be due to the requirement that the RoI cluster for the algorithm based on a step of 0.2×0.2 be at the centre of the 0.8×0.8 jet cluster. Relaxing this requirement would however require an algorithm window of 6×6 jet elements, and an unacceptable increase in signal fan-out.

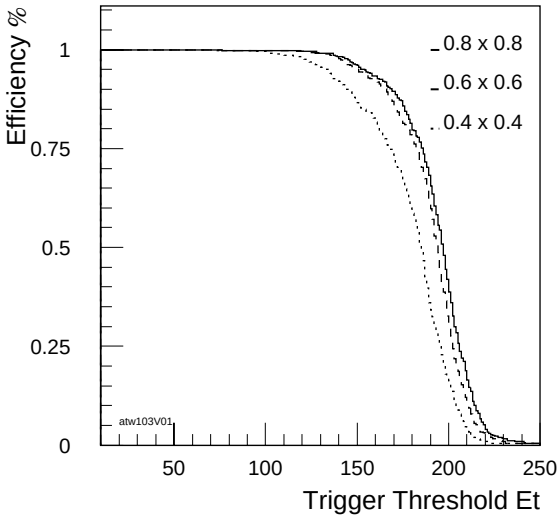

Figure 4-33 Trigger efficiency curves for 200 GeV E_T jets, for different cluster sizes, at luminosity 10^{34} cm^{-2}s^{-1}.

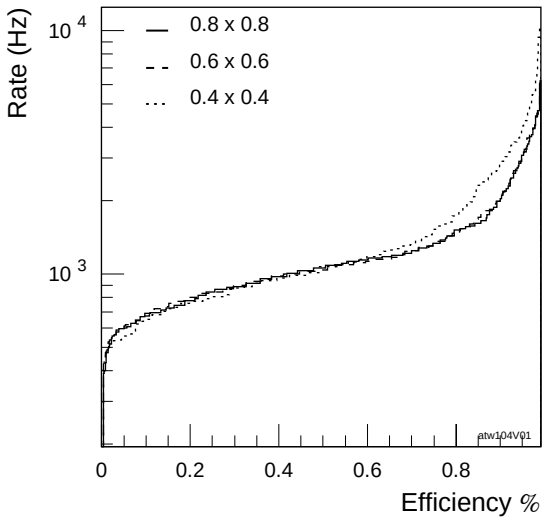

Figure 4-34 Trigger rate vs. efficiency for 200 GeV E_T jets, for different cluster sizes, at luminosity 10^{34} cm^{-2}s^{-1}.

Table 4-8 Inclusive jet trigger rates at luminosities of 10^{33} cm^{-2}s^{-1} and 10^{34} cm^{-2}s^{-1} for different algorithms.

Jet cluster size	Cluster step size	Rate for 100 GeV jets at 10^{33} cm^{-2}s^{-1}	Rate for 200 GeV jets at 10^{34} cm^{-2}s^{-1}
0.8×0.8	0.4	6.0 kHz	2.7 kHz
0.8×0.8	0.2	6.3 kHz	3.0 kHz
0.6×0.6	0.2	6.3 kHz	3.0 kHz
0.4×0.4	0.2	8.7 kHz	3.8 kHz

4.4.4 Spatial resolution and jet separation

While the resolution for inclusive high-E_T jets depends primarily on the trigger cluster size, the RoI coordinate resolution and the ability to resolve nearby jets depend on the step size and RoI definition.

Figure 4-35 shows the coordinate resolution of the three RoI definitions considered. All algorithms based on a step of 0.2×0.2 use a cluster of 0.4×0.4 as the RoI. For a step of 0.4×0.4, two RoI definitions are compared: requiring that an individual jet element be a local E_T maximum, or requiring that the 0.8×0.8 cluster be an E_T maximum. The quantity plotted is the distance $\Delta R = (\Delta\eta^2 + \Delta\phi^2)^{1/2}$ between the RoI centre and the axis of the reference jet. As can be seen, better resolution is obtained from both a smaller RoI cluster and a smaller step size.

Figure 4-36 illustrates the different abilities of the algorithms to resolve nearby jets. It shows the efficiency for the trigger to identify a reference jet as a function of its separation ΔR from a more energetic neighbour. As expected, the smaller RoI clusters and steps result in a higher efficiency to resolve nearby jets. This affects the acceptance of a multijet trigger, and the ability to count jets in events with complex topologies.

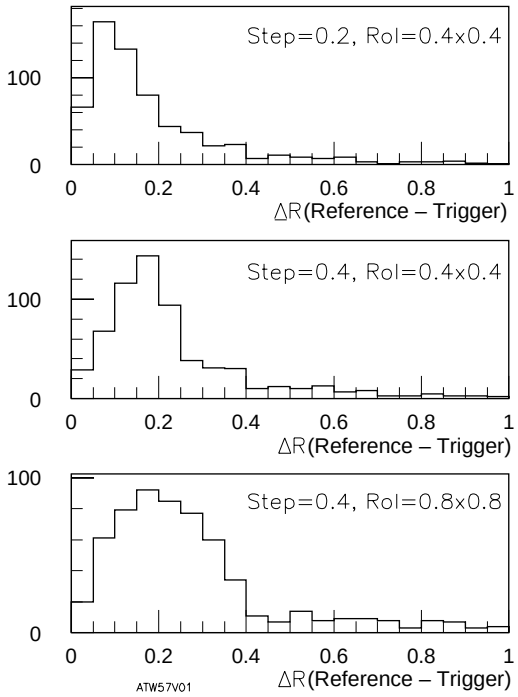

Figure 4-35 RoI coordinate precision for the three RoI algorithms, for jets of E_T > 100 GeV. Reference jets were required to be separated from their nearest neighbours by $\Delta R > 1.0$, to avoid confusion due to differing abilities of the trigger algorithms to resolve nearby jets.

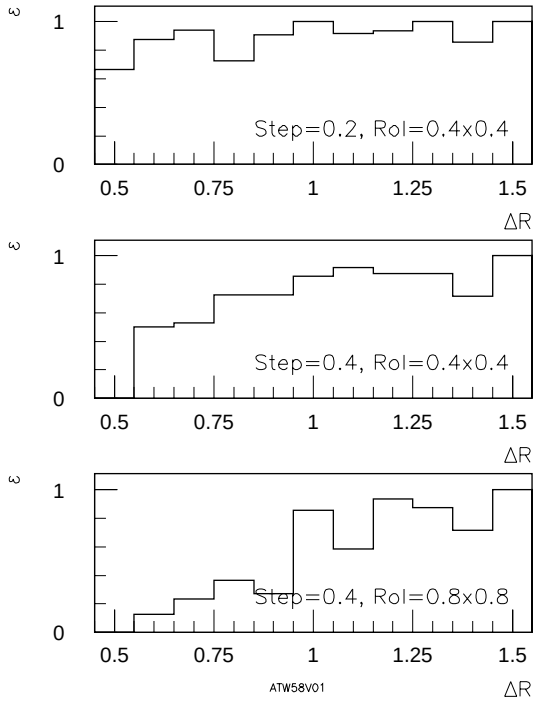

Figure 4-36 Efficiency to recognise two separate jets as a function of their separation, for the three RoI/decluster algorithms considered. In each case, the trigger was required to find distinct jets matching the two reference jets, with a separation between trigger and reference jet of $\Delta R < 0.6$.

4.4.5 Multijet trigger performance

Multijet trigger performance depends on both the resolution for isolated jets and the resolution for separating nearby jets. As an example, Figure 4-37 which shows the efficiency of a 4-jet trigger to select top events vs. the resulting trigger rate, for different algorithms. It can be seen that algorithms with a poorer two-jet separation produce a lower efficiency for a given trigger rate, even when such algorithms give slightly better inclusive jet trigger performance. Top production is used here as an example only — other triggers, such as inclusive leptons, are expected to be more important in top physics. Other multijet processes show these same effects to differing degrees.

The relative performance of the single-jet and four-jet triggers for a few physics processes are compared in Table 4-9. We see that the inclusive jet trigger performance is similar for algorithms based on steps of 0.4×0.4 and 0.2×0.2, but both a finer step and a smaller cluster size offer somewhat better multijet performance. While it is unclear whether any physics would necessarily be lost if the coarser step were used, the superior multijet performance of the finer-grained algorithms offers a more flexible set of trigger selections, and so is preferred.

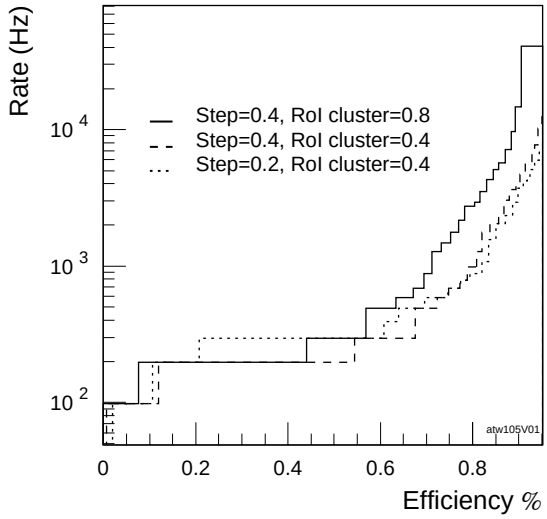

Figure 4-37 Efficiency vs. 4-jet trigger rate for top events at low luminosity, for three algorithms based on a cluster of 0.8×0.8. The differences are understood to be due to the different abilities of the algorithms to resolve nearby jets.

Table 4-9 One- and four-jet efficiencies and rates for some example physics processes, for different jet trigger algorithms. Thresholds chosen to give a trigger rate of ~1.5 kHz at luminosity 10^{33} $cm^{-2}s^{-1}$. For the algorithm based on a step of 0.4, the RoI was a single 0.4×0.4 jet element, and thus the same size as the RoI used in the other algorithms. For 0.8×0.8 jets the thresholds were 99 GeV for one-jet and 34 GeV for 4-jet triggers, and for 0.4×0.4 jets they were 87 GeV and 22 GeV, respectively.

		Top		Low-mass SUSY		hh —> bbbb	
Jet cluster size	**Step size**	**1 jet**	**4 jets**	**1 jet**	**4 jets**	**1 jet**	**4 jets**
0.8×0.8	0.4	61%	68%	75%	52%	60%	51%
0.8×0.8	0.2	58%	73%	72%	59%	58%	56%
0.4×0.4	0.2	56%	75%	71%	64%	59%	58%

4.4.6 Low-E_T RoI reconstruction

In addition to providing signals for use in inclusive jet triggers, multijet triggers, and combined triggers (such as jet and missing-E_T), the jet trigger system should flag 'secondary jet RoIs' which might be useful for more refined event selections at level-2. These would correspond to jets of lower E_T than those used in the level-1 event selections. For some processes, the offline analysis suggests that jets of 40 GeV E_T are of interest at LHC design luminosity, and as low as 15 GeV at lower luminosities. However, it has not been established at all yet that such jets would be needed at level-2. With the E_T of such jets being shared between many trigger towers (even the smallest jet cluster studied contains 32 towers), each of which has an r.m.s. noise level of several hundred MeV, it is clear that such low-E_T jets would present a formidable challenge to the level-1 calorimeter trigger. It is important though to understand what algorithm choices and parameters might optimize the ability to flag very low-E_T jets, and how low a jet E_T might still be separable from the noise in the level-1 trigger.

The problem of tagging low-E_T jet RoIs is illustrated in Figure 4-38. This shows an estimate of the mean RoI multiplicity as a function of jet cluster threshold, in a sample of events accepted

using an inclusive jet trigger. For each algorithm, there is a gradual increase in RoI multiplicity with decreasing threshold until the point is reached where noise fluctuations begin to become significant compared with the threshold, after which there is a catastrophic increase in RoI multiplicity. To optimize efficiency for low-E_T jets, the secondary RoI threshold should be set above this point, and the algorithm (as well as any other relevant parameters, e.g. thresholds on the E_T of trigger towers or jet elements) should be chosen to maximize efficiency for low-E_T jets while remaining above the catastrophe. Clearly, however, such a study is very sensitive to details of the detector response, namely the behaviour of small signals, the noise simulation and the processing of calorimeter signals, which are amongst the most difficult parts of the trigger chain to model accurately.

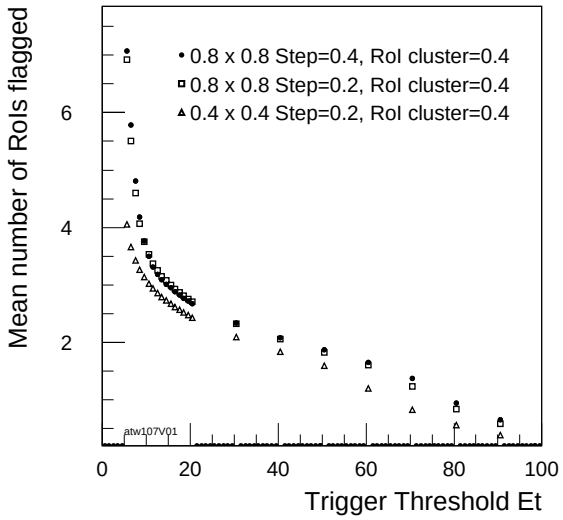

Figure 4-38 Mean RoI multiplicity vs. cluster threshold, for different algorithms. The events were selected by a 100 GeV inclusive jet trigger.

Using a ‘fast’ simulation, the ideal RoI performance of the different algorithms has been investigated for inclusive jet triggers. Table 4-10 shows the mean RoI multiplicities for different algorithms, for thresholds chosen to give 95% efficiency for the jet E_T values shown. These show that a smaller jet cluster results in a lower RoI multiplicity. This is understood to be due to the smaller area of the cluster resulting in a lower average noise E_T summed in each jet window. This is qualitatively different from the situation for high-E_T inclusive jet triggers, where resolution is dominated by jet containment rather than noise. In such an idealized detector model, even the lowest-E_T jets can be efficiently flagged with a reasonable RoI multiplicity. These results favour a system which allows the option of using a small jet cluster for RoI flagging.

Table 4-10 Mean RoI multiplicities for different algorithms, for low and high luminosity.

Jet cluster size	Cluster step size	Low luminosity E_T > 15 GeV	High luminosity E_T > 40 GeV
0.8×0.8	0.4	4.8	3.8
0.8×0.8	0.2	5.5	4.2
0.6×0.6	0.2	4.8	3.9
0.4×0.4	0.2	3.7	3.7

The situation is less favourable when the full (GEANT-based) detector simulation is used. Figure 4-39 compares trigger and reference jets in full and fast simulations. The two are in reasonable agreement. However, the trigger jet E_T in the full simulation is systematically slightly lower than in the fast simulation, and this effect becomes significant for low-E_T jets. This is shown in Figure 4-40. Here we see that for jets of less than 30 GeV (as found by the reference jet-finder), the trigger sees only a very low-E_T cluster, which is difficult to separate from the noise. The consequence is shown in Figure 4-41, which plots the efficiency for the trigger to find an RoI matched to a reference jet as a function of jet E_T, for a trigger threshold chosen to give an average RoI multiplicity in electron / photon-triggered events (assumed to dominate over jet triggers) of about three RoIs / event for a luminosity of $10^{33}\,cm^{-2}s^{-1}$. This suggests that efficient identification of 20 GeV jet RoIs might be possible at low luminosity, but lower-E_T jets would be problematic. Figure 4-42 shows similar plots at $10^{34}\,cm^{-2}s^{-1}$. Here we compare the use of a smaller cluster (0.4×0.4), or applying a threshold to the jet-element E_T, both done in order to suppress the contribution from pile-up. It can be seen that if a low RoI multiplicity is required, these techniques can improve the RoI efficiency, but the advantage is lost if a highter RoI multiplicity is tolerable.

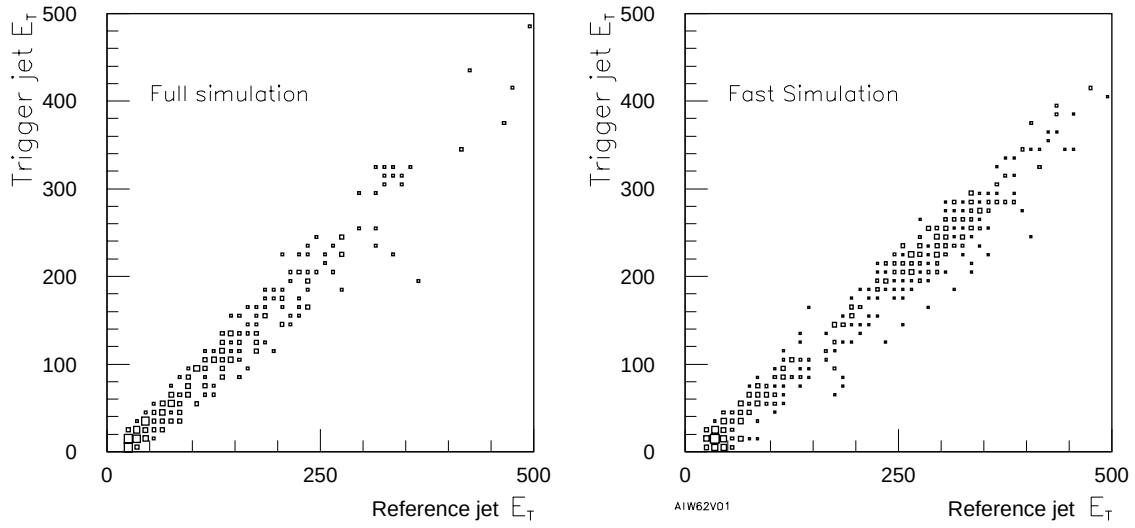

Figure 4-39 Trigger-jet E_T vs. reference jet E_T, for GEANT and fast simulations. The reference jet-finder used the KINE data as input in both cases for consistency.

4.4.7 Signal resolution and dynamic range

The jet-element signals are formed by summation of groups of trigger-tower signals in the Preprocessor, after digitization and bunch-crossing identification (BCID). The resolution of the signals cannot then be better than that of the trigger towers. In light of the anticipated importance of low-E_T jet tagging, it would seem prudent to use the full resolution of the trigger towers in jet reconstruction, and not to increase the least significant bit (which would allow a reduction in data transfer).

Since the currently-understood physics requirements include triggers on jets of E_T = 200 GeV, it would seem prudent to require that the trigger be able to set a meaningful threshold somewhat higher than this, preferably up to 300–400 GeV. Jet elements of 0.2×0.2 will be formed separately in the Preprocessor for the electromagnetic and hadronic calorimeters, and only combined in the jet processor itself. Thus there are two dynamic ranges to consider: that of the

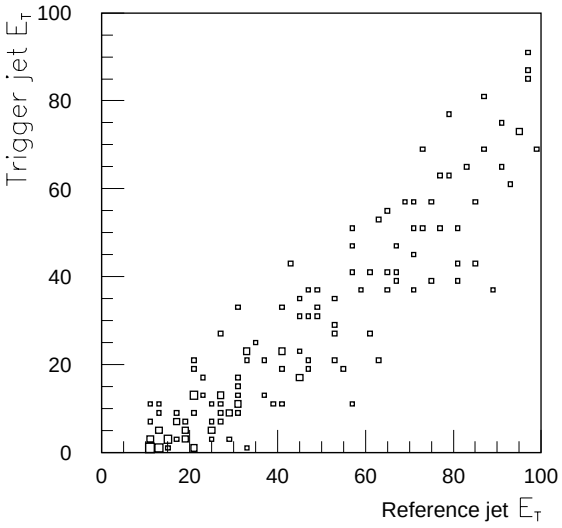

Figure 4-40 E_T of a reference jet compared with that reconstructed by the trigger. Jets were required to be well separated from their neighbours for this study.

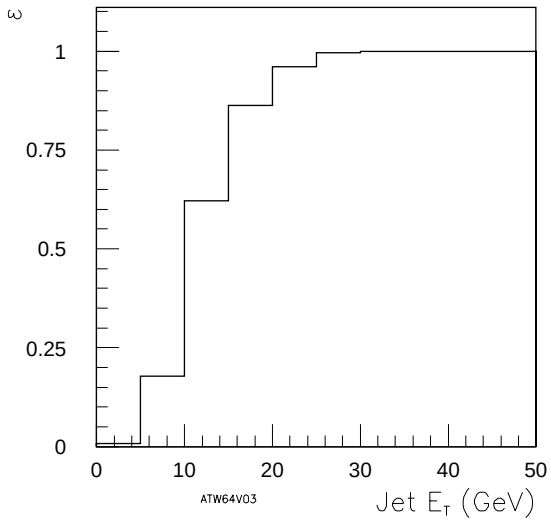

Figure 4-41 Efficiency to flag a jet RoI as a function of jet E_T. The trigger threshold was chosen to give an average RoI multiplicity in electron/photon-triggered events of about three per event. The algorithm used was a cluster of 0.8×0.8, sliding by 0.2.

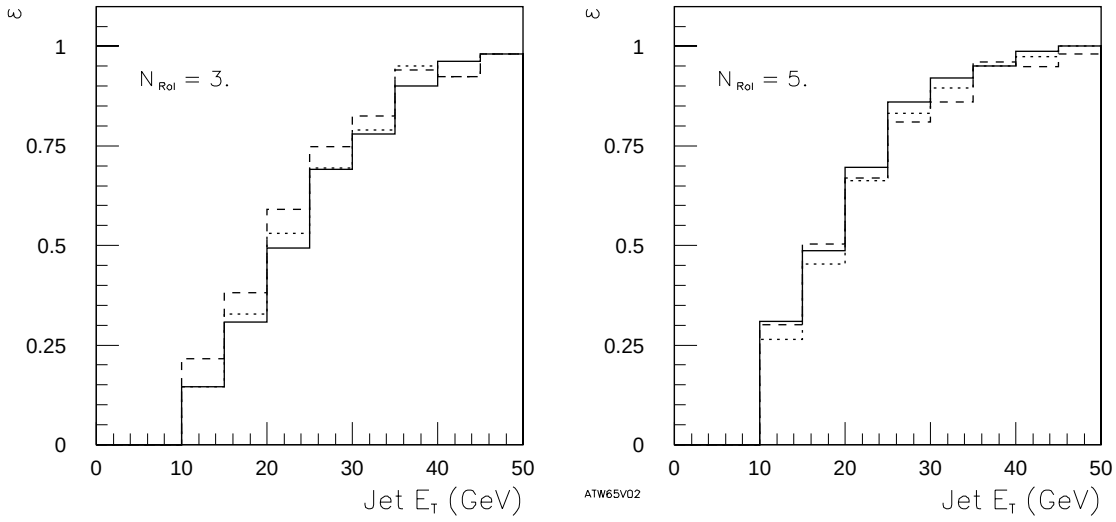

Figure 4-42 Efficiency to flag a jet RoI as a function of jet E_T, at luminosity 10^{34} cm^{-2}s^{-1}, for two different average RoI multiplicities. The plots compare a jet of 0.8×0.8 with no threshold on jet-element E_T (solid), the same cluster but using only elements with $E_T \geq 3$ GeV (dashed), and a jet of 0.4×0.4 (dotted).

separate e.m. and hadronic jet element signals (which must be transferred between crates) and that of the combined signals used within the jet processor. Figure 4-43 shows the correlation between the reference-jet E_T and trigger-jet E_T, for maximum counts for the separate (e.m. and hadronic) jet elements of 255 GeV and 511 GeV, and the degradation of the measurement compared with a maximum count of 1023 GeV. Figure 4-44 shows the same correlations for the maximum count of the combined jet elements. As can be seen, a maximum count of 255 GeV is clearly unsuitable when electromagnetic and hadronic data are combined, and causes some degradation of performance even when separate sums are formed. With a maximum count of 511 GeV, it is clear that the finite range of the jet element signals is not the limiting factor on

performance. In all plots the limit on a single trigger-tower E_T is 255 GeV, and the effect of this on trigger measurements for high-E_T jets is clearly visible.

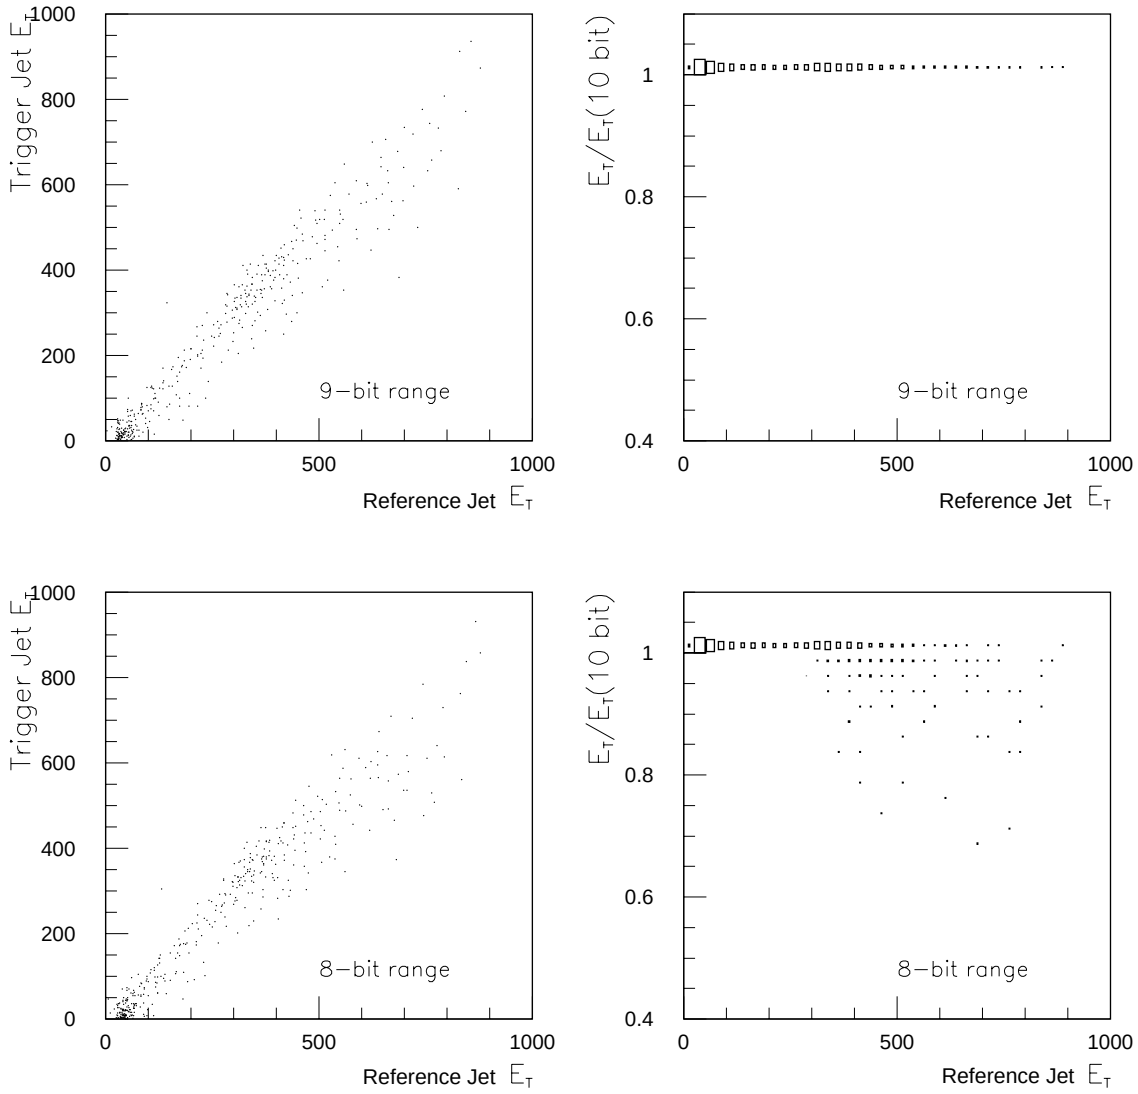

Figure 4-43 Effect of finite dynamic range of 0.2×0.2 jet elements, divided into e.m. and hadronic signals. Ranges are 1–511 GeV (9 bits) and 1–255 GeV (8 bits). Left-hand plots show trigger-jet E_T vs. reference-jet E_T, while right-hand plots show ratio of trigger-jet E_T to that found with a 10-bit (1–1023) range, as a function of reference-jet E_T.

From these results, the preference is that all jet element signals should have a dynamic range extending to approximately 500 GeV (or in any case significantly above the single trigger-tower limit of 255 GeV). It should be noted, however, that this may not be achievable throughout the calorimeter: for $|\eta| > 2.5$, the trigger towers themselves have a granularity of 0.2×0.2, and so the jet element (within a single calorimeter) is just equal to the trigger tower. If these towers have a maximum count of 255 GeV, that will be the maximum count of the corresponding jet element.

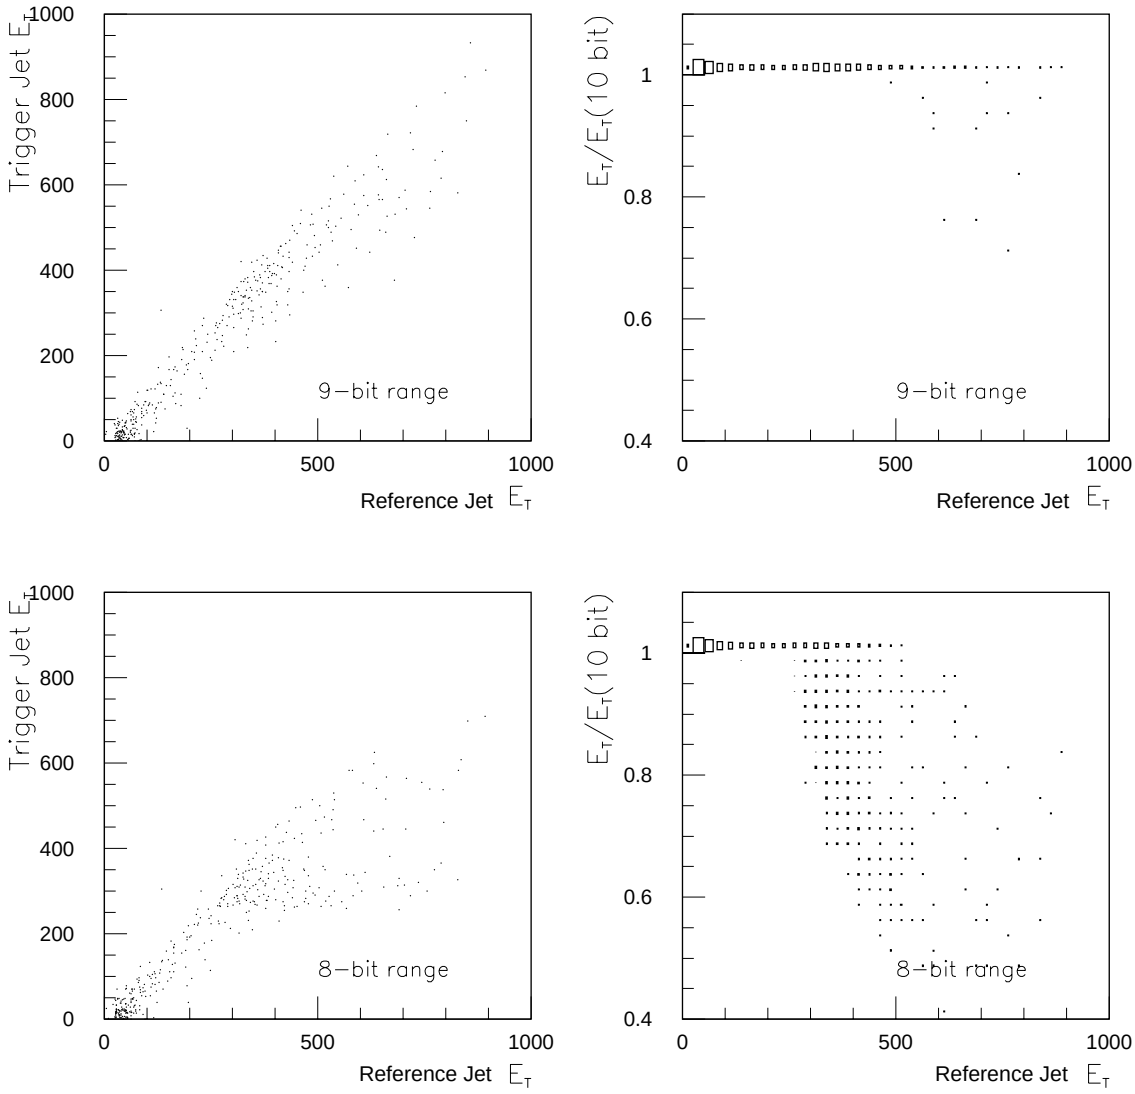

Figure 4-44 Effect of finite dynamic range of jet elements, after summation of e.m. and hadronic signals.

4.5 Missing transverse energy and total transverse energy triggers

Missing transverse energy will be one of the distinct signatures on which to select interesting physics processes at the LHC. Many extensions of the Standard Model include weakly-interacting particles which, if produced at the LHC, will escape detection. Their presence will, however, be signalled by an imbalance of transverse momentum.

Among the basic building blocks of the level-1 calorimeter trigger is the summation of the transverse energy deposited in the calorimeters. The total scalar E_T, as well as the components E_x and E_y in the plane transverse to the beam axis, are computed in the Jet/Energy-sum Processor of the calorimeter trigger. Although the missing-E_T trigger itself is not included in the basic level-1 inclusive triggers, its combination with the single jet, electron/photon, and hadron/tau triggers is important to allow triggering on interesting events with low jet or tau transverse-energy thresholds.

QCD jet events dominate the high-p_T cross-section at the LHC, and provide a range of total E_T without showing any intrinsic missing E_T since, apart from neutrinos in the fragmentation cascade, no other non-interacting particles are present.

A total transverse energy trigger may serve firstly as a kind of low-bias trigger, and secondly in order to be prepared to trigger on unexpected event shapes.

The two ‘energy-sum’ triggers are discussed together because many similar issues occur in both, and the simulations were done together.

4.5.1 Description of the simulation

The simulation results described below are based on the fast simulation program ATLFAST [4-6]. The physics events are generated by ISAJET and PYTHIA [4-7]. Most of the results obtained so far are based on the following processes:

- $W \rightarrow e\nu$;
- SUSY events with squark/gluino masses of 200, 400 and 1000 GeV;
- QCD events, mainly with $E_T(\text{jet}) > 20$ GeV, but some sub-samples at other thresholds were also used;
- fully simulated minimum-bias events.

Compared to the standard version, some important modifications have recently been made to ATLFAST, mainly concerning the response and the resolution of the calorimeters together with a correct description of the longitudinal energy-sharing between the electromagnetic and the hadronic calorimeters. The calorimeter response and resolution functions have been determined from a full ATLAS GEANT simulation for both single charged pions and photons in the full energy range of interest, from 200 MeV up to 100 GeV. In the same simulation the energy-dependent longitudinal energy sharing in the calorimeters was determined. It should be noted that all effects of energy losses in the tracking detectors and in the insensitive material are taken into account in the response and resolution functions. The deposited energies in the calorimeters are summed into trigger towers with granularities in $\Delta\eta \times \Delta\phi$ according to the proposed scheme over the full range of pseudorapidity (i.e. mostly 0.1×0.1, except in the endcap regions). Electronic noise is added for each trigger tower, assuming Gaussian

distributions with standard deviations as given in the ATLAS Calorimeter TDRs [4-8], [4-9]. For the high-luminosity studies, minimum-bias pile-up events are added.

After adding up the energy of the trigger towers in the event, including the contribution of pile-up, shaping is applied with a characteristic shaping function (as determined from beam tests of the different calorimeter types). The fact that the calorimeter signals extend over several bunch-crossings means that the contributions of energy depositions from previous events to the event of interest must be taken into account. In the simulation, this was done by tracing back the single-cell energies 25 bunch-crossings. The positive part of the signal extends over a few bunch-crossings, followed in the liquid-argon calorimeters by a negative undershoot which extends over 20 bunch-crossings. The contributions of events earlier in time were then added to the energy deposited in the event of interest.

Effects from the electronics chain of the calorimeter trigger Preprocessor are simulated. These include the upper limit on the dynamic range of the FADC, which is 255 GeV, bunch-crossing identification (BCID), and lookup tables which among other things apply noise thresholds. Although the ADC uses 10 bits, after BCID the signals are truncated to 8 bits with a step size of about 1 GeV. Finally, the calorimeter energies are summed into a map with a granularity of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$, which contains the basic energy information that has been used for studies of the performance of the Jet / Energy-sum Processor.

4.5.2 Results of the simulation

4.5.2.1 Missing-E_T resolution

The performance of the missing-E_T trigger can best be parametrized by its resolution function, i.e. in a way independent of special event types and physics processes considered. This resolution function is generally found to depend on the total transverse energy deposited in the calorimeters. Of particular importance to the level-1 trigger is the question of how much the missing-E_T resolution is degraded by the trigger hardware implementation, including digitization, threshold setting, signal transmission with a limited number of bits, and the adder tree in the summation chain.

The missing-E_T resolution as a function of the total transverse energy in the event was therefore determined at the various stages of the level-1 trigger hardware. Using QCD jet events, which should have little intrinsic missing E_T, the r.m.s. width of the distribution of the components E_x and E_y of the total energy sum is used as a measure of the resolution. As an example, the distribution of E_x for QCD events with a total transverse energy in the range between 300 and 350 GeV at the various steps of the simulation was considered. The missing-E_T resolution is dominated by the calorimeter resolution and response, and by the addition of electronic noise in the tower-builder electronics. The numbers are summarized in Table 4-11. Truncating the digitized values for the tower energies to eight bits effectively applies a 1 GeV threshold to each trigger tower, which reduces the noise contribution to the resolution. All results given here are for low-luminosity running.

For the E_T range considered, no degradation in resolution is seen in the subsequent summation steps if the summing into $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ is performed using at least nine bits.

The dependence of the resolution, represented by the r.m.s. in E_x, on the value of total E_T is shown in Figure 4-45. It should be noted that the r.m.s. values are computed including non-

Table 4-11 Values for the r.m.s. resolution of total E_x at various stages of the missing-E_T determination, for jets with transverse energy in the range 300–350 GeV.

Summation step	Resolution (r.m.s.)
Particle level, after fragmentation	3.8 GeV
Particle level, with cut $\|\eta\| < 5.0$	5.8 GeV
Calorimeter response and resolution	11.9 GeV
Non-compensating calorimeter	13.0 GeV
Addition of trigger tower noise	18.4 GeV
Digitized trigger towers, 8 bits, 1 GeV steps	14.7 GeV

Gaussian tails. Therefore they do not exactly follow the expected scaling with total E_T. The resolution obtained at the calorimeter level can almost be maintained at the trigger level up to very high values of total E_T. For total E_T below ~1000 GeV, which corresponds to a typical two-jet event with E_T values of ~500 GeV, no losses due to the ADC cutoff at 256 GeV are visible. Beyond that value the tails in the distributions increase significantly. This value is sufficiently high that the overall performance of the level-1 trigger is not affected.

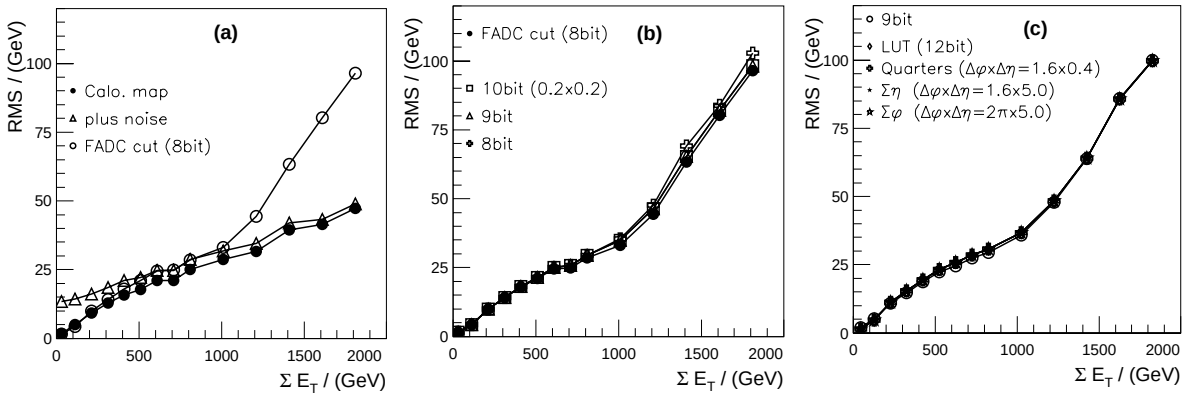

Figure 4-45 Dependence of the resolution of the E_x, E_y components of the total transverse momentum on total E_T. (a) after the trigger Preprocessor, in comparison to the resolution obtained at the calorimeter level, using the trigger tower granularity for the E_T calculation; (b) after the transmission of the summed 0.2×0.2 element energies to the Jet/Energy-sum Processor, using 8, 9 (default), or 10 bits; (c) after various summation steps in the Jet/Energy-sum Processor.

The degradation in resolution observed at high total E_T has no impact on the physics. Figure 4-46 shows an extreme case of top events with top $p_T > 500$ GeV, where a dynamic range of 1 TeV instead of 256 GeV would be beneficial. In these events we observed that the standard deviation of the difference between generated and trigger values increases from 36 GeV, which is the best we would expect for a dynamic range of 1 TeV, to 44 GeV, which is thus a reference for any further reduction of the dynamic range.

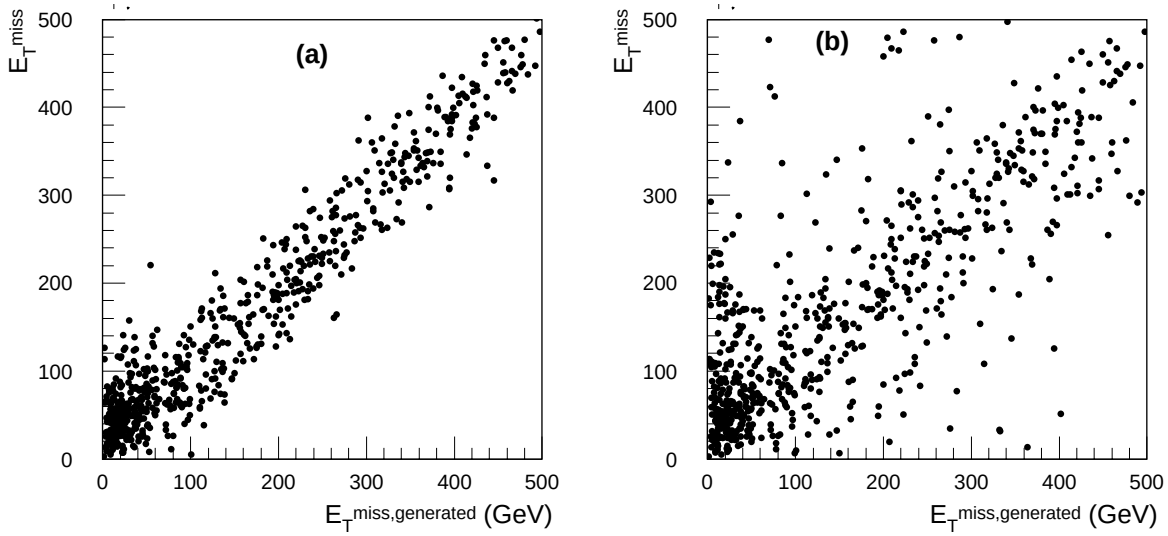

Figure 4-46 High-E_T top events (parton E_T > 500 GeV): missing E_T in trigger logic vs. generated. (a) dynamic range 1 TeV, (b) dynamic range 256 GeV.

4.5.2.2 Number of bits used for signal transmission

Figure 4-45b shows the dependence of the resolution on the number of bits used for signal transmission from the Preprocessor to the Jet/Energy-sum Processor. For the proposed 9-bit transmission no significant degradation is seen up to very high transverse energies.

4.5.2.3 Adder tree

Effects from the hardware summation tree are shown in Figure 4-45 for the various summation steps, which consist of:

- summation of the electromagnetic and the hadronic calorimeter energies with a granularity of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ (dynamic range 10 bits);
- summation of two cells in η to $\Delta\eta \times \Delta\phi = 0.4 \times 0.2$ (dynamic range 11 bits);
- transformation to E_x and E_y using 12-bit lookup tables;
- summation over ϕ on one module (see Section 6.4.2), i.e. $\Delta\eta \times \Delta\phi = 0.4 \times 1.6$;
- summation over the full range in η (12 bits);
- summation over the four quadrants in ϕ (12 bits), applying the appropriate signs.

As can be seen from the figure, the resolution is not further degraded by any of the summation steps, so the proposed hardware design for the missing-E_T summation is adequate.

4.5.2.4 Rapidity coverage of missing-E_T trigger

The performance of the missing-E_T trigger depends critically on the rapidity coverage. In Figure 4-47 the dependence on η coverage is shown for QCD events, which are expected to give the largest contribution to fake missing-E_T triggers. The η coverage was varied from 3 to 5, which requires use of signals from the forward calorimeters (FCAL). From these plots it is clear

that the largest possible η coverage is needed in order to maintain a good correlation between generated missing energy and the missing E_T as seen in the trigger (Figure 4-47a). With the FCAL included, the background due to QCD events will be low enough to allow use of low missing-E_T trigger thresholds. In Table 4-12 the mean value of missing E_T in QCD events is given as a function of the η range.

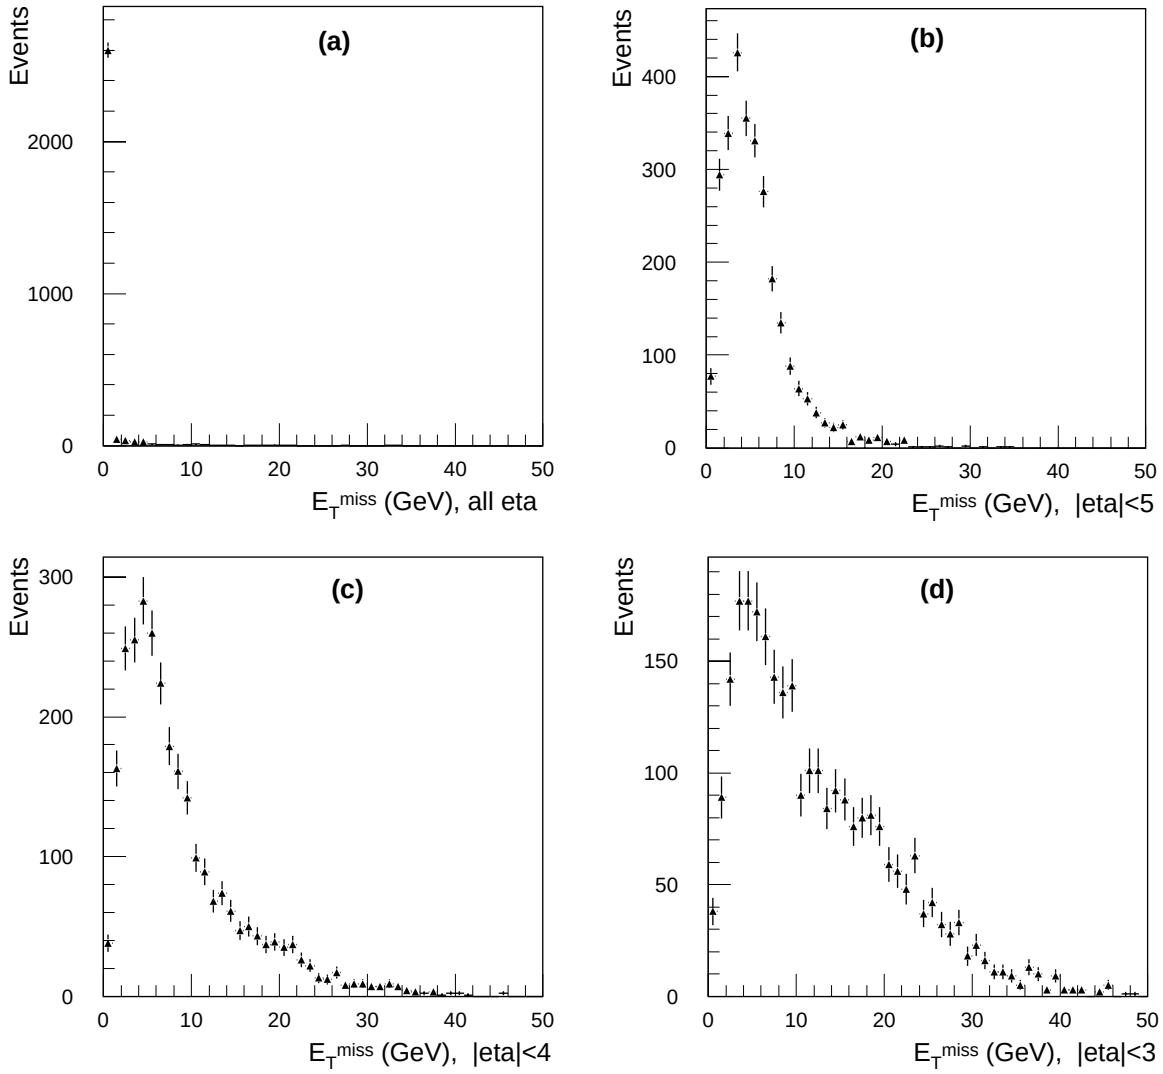

Figure 4-47 Missing-E_T distributions for four rapidity ranges. (a) generated, (b) |η| < 5, (c) |η| < 4, (d) |η| < 3.

Table 4-12 Mean value of missing E_T as a function of rapidity coverage.

Rapidity coverage	Mean value of missing-E_T
All	0.9 GeV
$\|\eta\| < 5$	5.6 GeV
$\|\eta\| < 4$	8.8 GeV
$\|\eta\| < 3$	12.5 GeV

4.5.3 Trigger rates at level-1

The calculation of the inclusive missing-E_T spectrum is a difficult task, because contributions from instrumental effects are important in addition to the contributions from physics channels. Among the effects which cannot be simulated reliably at present are machine-induced backgrounds, beam–gas interactions, as well as contributions resulting from extreme tails in the detector performance.

4.5.3.1 Rates in different processes

The relative rates for minimum-bias events and W decays as a function of total E_T are compared in Figure 4-48. The signals before digitization are shown in (a), and in (b) we show the effect of the level-1 trigger logic. After BCID the contribution of minimum-bias events to the rates of both the missing-E_T trigger and the total E_T triggers are expected to be negligible when reasonable thresholds (missing E_T > 50 GeV, total-E_T > 600 GeV) are applied. One can conclude that the contribution to the total E_T of the minimum-bias events, which is of the order of 250 GeV before signal handling, is expected to be only 20–30 GeV after including the effects of the signal treatment in the trigger logic mentioned above. For the missing-E_T trigger the signals before digitization show significant tails, which extend to missing-E_T = 100 GeV. On the other hand, the signals used to form a level-1 trigger decision (i.e. signals after BCID) show significantly-reduced tails, hopefully allowing a threshold setting much below 100 GeV. We claim that the reduction of background contributions and the much sharper thresholds after BCID are due to the long negative undershoot of the of the signals, which provides a very effective filter of pile-up noise from previous events. In the event of interest itself, the BCID, removes out-of-time energy deposits, and the noise cut within the lookup table, set to 1 GeV in the current simulations, also removes much of the problem.

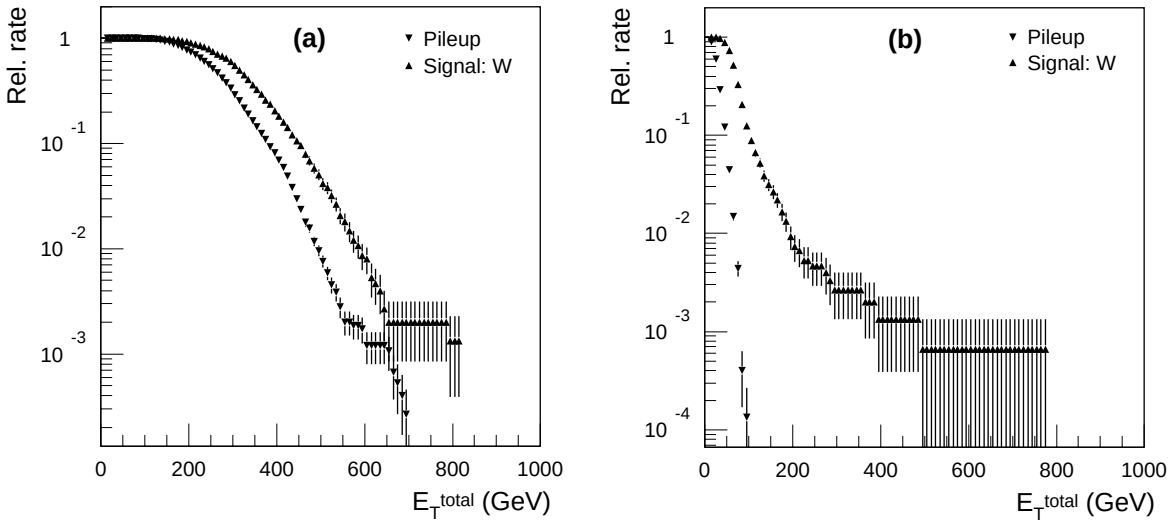

Figure 4-48 Rate vs. total E_T for minimum-bias and W events. (a) energy before digitization, (b) after calorimeter trigger logic.

4.5.3.2 Contributions to missing-E_T from QCD jet events

The same comparison as in Section 4.5.3.1 has been made for a sample of two-jet events, where each jet was required to have an E_T larger than 100 Gev at the generator level. When looking at

the input signals before digitization, the total E_T in these events extends to 1 TeV, while after BCID the mean value of the total E_T is of the order of 300 GeV. A threshold giving a total rate for these events below 1 kHz would be of the order of 600 GeV. It was found that the contribution to missing E_T from QCD events is expected to extend to 50–150 GeV after BCID. If a rate of 1 kHz is acceptable for such a trigger, thresholds of 50–100 GeV should be possible, which would allow a missing-E_T trigger on $W \rightarrow e\nu$.

Minimum-bias and QCD jet events will contribute significantly to the missing-E_T due to their high rate. For the absolute rate computation, the total inelastic proton–proton cross section of ~70 mb is simulated by applying a very low-p_T cutoff in the PYTHIA simulation of two-jet events in the region of ~4 GeV.

A combination of the missing-E_T trigger with the single-jet trigger is among the basic level-1 triggers of ATLAS. The rate estimates start from the missing-E_T spectrum reconstructed at trigger level. These are shown in Figure 4-49 as a function of the threshold applied in missing E_T, with curves showing the result of combining missing-E_T with different jet thresholds. Further results on combinations of the missing-E_T signature with other triggers are given in [4-10].

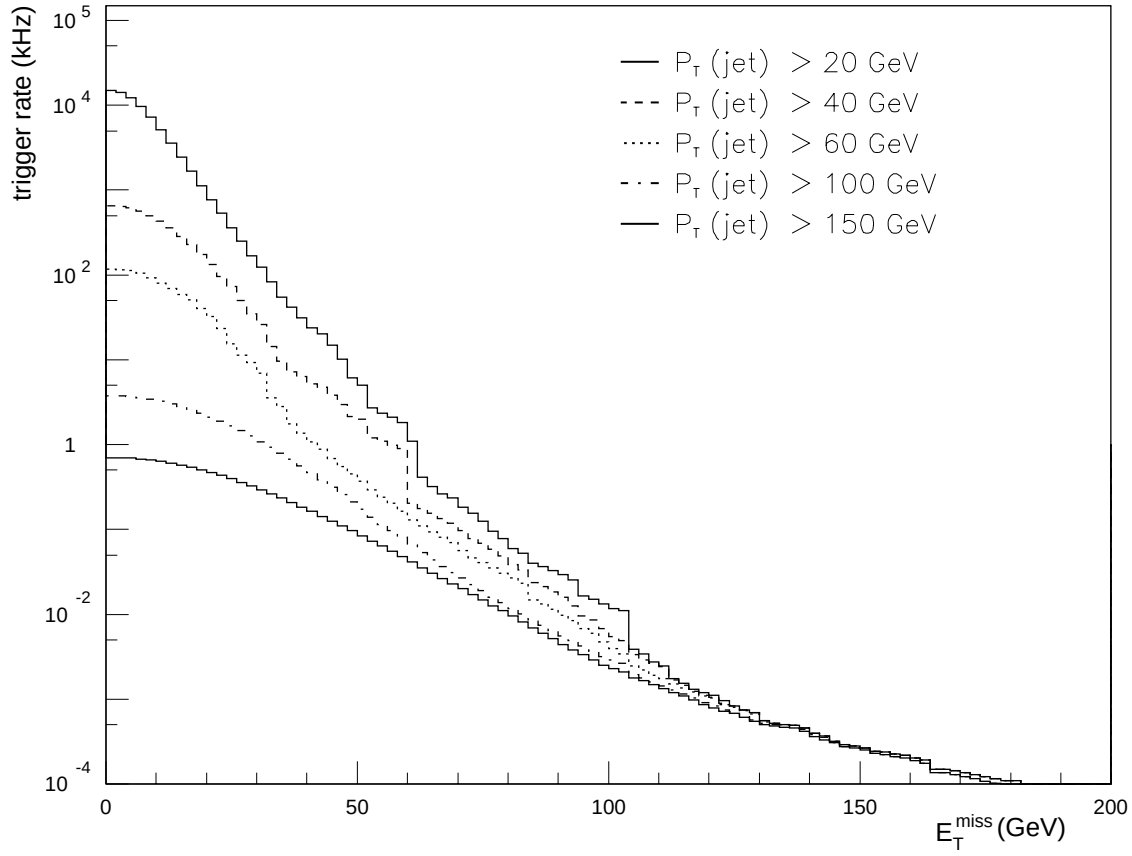

Figure 4-49 Level-1 trigger rates using combined single-jet and missing-E_T signatures. The rates are given as a function of the missing-E_T threshold for different jet thresholds. On average 2.3 minimum-bias events have been added to the hard collision.

4.5.3.3 Rates for physics processes with missing-E_T

Rates for events containing W bosons for an exclusive trigger on missing E_T are overlaid in Figure 4-50. The rate of QCD events triggered by the missing-E_T trigger, calculated for the limited range $|\eta| \leq 3.2$, is far higher than the rate for the range $|\eta| \leq 5.0$. The rate of minimum-bias pile-up is dropping rapidly, and we estimate negligible rates with the missing-E_T trigger threshold as low as 50–100 GeV, depending on the details of the pile-up simulation. Such physics processes can be triggered with efficiencies as shown in Table 4-13 for missing E_T and Table 4-14 for total E_T at acceptable trigger rates of 1 kHz. We observe that if a rate of 1 kHz is acceptable, we might be able to trigger on a significant number of Ws, resulting in a sample of missing-E_T events, which might be used for calibration purposes.

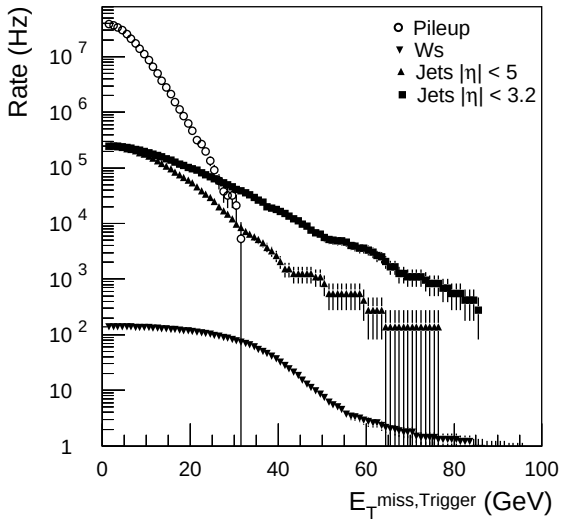

Figure 4-50 Rate vs. missing-E_T threshold for various processes; QCD events with a parton p_T > 20 GeV.

Table 4-13 Efficiencies for W and SUSY decays at different stages of missing-E_T processing, for a missing-E_T rate of 1 kHz.

	Missing-E_T threshold	W decays	SUSY low	SUSY medium	SUSY high
Before digitization	87 GeV	1.2%	37.8%	70.7%	97.5%
After BCID	51 GeV	4.3%	66.7%	86.9%	98.8%

Table 4-14 Efficiencies for SUSY decays at different stages of total-E_T processing, for a total-E_T rate of 1 kHz.

	Total-E_T threshold	SUSY low	SUSY medium	SUSY high
Before digitization	879 GeV	19.0%	63.8%	97.6%
After BCID	336 GeV	61.6%	96.6%	99.7%

4.6 Other trigger possibilities

In addition to the triggers discussed above, a number of other possibilities have been proposed and are being investigated. Two candidates that may be regarded as realistic options will be mentioned below. These are triggers on localized energy in the forward calorimeters, and a trigger on total jet E_T.

4.6.1 Localized energy in the FCAL

A suggestion to trigger on energy deposition in the FCAL has been made. One of the main goals would be to study diffractive physics at low luminosity. This is not difficult to implement, provided the E_T summing logic is organized in a convenient manner. A clear physics case for this trigger has yet to be made, and at present it is being kept in mind for the hardware design.

4.6.2 Total jet transverse energy

4.6.2.1 Introduction

Total scalar E_T, particularly under LHC conditions with many interactions per bunch-crossing, is very likely to be a trigger that selects noisy and problematic events as well as machine background. A trigger on the total transverse energy in jets is more likely to be of use. Furthermore, many ATLAS physics analyses, particularly in the SUSY field, use the total E_T of reconstructed jets and leptons as an event-selection variable rather than the total scalar E_T in the calorimeters. The sum of the transverse energy of all jets found by the level-1 calorimeter trigger would be a reasonable analogue to this quantity, since the jet trigger will also be sensitive to electrons and taus (though not muons), provided these are energetic enough to pass one or more jet-trigger thresholds. A trigger based on such a variable would be complementary to the total scalar-E_T trigger described above.

The ‘obvious’ way to construct such a trigger would be to output the E_T of each accepted jet candidate, and to sum these over the entire acceptance of the jet trigger, or some region within it. This would incur the overheads of outputting the jet E_T values from the processors, some additional data transfers, and a summing tree to add them. This ‘full’ algorithm does not fit in very well with the current design for the calorimeter trigger. However, a simple algorithm which approximates such a trigger seems to work surprisingly well in simulations, and can be very easily implemented. It is currently proposed to do this in the Central Trigger Processor, as described in Section 15.2.3.1. An alternative approach within the calorimeter trigger, using the same estimator, is briefly mentioned in Section 6.4.8.

4.6.2.2 Multiplicity-based estimator

The simple approximation algorithm uses the multiplicities of jets passing the different jet-E_T thresholds to estimate the total E_T of the jet system. To see how this works, consider an event containing n jets with $E_T > x$ GeV and m jets with $E_T > y$ GeV, with $n > m$ and $y > x$. There are then $n - m$ jets in the E_T range x GeV $\leq E_T \leq y$ GeV. These jets then have a total E_T of at least $(n - m)x$ GeV and at most $(n - m)y$ GeV. A number of possible estimators of the total jet E_T can be constructed in this way, such as lower limits on jet E_T (assuming each jet lies exactly at the

lower end of the range, upper limits, or indeed both together. For comparison with the 'true' total jet-E_T algorithm, the following estimator was used:

$$\sum E_T^{jet} = N(j_{max})\, E_T^{thresh}(j_{max}) + \sum_{j=2}^{j_{max}} \left(N(j-1) - N(j)\right) \frac{E_T^{thresh}(j-1) + E_T^{thresh}(j)}{2}$$

This variable treats any jets between $E_T(j-1)$ and $E_T(j)$ as lying half-way between the two thresholds, while any jet passing the highest threshold is assumed to have E_T equal to that threshold. For events with no jets passing the highest threshold, this is the average of the 'lower limit' and 'upper limit' approaches

The first term, before the summation, illustrates a limitation of this method: for jets passing this highest threshold, all we know about their E_T is a lower limit, and hence this approach must underestimate the E_T of events containing such jets, whatever estimator is actually used. In practice this does not affect trigger efficiency, since any event with a jet passing the highest jet threshold will be accepted. Nor is this necessarily a problem for the level-2 trigger, since the level-2 processors would know that the events contained such a jet and hence that the total jet-E_T estimator was unreliable.

4.6.2.3 Performance of the estimator

For both the 'full' and the simple 'multiplicity-based' total jet-E_T algorithms, the following questions must be addressed:

1. How sensitive is the trigger to the jet cluster size used?
2. How sensitive is the trigger to the lowest jet-E_T threshold used?

The multiplicity-based estimator might in addition be sensitive to the actual values of the jet-E_T thresholds. If this sensitivity were too strong, we might require additional jet thresholds purely for the operation of this trigger, which would be a significant overhead.

The basic properties of the two algorithms, and their potential advantages and disadvantages, are illustrated in Figures 4-51 and 4-52. These plot the 'true' total E_T in the event (sum of final-state 4-vectors, excluding muons, neutrinos and pile-up) against the different total jet-E_T estimators, at low and high luminosity. A 'fast' simulation of the ATLAS detector was used in these studies. The most significant observations are:

- The total jet E_T has a larger dynamic range than the multiplicity-based estimator; as noted above, this has no effect on trigger efficiency.
- The total jet E_T is less sensitive to pile-up if a smaller jet cluster is used.
- The multiplicity-based estimator is comparatively insensitive to pile-up.

These plots are only for one set of parameters, though. In order to make a well-founded decision, the sensitivities to different parameters must be investigated.

In Figures 4-51 and 4-52, we see that the total jet E_T is potentially very sensitive to pile-up, especially if a large jet cluster is used. However, this sensitivity is, as would be expected, easily controlled by adjusting the threshold for the lowest-E_T jets contributing to this trigger. This is illustrated in Figure 4-53.

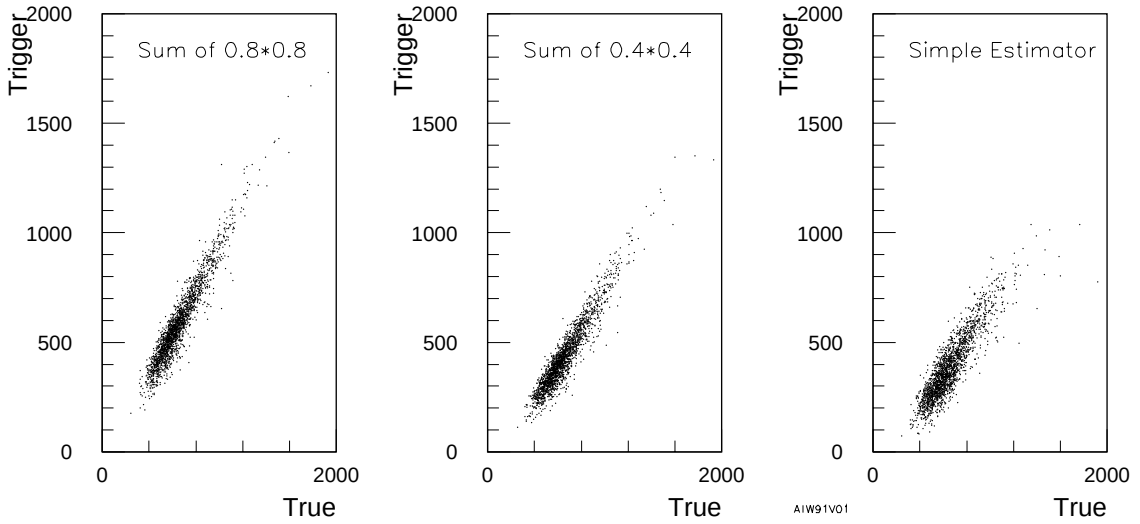

Figure 4-51 Total event E_T vs. different trigger estimators, for low luminosity. All E_T sums are within $|\eta| < 3.2$. The minimum jet-E_T threshold was 5 GeV in all cases. For the simple multiplicity-based estimator the maximum jet-E_T threshold was set at 250 GeV.

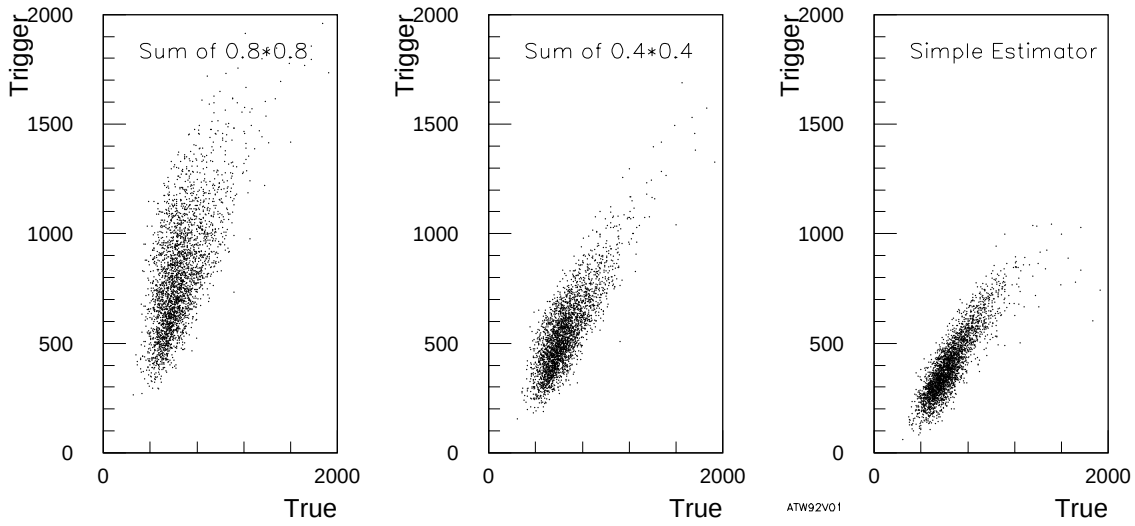

Figure 4-52 Total event E_T vs. different trigger estimators, for high luminosity. All other parameters as in Figure 4-51.

The jet trigger design foresees the possibility of using different jet cluster sizes for different purposes in the trigger, as discussed in Section 4.4.1. The sensitivity of the total jet-E_T trigger to cluster size is illustrated in Figure 4-54, for both low and high luminosity. While there is some sensitivity to luminosity (and to a lesser extent to cluster size), it can be seen from these figures that this is much less strong than for the 'full' total jet-E_T algorithm. Thus a multiplicity-based total jet-E_T trigger should not impose any significant constraints on the choice of cluster sizes for the jet trigger. In both cases, jet thresholds of 150, 100, 75, 50, 30, 20, 10, and 5 GeV were used, though the trigger does not appear to be very sensitive to the precise values of the jet thresholds.

The relative insensitivity of the multiplicity-based estimator to pile-up is understandable in that, while the full total jet-E_T trigger sees every fluctuation in pile-up noise within a cluster

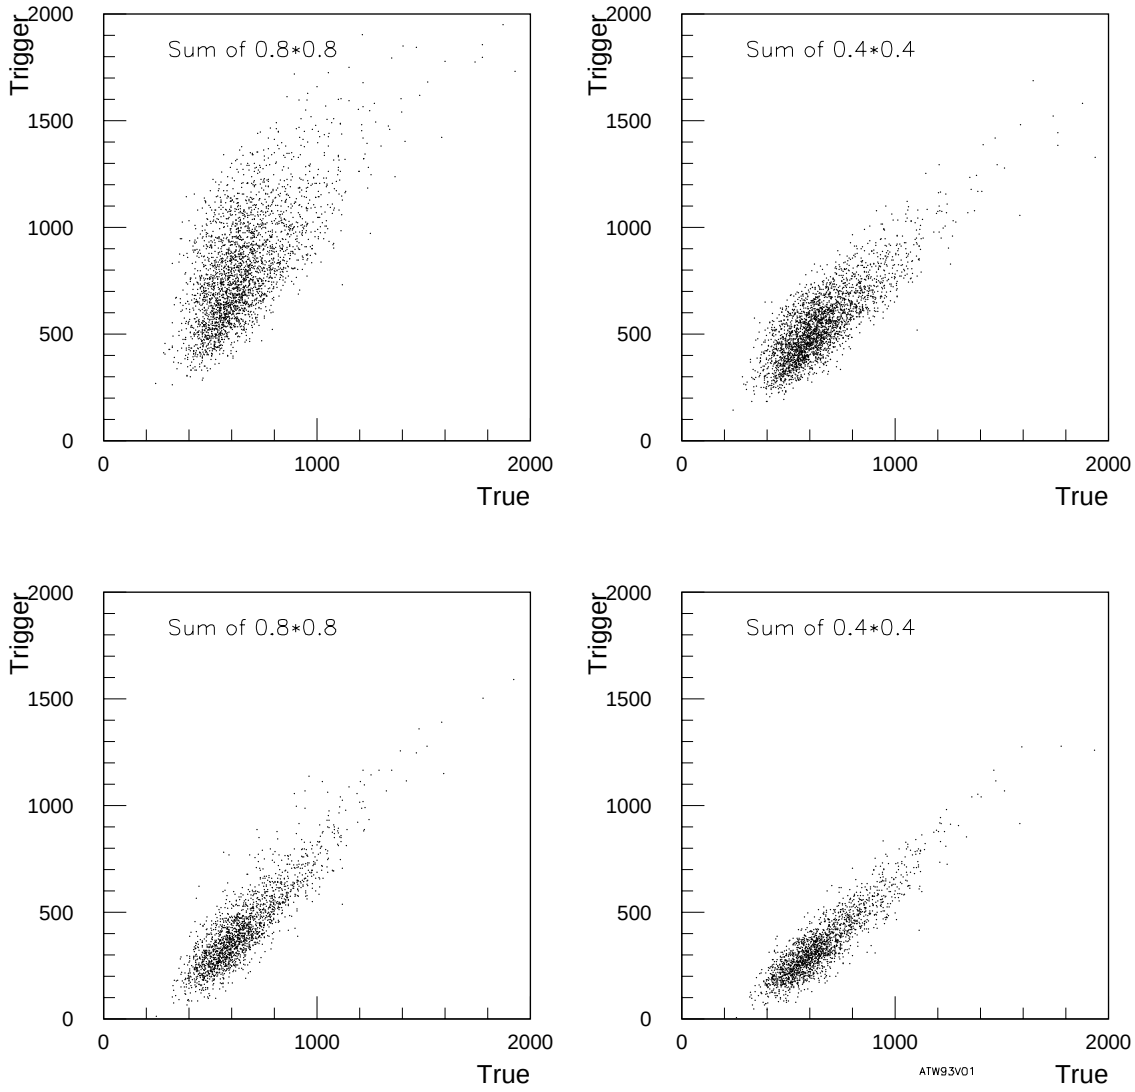

Figure 4-53 Total event E_T vs 'true' jet E_T at high luminosity, for a minimum jet E_T of 5 GeV (upper plots) and 10 GeV (lower plots).

passing the lowest jet-E_T threshold, the multiplicity-based estimator is sensitive only if those fluctuations change the highest threshold that a given jet passes. In a similar manner, the multiplicity-based estimator is found to be less strongly sensitive to the lowest jet-E_T threshold used.

The above plots are a qualitative guide to the eye, but are not easy to base quantitative judgements on. What really matters for the trigger is the trigger efficiency for a given rate. This is illustrated in Table 4-15 for low-luminosity top and SUSY selections. The SUSY data were for SUGRA Point 4 of the LHCC SUSY workshop [4-11], which gives an LSP mass of 80 GeV and squark / gluino / slepton masses in the range 600–900 GeV. Here, little difference is seen between the different approaches, though if anything the multiplicity-based approach might offer better rate vs. efficiency in some cases (especially when a lower rate is required).

For high luminosity the story is similar: in all cases a multiplicity-based estimator is either better, or at least no worse, than a full total jet E_T (Table 4-16). This remains true for other sets of jet-E_T thresholds considered for the multiplicity-based algorithm. Further study is needed,

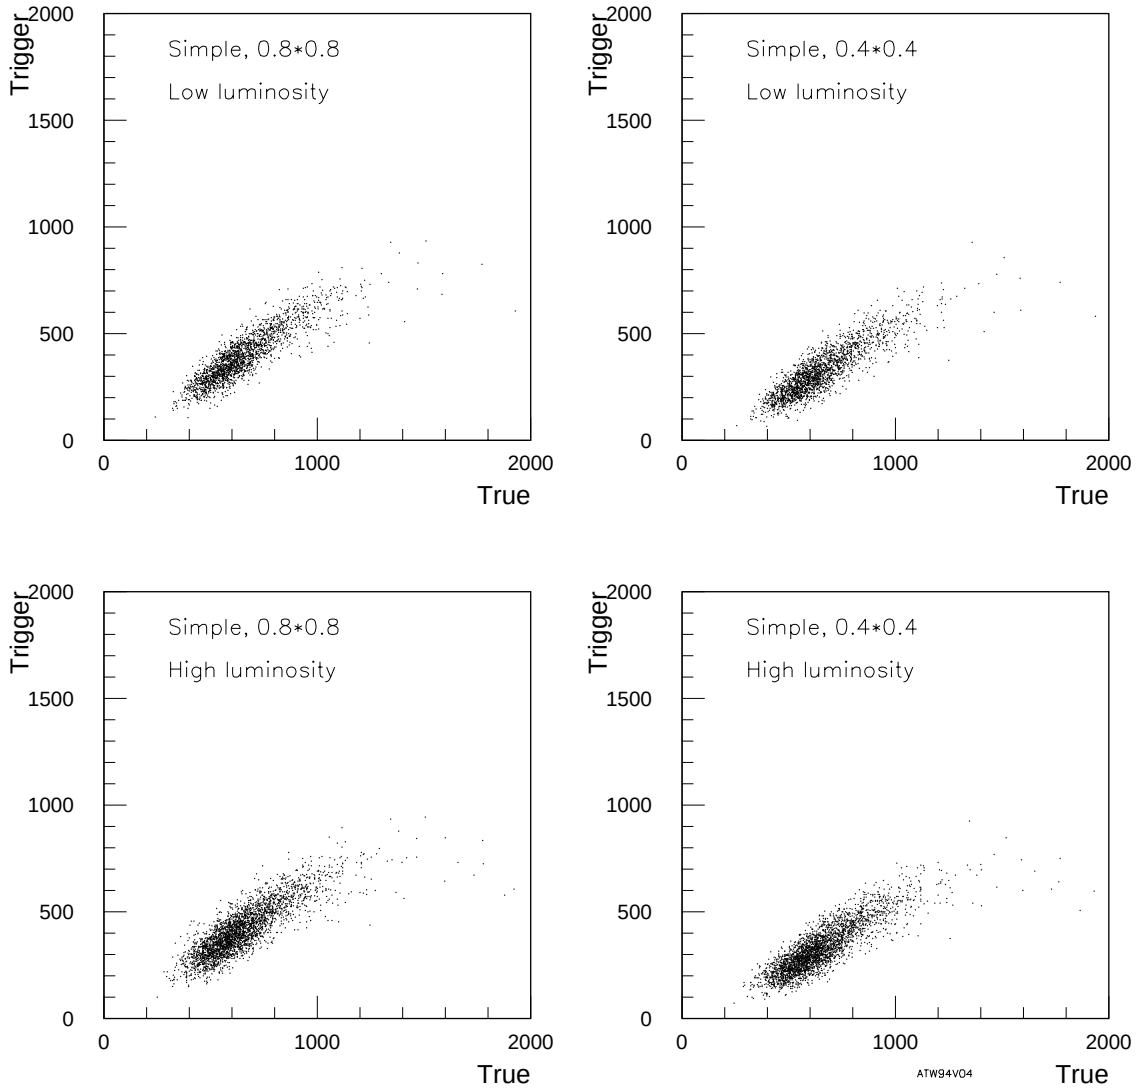

Figure 4-54 Multiplicity-based total jet-E_T trigger vs. total event E_T, for jets of 0.4×0.4 and 0.8×0.8, at low luminosity (top) and high luminosity (bottom).

Table 4-15 Physics efficiencies of 'full' and multiplicity-based total jet-E_T triggers for different trigger rates, at low luminosity. The lowest jet-E_T threshold was 5 GeV in all cases.

	Top events			SUSY events		
Rate	**Full algo. 0.8×0.8**	**Full algo. 0.4×0.4**	**Multiplicity algorithm**	**Full algo. 0.8×0.8**	**Full algo. 0.4×0.4**	**Multiplicity algorithm**
1 kHz	70%	68%	66%	89%	88%	87%
500 Hz	48%	48%	51%	86%	86%	86%
250 Hz	31%	31%	38%	85%	85%	85%
100 Hz	17%	16%	22%	80%	81%	81%

however, to be sure that such an algorithm is compatible with all other likely requirements on the jet trigger without requiring the addition of extra jet thresholds.

Table 4-16 Physics efficiencies of 'full' and multiplicity-based total jet-E_T triggers for different trigger rates, at high luminosity. The lowest jet-E_T threshold was 10 GeV in all cases.

	Top events			SUSY events		
Rate	**Full algo. 0.8 × 0.8**	**Full algo. 0.4 × 0.4**	**Multiplicity algorithm**	**Full algo. 0.8 × 0.8**	**Full algo. 0.4 × 0.4**	**Multiplicity algorithm**
1 kHz	19%	19%	21%	81%	82%	80%
500 Hz	11%	11%	15%	74%	75%	75%
250 Hz	7%	6%	9%	66%	67%	67%
100 Hz	3%	3%	4%	54%	57%	55%

4.6.2.4 Conclusions and preferred solution

Other studies (see Section 4.5) indicate that good performance can be obtained from a conventional total-E_T trigger, even at LHC design luminosity. Nonetheless, a total jet-E_T trigger has different sensitivities, and is a sensible complementary trigger if costs of implementation are not excessive. From the studies to date, it would seem that a solution based on the multiplicities of jets passing different thresholds can provide as good a physics performance as a full total jet-E_T sum without significantly constraining the operation of the jet trigger, and thus with a very low impact and cost of implementation. This then seems an attractive solution for such a trigger.

4.7 References

4-1 *ATLAS Trigger Performance Status Report.* CERN / LHCC / 98–15, June 1998.

4-2 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, 1994

4-3 *ATLAS Letter of Intent*, CERN / LHCC / 92–4, 1992.

4-4 *CMS Technical Proposal*, CERN / LHCC / 94–38, 1994.

4-5 S.Catani et al., Nucl. Phys. B406 (1993) 187.

4-6 E. Richter-Was et al., *ATLFAST: a package for particle-level analysis,* ATLAS note PHYS–NO–79, March 1996.

4-7 T. Sjöstrand, *PYTHIA 5.7 and JETSET 7.4 Physics and Manual*, CERN–TH.7112 / 93, 1993.

4-8 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN / LHCC / 96–41, December 1996

4-9 *ATLAS Tile Calorimeter Technical Design Report*, CERN / LHCC / 96–42, December 1996.

4-10 R. Dubitzky et al., *Level-1 Rates for Triggers Using the $E_T$miss Signature.* ATLAS note DAQ–98–011, August 1998.

4-11 F. Gianotti, *Precision SUSY measurements with ATLAS: SUGRA "point 4",* ATLAS note PHYS–NO–110, October 1997.

5 Calorimeter signals and bunch-crossing identification

5.1 Introduction

In the first part of this chapter (Section 5.2) we describe how the analogue trigger-tower signals from the various calorimeters are processed before they reach the level-1 calorimeter trigger system. This includes the generation and transmission of the signals from the detector to the trigger counting room (USA15), and the organization of signals for cabling to the trigger Preprocessor. The Level-1 Calorimeter Trigger makes strong demands on the addition of calorimeter cells to form trigger towers, and on the properties and handling of the resulting trigger tower signals. One of the key remaining questions is how to form and use trigger towers at the boundaries between the different types of calorimeters used in ATLAS. Simulation studies of how various options affect the performance of the trigger have therefore been carried out, and are described below in Sections 5.2.2.3, 5.2.3.2, and 5.2.3.4.

The second part of the chapter (Section 5.3) is devoted to bunch-crossing identification (BCID), which is crucial to the successful and efficient operation of the trigger. We must be able to associate broad signals uniquely with LHC bunch-crossings to a high degree of accuracy. At the same time, we must extract the transverse energy of those signals correctly, and in such a way as to minimize the effects of noise and pile-up background. An important issue is BCID for saturated pulses, and current ideas for dealing with these will be described.

5.2 Calorimeter signals

5.2.1 Level-1 signal requirements

The Level-1 Calorimeter Trigger has three main functional components: the Preprocessor, the Cluster Processor and the Jet/Energy-sum Processor. The Preprocessor digitizes the signals, performs BCID, and converts the trigger-tower signals into E_T values. It also constructs the $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ elements needed by the jet and energy-sum triggers. The Cluster Processor searches for electron, photon, and hadron candidates. The Jet/Energy-sum Processor looks for jet candidates, and calculates the missing and total transverse energy for each bunch-crossing.

Physically, these processing systems are composed of printed-circuit modules mounted in crates. To keep the system interconnections conceptually simple, there will ideally be a clear correspondence between modules in the Preprocessor and modules in the Cluster and Jet/Energy-sum Processor in terms of the region of calorimeter space they handle. The baseline design of the Preprocessor uses 64-channel modules. To match the preferred 'ϕ-quadrant' architecture of the Cluster and Jet/Energy-sum processors (see Section 6.1), each PPM will handle signals from a $\Delta\eta \times \Delta\phi = 0.4 \times 1.6$ region of the calorimeters. Hence the processor architecture constrains the allocation of analogue trigger towers to PPMs. There is no need to mix signals from the electromagnetic and hadronic calorimeters on the same PPM.

There are several regions of rapidity in ATLAS where signals from different calorimeters overlap. These include the boundaries of the liquid-argon (LAr) electromagnetic (e.m.) barrel and endcap (see Section 5.2.2.3), the scintillating-tile hadronic barrel and extended-barrel (see Section 5.2.3.2), and the Tile extended-barrel and LAr hadronic endcap (see Section 5.2.3.4). Figure 5-1 shows a simplified cross-section of one quadrant of the calorimeters. Ideally, the relatively small number of signals in these overlapping regions would be summed across the boundaries to form projective trigger towers. The summation could be done either using the analogue signals before they reach the PPMs, or using the post-BCID digital signals on the PPMs. The first option is difficult, since signals from different calorimeters have different timing characteristics, and they may also saturate at different transverse energies. On the other hand, summing signals on the PPMs would require special versions of PPMs, which is undesirable and would add latency. Therefore, compromises must be found, and these are discussed below. These compromises at level-1 may require extra processing at level-2 in the transition regions.

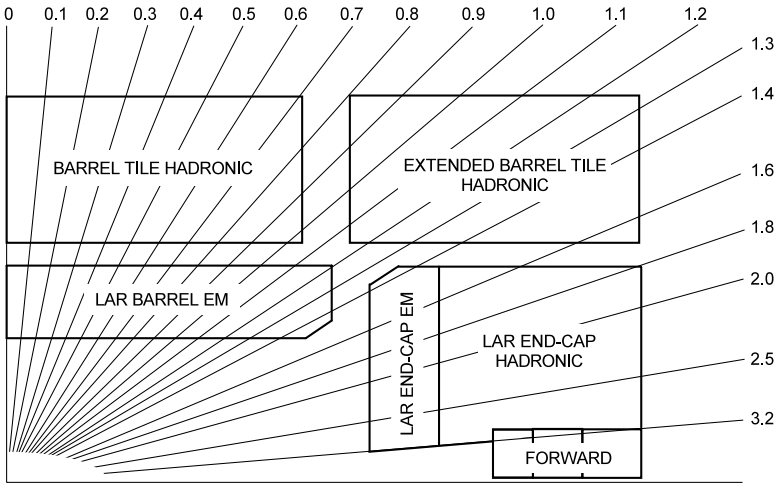

Figure 5-1 Layout of calorimeter boundary regions.

The analogue signal processing of the calorimeter signals and the formation of trigger towers is described in detail in the calorimeter *Technical Design Reports* [5-1][5-2]. The important characteristics of these signals from the point of view of the level-1 trigger are described in the *Level-1 Calorimeter Trigger User Requirements Document* [5-3]. A brief summary is presented here.

5.2.1.1 Cables and line receivers

All calorimeter signals will be transmitted to USA15 using individually-shielded multiway twisted-pair cables. The number of channels per cable and their organization in η–ϕ space has not yet been decided for all of the calorimeters. Where these details are known at present, the grouping of signals on cables is not matched to the PPM requirements, and we need to regroup the signals before connecting them to the PPMs. Patch-panels could introduce crosstalk and reduce the overall system reliability, and so should be avoided where possible and carefully-designed where necessary.

The differential signals from the Tile Calorimeter are fed directly to the trigger Preprocessor Modules, but those from the LAr calorimeters are first fed to separate ‘Receiver Stations’. This gives the LAr calorimeter group access to the analogue waveform, which is an important

diagnostic facility. The outputs of these receiver modules are differential signals. The inputs of the trigger Preprocessor Modules are active line-receivers, which terminate the cables with the appropriate impedance and give some common-mode rejection to reduce the effects of coherent noise.

The boundary of responsibility between the calorimeter groups and the trigger group is at the output connectors of the Receiver Stations for all liquid-argon calorimeters, and at the ends of the long cables to USA15 in the case of the Tile Calorimeter.

5.2.1.2 Conversion to transverse energy

The energy signals from the calorimeters must be converted to transverse energy, within a ±10% band. This is done in the Receiver Stations for all liquid-argon calorimeters, and is therefore the responsibility of the liquid-argon group. For the Tile Calorimeter, the conversion is done in the Preprocessor line receivers, and so is the responsibility of the calorimeter trigger group.

5.2.1.3 Incoherent noise

There are two main sources of incoherent noise —the thermal noise of the amplifiers and cables, and the pile-up noise. In terms of transverse energy, the latter is more or less constant with rapidity but varies with sampling depth and machine luminosity. Obviously the noise associated with the trigger-tower signals should be as small as possible, and noise from unexpected sources avoided by careful design and installation.

From the point of view of the trigger system the incoherent noise affects a number of key areas. Firstly it reduces the efficiency of the BCID algorithm for transverse energies of less than about 3 GeV. This affects all types of trigger. High levels of noise will also reduce the effectiveness of the isolation cuts applied to candidate electromagnetic showers. Finally, noise forces up the minimum threshold for jet triggers in order to avoid a high rate of fake jets. This is due to the large number of trigger towers that are added together.

In the LAr calorimeters the pile-up noise at a luminosity of 10^{34} $cm^{-2}s^{-1}$ is expected to be about 400 MeV per trigger tower. This is non-Gaussian. The thermal noise will be 410 MeV at $\eta = 0$ and 217 MeV at $\eta = 2$, giving total noise levels at these rapidities of 570 MeV and 450 MeV respectively. In the scintillating-tile calorimeter the pile-up noise at a luminosity of 10^{34} $cm^{-2}s^{-1}$ is expected to be about 90 MeV per trigger tower. The thermal noise will be about 35 MeV independent of rapidity, giving a total noise level of about 97 MeV. Assuming an ADC least-significant bit of about 0.25 GeV (see Section 6.2.1.1), the quantization noise is about 75 MeV. Hence the noise of the LAr calorimeters will dominate over the quantization noise, while the noise of the scintillating-tile calorimeters will be comparable to the ADC noise.

5.2.1.4 Coherent noise

Coherent noise due to electromagnetic interference and ground-loop currents is a potentially very serious problem for the trigger processor. Perhaps most worrying is the fact that this form of noise cannot be quantified until ATLAS is completely installed, by which time it may be too late to correct. Given that the total current supplied to the LAr calorimeter readout may be as much as 20 kA, the problem could be severe. The use of analogue-optical signal transmission between the experiment and the trigger processor would largely eliminate this problem. Prototype links have been tested and found to have adequate dynamic range and linearity, but

poor channel-to-channel reproducibility and temperature stability [5-4]. These effects could be reduced by frequent calibration. Radiation damage would need to be investigated further before this solution could be used. The extra cost of such a system would also have to be studied carefully.

5.2.1.5 Crosstalk

Crosstalk between the analogue calorimeter signals could have severe consequences for the performance of the trigger system, particularly in the electron / photon and hadron / tau triggers where large genuine signals occur next to trigger towers with very low isolation-veto thresholds. It is hard to quantify exactly what is an acceptable level of crosstalk, but less than 1% should certainly be achievable and is probably acceptable. The major source of crosstalk is expected to be at the front-end, but it can also occur on the cables between the detector and trigger processor, and on any patch-panels needed to achieve the appropriate grouping of signals on multiway cables. Note that BCID and timing can help to reject crosstalk, much of which is out-of-time since it is expected to be capacitively coupled.

5.2.1.6 Saturation

The BCID system (Section 5.3) has to determine the timing and amplitude of the analogue pulses from the calorimeters. It uses one of two algorithms depending on whether or not the pulse being processed causes saturation in the system. The BCID system puts severe constraints on the performance of the analogue processing chain used to produce the trigger towers.

For non-saturated pulses, a matched finite-impulse response filter (FIR) is used. The filter coefficients are programmable and so the filter can be tuned to a variety of pulse shapes. However, if the shape or timing of the pulse from a given trigger tower varies significantly from event to event then the filter coefficients will no longer be ‘matched’ and the BCID efficiency can fall. This implies that all of the calorimeter cells which are combined to make a trigger tower must have the same pulse shape and must be synchronized in time before being summed. Some residual variation in pulse shape is unavoidable because of pile-up.

The dynamic range of trigger-tower signals is at least 256 GeV in E_T, up to which point they should be reasonably linear. For signals which do cause saturation, the pulse shape will inevitably change, and the matched filter approach fails. For these pulses a system based on a lookup table is used. The nature of the saturation can be either digital or analogue. In digital saturation, the signal exceeds the full scale ADC voltage (about 2.5 V) and is simply truncated. In analogue saturation, the amplifiers which supply the trigger-tower signals saturate. Clearly, the analogue saturation must not occur at a lower signal level than the digital saturation.

Some studies of saturation energies have been carried out for the LAr electromagnetic calorimeters. A trigger tower is built from up to 60 calorimeter cells (see Section 5.2.2.1 below). The signals from each cell are fed through a number of amplifiers, shapers, buffers, variable delay lines, and summing circuits before reaching the trigger Preprocessor. Figure 5-2 is a simplified schematic showing the key summing stages. Using information from a number of sources, the saturation energy per cell as a function of η for the various calorimeter depth samplings has been calculated. The results are shown in Figure 5-3. Part (a) shows the saturation transverse energy in the preamplifiers for the presampler and the front, middle and back layers. A horizontal line has been drawn at E_T = 256 GeV. Part (b) shows the saturation energy in the shapers, (c) the layer sums, and (d) the tower builders.

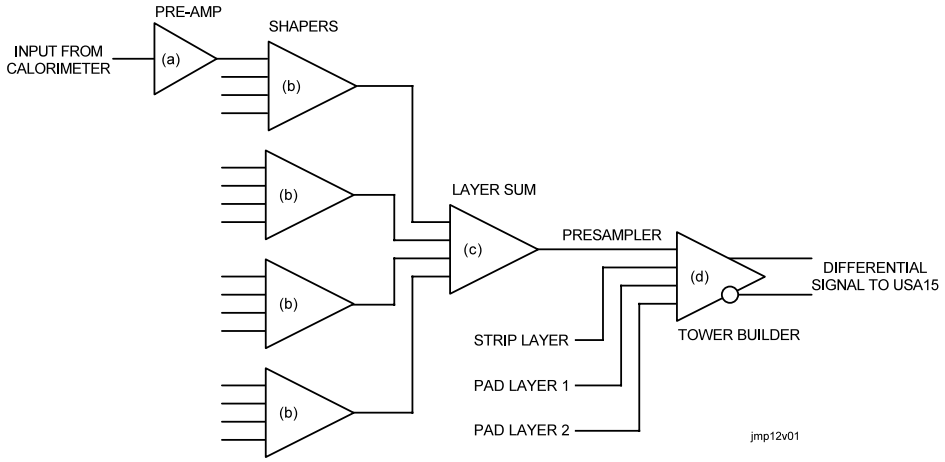

Figure 5-2 Simplified diagram showing the key summing stages in the LAr trigger-tower processing.

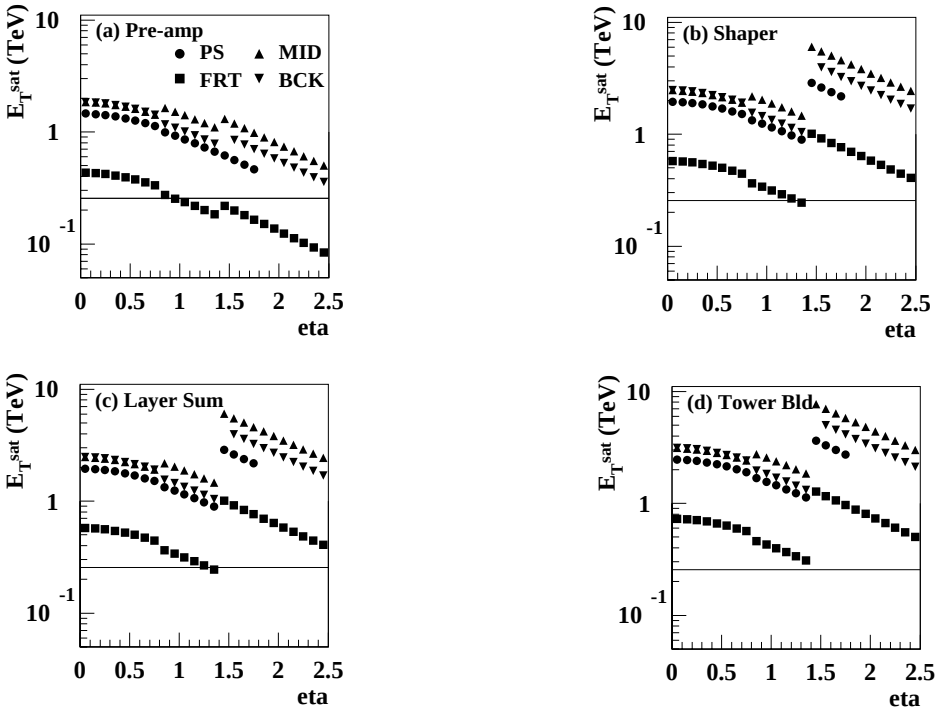

Figure 5-3 Saturation energy per cell as a function of η for various LAr calorimeter depth samplings, at different stages in the summation chain.

The front layer has the lowest saturation energy for all values of η, but it also has the highest granularity and thus on average will see smaller energy deposits per cell. The middle layer is expected to see the largest energy deposits on average because of its thickness. From these plots it can be seen that, since the trigger-summing electronics has a guaranteed dynamic range of 256 GeV, saturation is an issue that must be very carefully studied. This is discussed further in section 5.3.2.

In the Tile Calorimeter, it is believed that very big pulses will be rare and saturation is much less of a problem. In addition, the much smaller number of cells added to form trigger towers reduces the possible variation in different modes of saturation along the electronics chain.

5.2.1.7 Pulse timing and shape stability

Fluctuations in the timing or shape of the pulses will reduce the efficiency of the BCID algorithm since they mean that the filter coefficients will not match the pulse. A trigger-tower signal should have less than ±3 ns timing jitter, and the individual summands for each trigger tower must be aligned to within ±2.5 ns before being summed. The peaking-time should not drift by more than 1 ns as a result of temperature variations and radiation damage. Simulation results in support of these required timing values are presented below in section 5.3.1.4.

For signals that cause saturation in the analogue signal-processing chain which constructs a trigger tower, the rise and fall times of the pulse must be preserved as well as possible. This means that amplifiers which exhibit 'latch-up' behaviour under saturation must be avoided.

The effect of the long twisted-pair cables on saturated pulses is also a concern. The frequency-dependent attenuation of the cable alters the shape of the pulse. This could also be avoided by using analogue optical transmission, but this can have other problems.

5.2.2 The electromagnetic calorimeters

The granularity of the trigger towers in the LAr electromagnetic calorimeter varies with rapidity, and there is one region of overlap between the barrel and endcap. Figure 5-4 shows the acceptance in η of each calorimeter, for a section of $\Delta\phi = 0.2$.

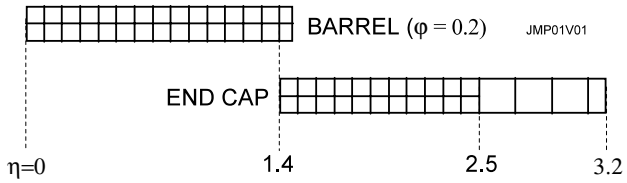

Figure 5-4 Coverage in η of LAr e.m. barrel and endcap calorimeters.

5.2.2.1 Barrel e.m. calorimeter trigger-tower formation

The barrel electromagnetic calorimeter covers the rapidity range $|\eta| < 1.5$ in two halves split at $\eta = 0$. It is divided into four depth segments: the preshower region, the η-strip front layer and then two pad layers. The total number of calorimeter cells that must be summed to form a trigger tower varies with depth and with η as shown in Table 5-1.

Table 5-1 Number and layout of calorimeter cells forming trigger towers in the barrel e.m. calorimeter.

Rapidity range	Presampler	Strip layer	Pad layer 1	Pad layer 2	Total
$\|\eta\| < 1.3$	4	32	16	8	**60**
$1.3 < \|\eta\| < 1.4$	4	32	16	4	**56**
$1.4 < \|\eta\| < 1.5$	4	12	4	0	**21**

The trigger-tower granularity of the barrel electromagnetic calorimeter is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. Trigger towers are formed in three stages: linear-mixer sums, layer sums, and finally tower-builder sums.

Tower-builder cards can produce up to 32 trigger-tower signals. The outputs are differential signals which are fed to a receiver station in USA15. Each receiver-station module (RXM) handles 64 trigger towers and has dedicated outputs to feed signals to the PPMs.

Conversion of the calorimeter signal into a transverse energy is done in two stages. A rough correction is done at the calorimeter front-end, where the gain of the linear mixers can be switched between 1 and 3.5 depending on the rapidity of the channel. A more accurate programmable adjustment is made on the RXM by tuning the gain of the final stage of amplification.

The structure of the calorimeter constrains the allocation of calorimeter signals to tower-builder cards. The proposed arrangement is shown in Figure 5-5 for one quadrant of the barrel electromagnetic calorimeter. Each tower-builder board handles 30 towers from a region of $\Delta\eta \times \Delta\phi = 1.5 \times 0.2$ (illustrated by the heavy lines). Fifteen trigger towers corresponding to $\Delta\eta \times \Delta\phi = 1.5 \times 0.1$ are grouped on a single 16-way cable for transmission to USA15.

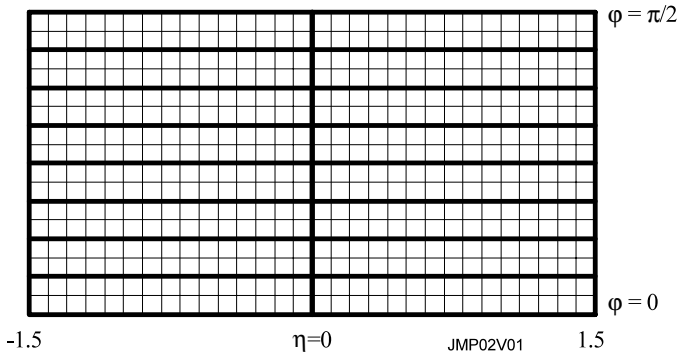

Figure 5-5 Arrangement of tower-builder cards for the LAr e.m. barrel calorimeter.

5.2.2.2 Endcap e.m. calorimeter trigger-tower formation

The endcap electromagnetic calorimeters cover the rapidity ranges $1.4 < |\eta| < 3.2$. Each is divided into up to four depth segments: the preshower region, the η-strip first layer, and then two pad layers. The endcaps overlap with the barrel electromagnetic calorimeter in the regions $1.4 < |\eta| < 1.5$. The number of cells that are summed to make a trigger tower is summarized in Table 5-2.

Table 5-2 Number and layout of calorimeter cells forming trigger towers in the endcap e.m. calorimeter.

Rapidity range	Presampler	Strip layer	Pad layer 1	Pad layer 2	Total
$1.4 < \|\eta\| < 1.5$	0	4	16	0	**20**
$1.5 < \|\eta\| < 1.8$	4	32	16	8	**60**
$1.8 < \|\eta\| < 2.0$	0	24	16	8	**48**
$2.0 < \|\eta\| < 2.4$	0	16	16	8	**40**
$2.4 < \|\eta\| < 2.5$	0	4	16	8	**28**
$2.5 < \|\eta\| < 3.1$	0	4	4	0	**8**
$3.1 < \|\eta\| < 3.2$	0	2	2	0	**4**

The granularity of the trigger towers in the endcaps is not uniform. For $1.4 < |\eta| < 2.5$ the granularity is 0.1×0.1 while for $2.5 < |\eta| < 3.1$ it is 0.2×0.2 (to match the hadronic endcap), and for $3.1 < |\eta| < 3.2$ it is 0.1×0.2. As in the barrel e.m. calorimeter, the trigger towers are formed in three stages: linear mixer, layer sum, and then tower-builder sum. Figure 5-6 shows how the trigger towers are allocated to tower-builder cards for one quadrant of an endcap. As with the e.m. barrel calorimeter, the outputs of the tower-builder cards are sent first to a receiver station in USA15, where they are converted to E_T.

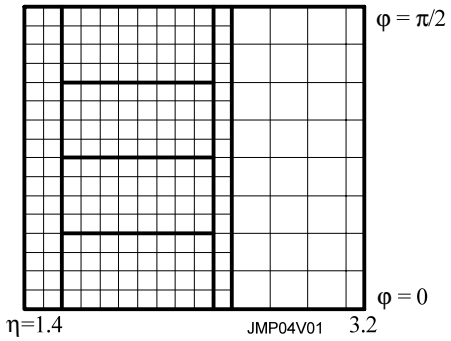

Figure 5-6 Arrangement of tower-builder cards for the LAr e.m. endcap calorimeter.

5.2.2.3 Treatment of the e.m. barrel/endcap boundary

In this e.m. boundary region the analogue addition of trigger towers upstream of the trigger Preprocessor is feasible; one possibility is to do this in the LAr Receiver Stations. The overlap between the e.m. barrel and endcap occurs within the trigger tower covering $1.4 < |\eta| < 1.5$. The on-detector electronics will form two trigger towers in this region: a 'thin' tower from the end of the barrel calorimeter, and a full-depth tower from the endcap calorimeter. Since the trigger electronics require a single tower from this region, there are then three options possible:

1. The thin end-of-barrel tower could be excluded from the trigger.
2. The end-of-barrel signals could be added (off-detector) to those from the last full tower in the barrel ($1.3 < |\eta| < 1.4$).
3. The end-of-barrel signals could be added to the first endcap tower (with which they overlap in space).

Clearly the third option would be closest to the optimum, in that it forms a single, projective tower across the transition, and equally the first would be simplest to implement. Both electromagnetic cluster and isolation performance are sensitive to the treatment of this transition. The effect on cluster performance is illustrated in Figure 5-7. As expected, a clear loss in signal is seen if the partial tower is simply excluded from the trigger. A smaller, but still visible degradation is seen when the partial tower is added to the last full barrel tower, rather than to the geometrically-correct tower. In both cases, there might be some improvement obtainable by re-calibration, but there would be a loss of resolution, and perhaps of linearity, in either case.

Since excluding signals from the trigger would also have similar effects on the hadron / tau, jet, missing-E_T and total-E_T triggers, this is not considered further.

Because isolation relies on signals which are small compared with the cluster E_T (typically a few GeV), isolation performance is potentially even more sensitive to the handling of the transitions

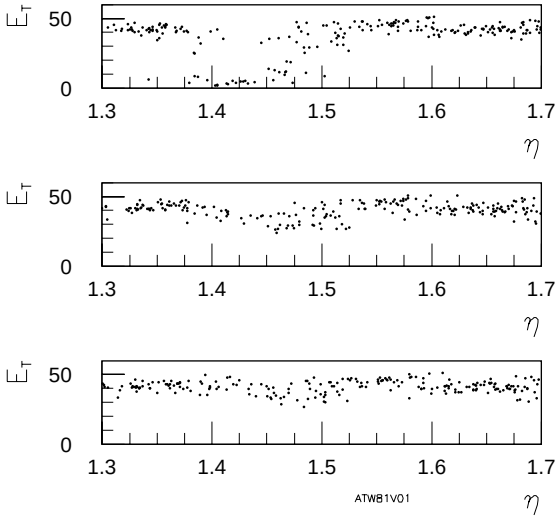

Figure 5-7 Cluster E_T vs. electron η, for different options for the e.m. barrel/endcap transition. The top plot shows the effect of excluding the end of the barrel calorimeter from the trigger. In the middle plot the signal from the end of the barrel is added to the last barrel tower, while in the bottom plot the signals are summed across the transition.

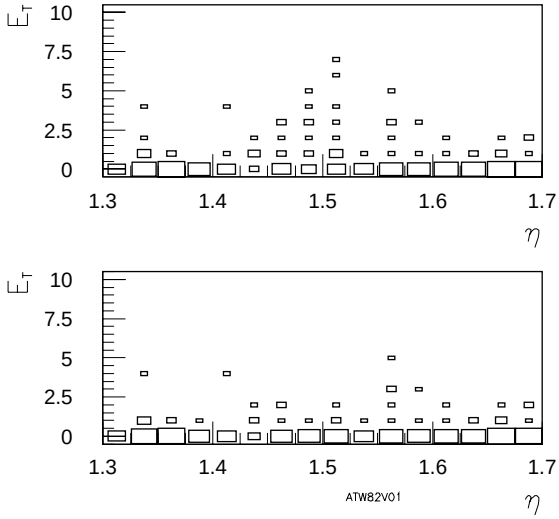

Figure 5-8 Electromagnetic isolation E_T vs. electron η, for 40 GeV electrons. The top plot shows the effect of adding the end-of-barrel signals to the tower covering η = 1.3–1.4, while in the bottom plot these signals are added to the tower covering 1.4–1.5, i.e. summed across the transition.

than the trigger clusters. Figure 5-8 compares the electromagnetic isolation sums for 40 GeV electrons near the transition, for the two different ways of adding in the end-of-barrel tower. Adding to the 'wrong' tower in space (option 2) results in a tail in the isolation E_T for electrons 'clipping' the end of the barrel. This would result in a reduction in efficiency, and re-calibrating to mitigate the loss of E_T in the trigger clusters could only make this worse.

In the case of the e.m. barrel/endcap transition, it is then seen that summing across the boundary between calorimeter modules would be the clearly preferred solution.

5.2.3 The hadronic calorimeters

The granularity of the trigger towers in the hadronic calorimeters varies with rapidity and sampling layer. There are two regions of overlap: between the Tile Calorimeter barrel and extended barrel, and between the extended barrel and the LAr hadronic endcap calorimeter. Figure 5-9 shows the acceptance in η of each calorimeter and the granularities of the various sampling layers for a section $\Delta\phi = 0.2$. The granularity of the trigger towers in the hadronic calorimeters is essentially the same as that of the readout cells (the main exception being the Tile Calorimeter tail-catcher).

5.2.3.1 Tile Calorimeter trigger-tower formation

The barrel and extended-barrel scintillating-tile hadronic calorimeters are divided into three depth samples. The first and second depth samples have a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. The third depth sample has a coarser granularity of $\Delta\eta \times \Delta\phi = 0.2 \times 0.1$. The trigger towers formed in the Tile Calorimeter are therefore a complex shape. The barrel calorimeter covers

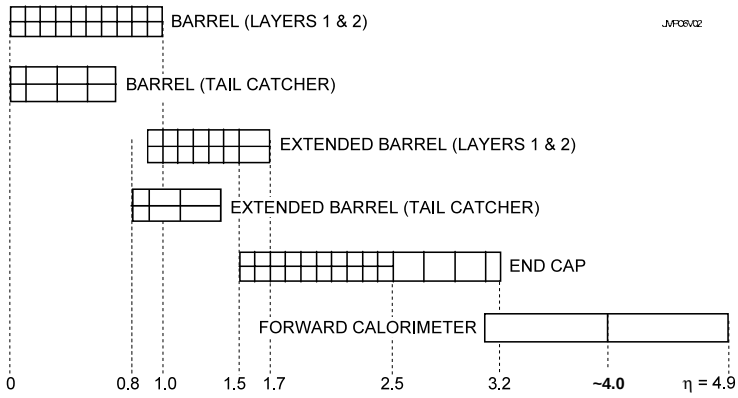

Figure 5-9 Coverage in η of Tile Calorimeter, LAr hadronic endcap calorimeter, and forward calorimeter.

$|\eta| < 1.0$ and the extended barrel covers $0.8 < |\eta| < 1.7$, so there is a substantial overlap region.

Scintillation light from each tile is fed to two photomultiplier tubes (PMTs) mounted in drawers at the back of the calorimeter. The readout electronics are also mounted in these drawers. The PMT signals are amplified and shaped to give a unipolar pulse with a full width at half maximum of 50 ns. A linear output from each shaper is fed to a tower-builder card which is also mounted in the drawer. The number of PMT signals which the tower builder must sum depends on the rapidity. Table 5-3 details the number of PMTs which are summed to make trigger towers in the barrel, and Table 5-4 gives the same information for the extended barrel. Note that no summation of signals between barrel and extended barrel takes place on the detector.

Table 5-3 Number and layout of PMTs forming trigger towers in the barrel hadronic calorimeter.

Rapidity range	Layer 1	Layer 2	Layer 3	Total
$\|\eta\| < 0.7$	2	2	1	**5**
$0.7 < \|\eta\| < 1.0$	2	2	0	**4**

Table 5-4 Number and layout of PMTs forming trigger towers in the extended-barrel hadronic calorimeter.

Rapidity range	Layer 1	Layer 2	Layer 3	Total
$0.8 < \|\eta\| < 1.1$	2	2	0	**4**
$1.1 < \|\eta\| < 1.3$	2	2	1	**5**
$1.3 < \|\eta\| < 1.6$	0	2	1	**3**

The signals from each layer are synchronized to within 2 ns before being summed by the tower-builder cards. Each tower-builder card is responsible for a single trigger tower.

The outputs of the tower-builder cards are connected to patch panels at the outer end of the electronics drawers via shielded twisted-pair cables. From the patch panel the signals will be fed to USA15, but the details of how signals are allocated to cables are not yet specified. Since the trigger towers are formed in the electronics drawers, which are long in η, it is reasonable to assume that the signals will be grouped on multiway cables in η strips. Conversion to E_T will be done in the line receivers on the Preprocessor Modules.

5.2.3.2 Treatment of the barrel/extended-barrel boundary

The Tile barrel / extended-barrel transition spans the region $0.8 < |\eta| < 1.0$ (see Figure 5-1), thus affecting two trigger towers in rapidity. At this boundary the only possibility for adding signals, which might be out of time with each other, is to add them digitally on the Preprocessor Modules. This would require special one-off modules and an increase in latency. Again, there are three ways in which this transition could be handled:

1. Partial towers could be formed in both calorimeter modules, each covering part of the depth of the calorimeter. These could then be summed off-detector. As noted, the signals would be out-of-time with each other, which would complicate the summation.
2. 'Enlarged' towers could be formed in both barrel and extended-barrel modules, each spanning the entire transition (0.8–1.0 in η). To the trigger, the barrel signal would appear to be the tower covering 0.8–0.9 and the extended-barrel signal to cover 0.9–1.0.
3. Partial towers could be formed in each calorimeter module, as in 1 above, but instead of adding the signals, only one or other partial tower would be used in the trigger. For the region 0.8–0.9 this would clearly be the tower formed from the Tile barrel calorimeter. For the region 0.9–1.0 one could either use the tower from the barrel calorimeter (only 1 cell deep) or from the extended barrel (physically deeper, but behind a lot of material). Both possibilities are considered here.

While the enlarged towers of the second option might appear unattractive, in fact only the hadron / tau trigger is potentially sensitive to this: since the entire transition lies within a single jet element, these signals are added together in that single element no matter how they are allocated between trigger towers. It should also be noted that the Tile towers are only approximately projective in any event, especially in the extended-barrel, and that the boundaries between trigger towers in the rapidity direction will not be sharp in any case. Excluding part of the depth of the calorimeter (option 3) would be expected to degrade both e.m. isolation and tau and jet clusters. The loss of E_T is visible in Figure 5-10, which shows the hadronic component of tau clusters in the transition region. A significant loss of E_T is seen for clusters centred on the transition if only the barrel sample is used for the tower 0.8–0.9, and only the extended barrel sample for 0.9–1.0. In contrast, the loss is smaller if the barrel sample only is used for both transition towers, though there is still a loss of E_T. This can also be seen in jets, as shown in Figure 5-11.

Forming over-sized towers within the two modules separately (option 2) results in no loss of E_T to the jet trigger, as can be seen from Figure 5-11. However, it does have adverse consequences for tau isolation, as shown in Figure 5-12. In these events, some hadronic E_T is moved from the 'cluster' into the 'isolation' region. This also slightly degrades the cluster E_T measurement, making it more comparable to that when the extended barrel is excluded from the transition region. Since the usefulness of hadronic isolation in the hadron / tau trigger is in any event questionable (see Section 4.3), and since the effect is only significant in relatively high-E_T clusters (which will be less reliant on isolation), the contamination of the hadron / tau trigger's hadronic isolation in this scenario is felt unlikely to have a significant physics impact.

If the extended-barrel signals were instead to be excluded from the transition, it might still be desirable to include these in the missing-E_T trigger. This will need further investigation. However, the loss of E_T in this scenario to both hadron / tau and jet triggers is more likely to damage physics performance than the degradation of the hadron / tau hadronic isolation would.

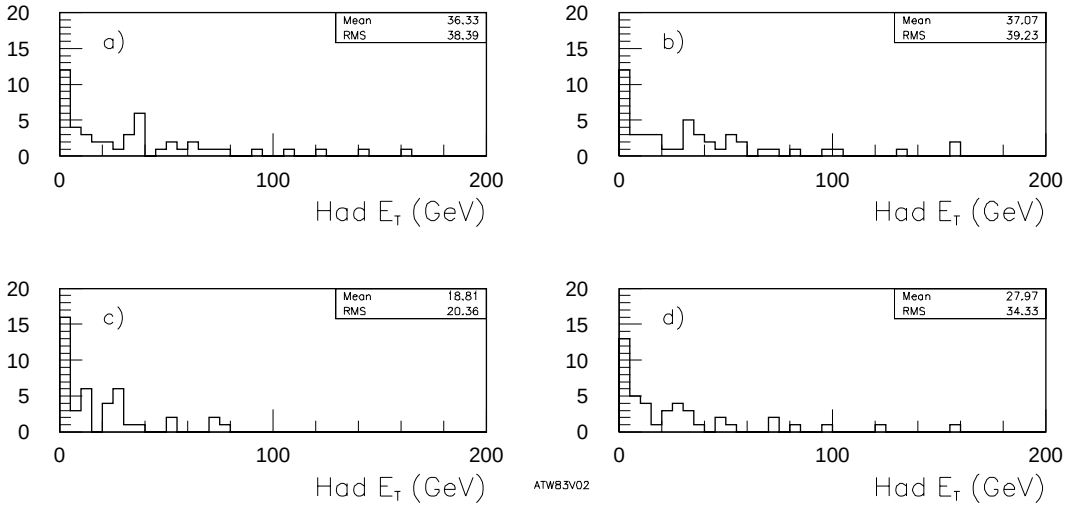

Figure 5-10 Hadronic component of tau cluster, for tau RoIs in the region $0.8 < |\eta| < 1.0$. (a) with signals summed across the transition, (b) with extended towers formed within the modules, (c) using only barrel signals for $|\eta| < 0.9$ and extended barrel for $|\eta| > 0.9$, and (d) using only barrel signals for $|\eta| < 1.0$.

5.2.3.3 Hadronic-endcap calorimeter trigger-tower formation

The hadronic endcap calorimeter (HEC) covers the rapidity range $1.5 < |\eta| < 3.2$. It is divided into four depth segments, of which only the first three are used in the trigger. It overlaps with the Tile extended-barrel hadronic calorimeter in the range $1.5 < |\eta| < 1.7$. The trigger tower dimensions, which are essentially the calorimeter cell dimensions, are $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for $1.5 < |\eta| < 2.5$, 0.2×0.2 for $2.5 < |\eta| < 3.1$, and 0.1×0.2 for $3.1 < |\eta| < 3.2$. Trigger towers are summed on the front-end boards.

Tower Driver boards are used to send 96 trigger towers each to USA15. Figure 5-13 shows how trigger towers might be mapped to these boards. The signals are grouped in ϕ to match the required grouping at the PPMs. As with the other LAr calorimeters, the outputs of the Tower Drivers will go first to a receiver station, where they will be converted to E_T, before transmission to the PPMs.

5.2.3.4 Treatment of the Tile/hadronic-endcap boundary

Across this boundary, shown in Figure 5-1, there is no prospect, even in principle, for analogue summation across the transition — as well as being out-of-time with each other, the signals from the two calorimeters will have quite different pulse shapes. Hence no ‘ideal’ solution exists for this transition. The following options have been considered:

1. Add all Tile signals from the transition region, $1.5 < |\eta| < 1.7$, to the last tower before the transition, thus creating an extended tower containing signals from 1.4–1.7, but appearing to the trigger to cover 1.4–1.5.
2. Exclude all Tile signals from $|\eta| > 1.5$ from the trigger.
3. Add signals within the Tile and HEC calorimeters to form two ‘extended’ towers covering the region $1.5 < |\eta| < 1.7$, with the Tile tower appearing to the trigger as 1.5–1.6 and the HEC tower appearing to the trigger as 1.6–1.7.

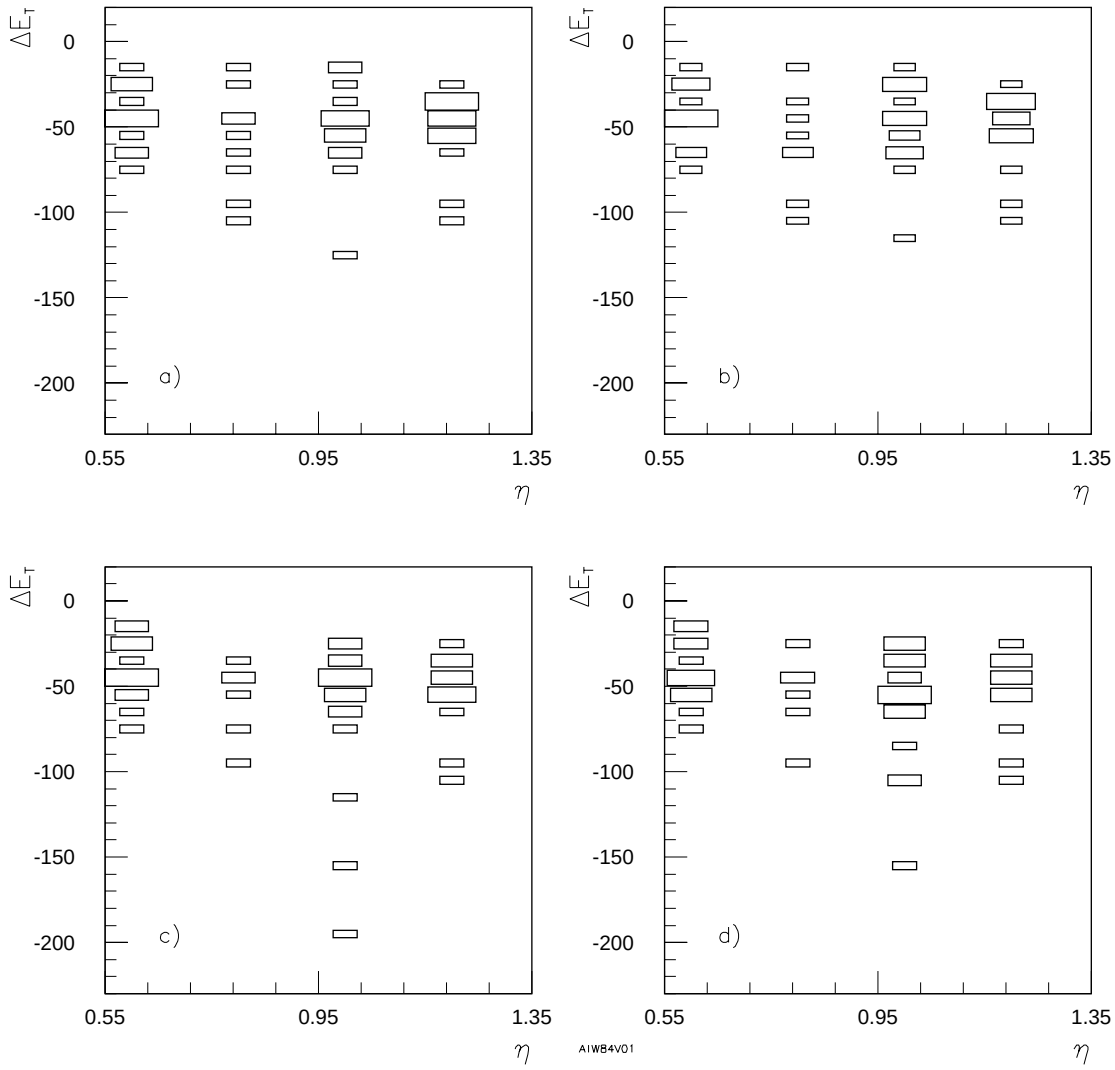

Figure 5-11 Trigger jet E_T (0.4 × 0.4 cluster) – k_T reference jet vs. η of jet. Jets were from decay of 400 GeV Higgs, and all jets had E_T > 50 GeV. The four plots are for the same treatments of the transition as in Figure 5-10.

It might, in principle, be possible to accommodate all of these options with on-detector summing. One unfortunate aspect to this transition is that it is not completely contained within a single jet element. However, for $|\eta| > 1.6$ there is a considerable depth of material in front of the Tile Calorimeter, and so in practice this should not prove problematic for the jet trigger.

For taus, the different options produce different results for RoIs in the range $1.4 < |\eta| < 1.6$. The hadronic component of such clusters is plotted in Figure 5-14 for the three options above. Some loss of E_T is seen if the last Tile cells are excluded. The rapidity dependence of the hadronic response is shown in Figure 5-15, which plots the hadronic component of the cluster against the RoI η coordinate. Note that some of the taus reconstructed at $|\eta| = 1.5$ in option 1 (Figure 5-15a) will be found at $|\eta| = 1.6$ in option 3 (Figure 5-15c), due to the different assignments of the calorimeter cells to trigger towers. Hence depletion of a particular bin does not necessarily mean that taus in this region are lost, though it does indicate a problem with reconstruction.

There is also some loss of jet E_T seen if the end of the Tile Calorimeter is excluded. This is shown in Figure 5-16. Little difference is seen between the two other options, which include all signals,

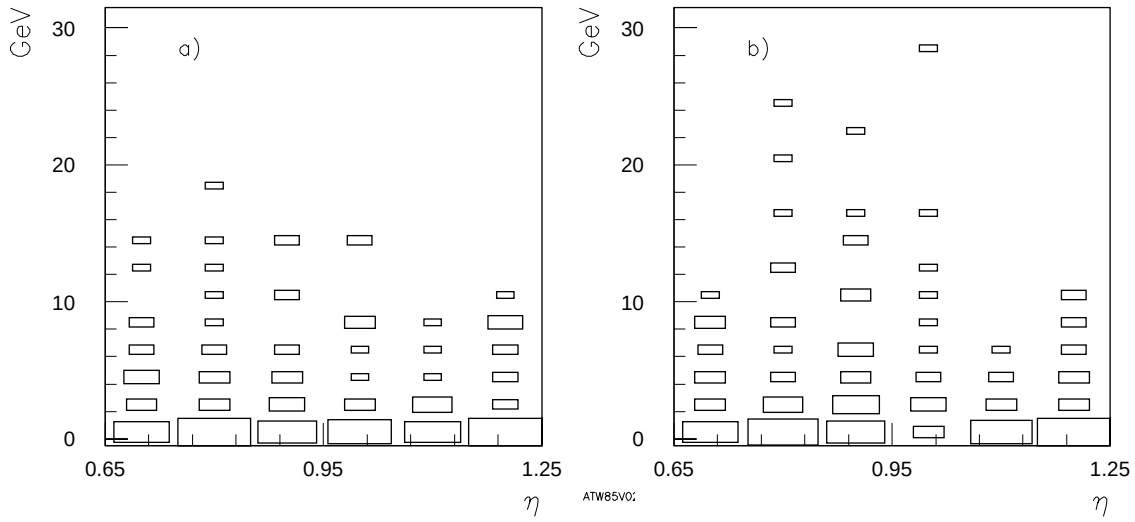

Figure 5-12 Hadronic isolation E_T for tau clusters near Tile barrel/extended-barrel transition. (a) with signals summed across the transition, and (b) with towers formed within the two detector modules, each spanning the transition.

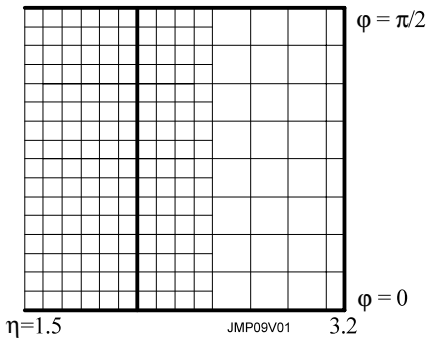

Figure 5-13 Possible arrangement of Tower Driver cards for a quadrant of the LAr hadronic-endcap calorimeter

though in principle option 1 (Figure 5-16a) should be preferable due to better mapping of towers onto jet elements.

The most sensitive quantity is, once again, hadronic isolation in the hadron/tau trigger. The effects of different options are shown in Figure 5-17. As with the Tile barrel/extended-barrel transition, the effect of adding signals to what is geometrically the 'wrong' tower is to produce a tail in the isolation distributions. The alternative tower arrangement of option 3 does seem to suffer less in this respect, and thus would be preferable. As noted previously, however, the hadron/tau hadronic isolation is also arguably one of the less important variables in the trigger algorithms.

In conclusion, there might be reasons to prefer a scheme where the transition between Tile and HEC occurs at $|\eta| = 1.6$ (as seen by the trigger) to a transition at $|\eta| = 1.5$. However, it is unlikely that there would be a major impact on trigger performance, and so this solution is perhaps only worth considering if the appropriate sums can be implemented with little additional cost.

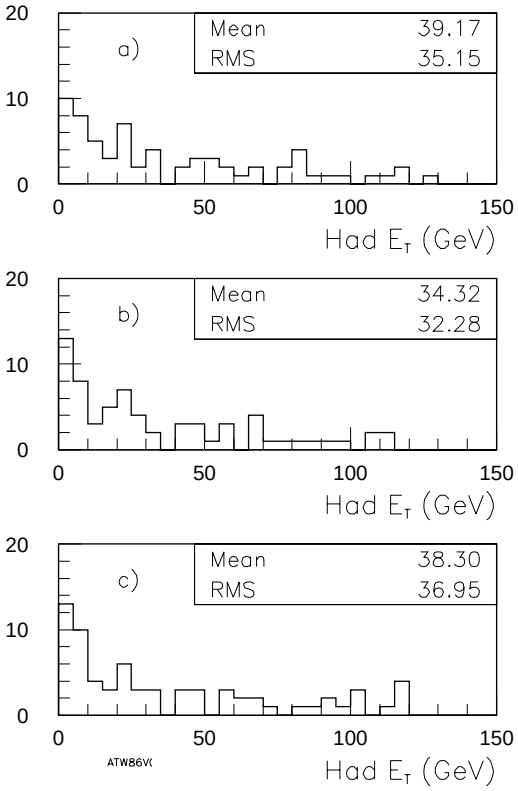

Figure 5-14 Hadronic component of tau cluster, for tau RoIs in the region 1.4 < |η| < 1.6. (a) with all Tile signals for |η| > 1.4 added to one tower, (b) with Tile signals from |η| > 1.5 excluded, and (c) with extended towers in both Tile and HEC covering 1.5 < |η| < 1.7.

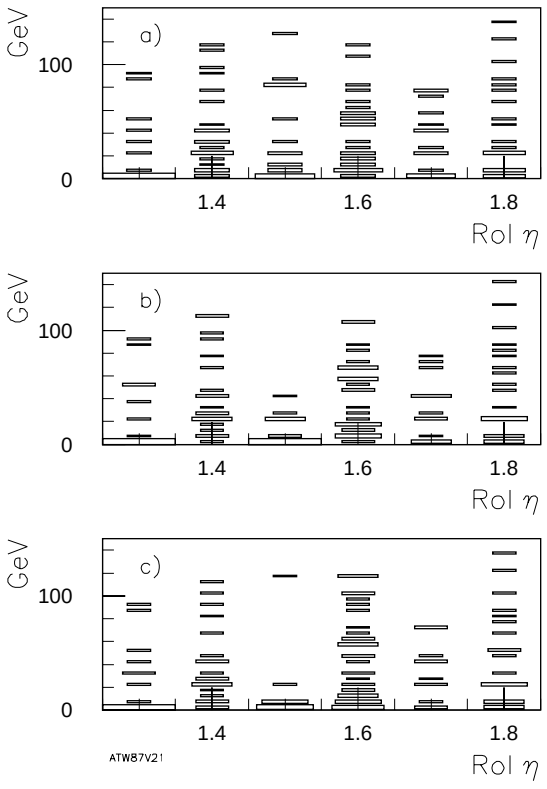

Figure 5-15 Hadronic E_T of tau clusters vs. RoI η. The three plots are for the same treatments of the transition as in Figure 5-14.

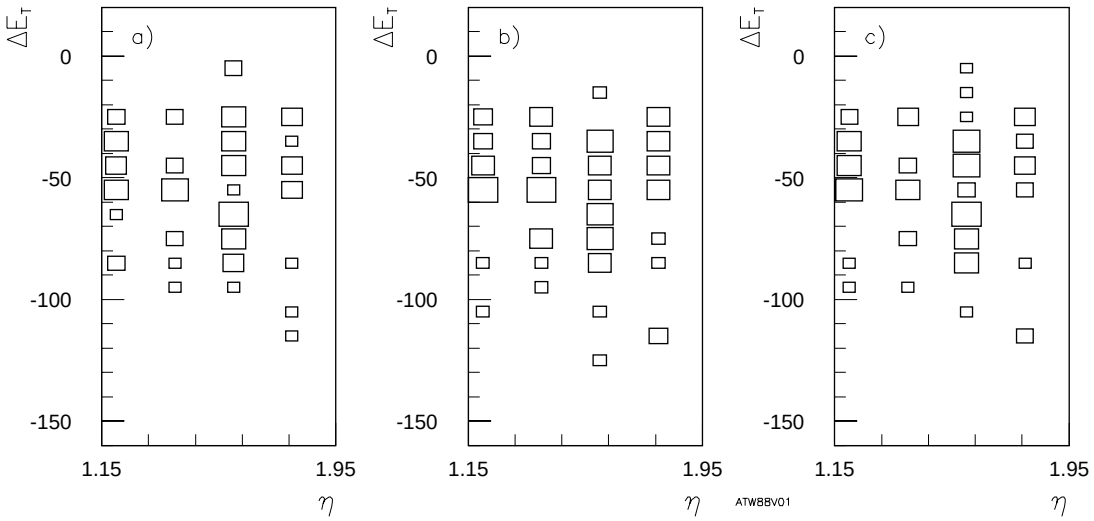

Figure 5-16 Trigger jet E_T – k_T reference jet, vs. η of jet. Jets were from decay of 400 GeV Higgs, and all jets had E_T > 50 GeV. The three plots are for the same treatments of the transition as in Figure 5-14.

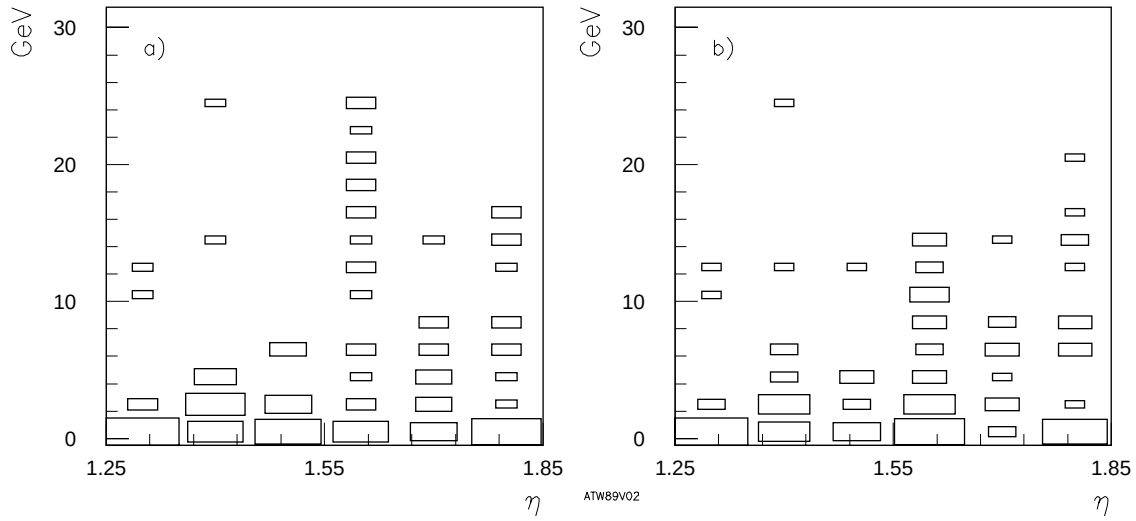

Figure 5-17 Hadronic isolation E_T for tau clusters near the Tile–HEC transition. (a) with all Tile signals for |η| > 1.4 added to one tower, and (b) with extended towers in both Tile and HEC covering 1.5 < |η| < 1.7.

5.2.4 The forward calorimeter

The forward calorimeter covers the rapidity range $3.0 < |\eta| < 4.9$. It is divided into three depth segments and two η ranges. Trigger towers are formed in three stages: preamplifier sums, layer sums, and finally tower-builder sums. The inputs to the layer sums are weighted to give a more accurate value of transverse energy. The allocation of forward calorimeter signals to tower-builder cards is not yet decided. Conversion to E_T will once more be done in the RXM.

5.3 Bunch-crossing identification and energy-extraction algorithms

Bunch-crossing identification (BCID) is an essential requirement for the successful implementation of the level-1 calorimeter trigger. The multiple inelastic pp interactions at each bunch-crossing, every 25 ns, and the high particle-multiplicity of the resulting events, make it impossible to decipher detector data if it is also integrated over several 25 ns periods. Information from each subdetector must therefore be correctly associated in time with the bunch-crossing from which it originated. The level-1 trigger does this by identifying a specific bunch-crossing as the source of a triggered event, enabling the selection of relevant data from the pipelines of the various ATLAS detectors.

The ATLAS LAr Calorimeters generate analogue pulses of several hundred nanoseconds duration, determined by the charge-collection time [5-1]; the ATLAS Tile Calorimeter generates narrower pulses from the rapid decay of the scintillation light [5-2]. However, in both cases the calorimeter pulses are modified in order to achieve improved signal-to-noise performance; the 'shaped' pulses have a peaking-time of ~45 ns. The LAr pulses are bipolar, with the positive part typically spanning three to four bunch-crossings. A Tile Calorimeter pulse after shaping is shown in Figure 5-18 and a simulated LAr pulse is shown in Figure 5-19.

A crucial prerequisite for the use of these shaped pulses in the calorimeter trigger is the conversion of broad pulses to a single digitized value, related to the true bunch-crossing time by a fixed latency. The digital value extracted from each pulse must also describe the E_T of the pulse with good precision. These two operations are jointly known as bunch-crossing identification in the context of the level-1 calorimeter trigger, and are performed by digital algorithms within the calorimeter trigger Preprocessor. The remainder of this section will describe these algorithms, their performance, and the properties required of the analogue pulses in order to operate successfully.

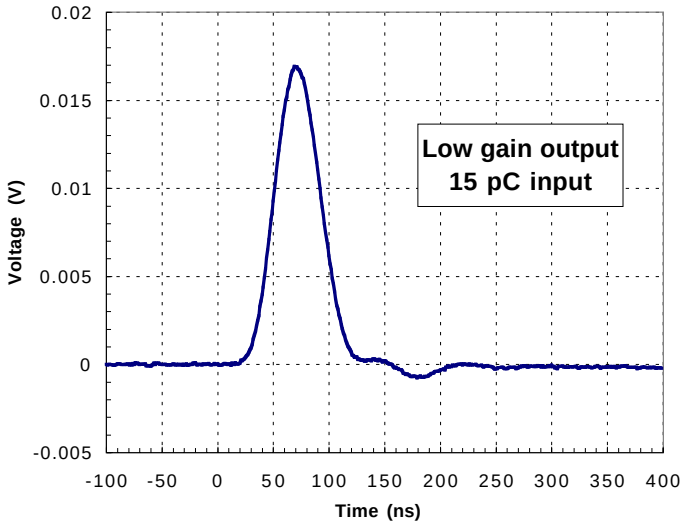

Figure 5-18 A typical Tile Calorimeter pulse, after shaping.

For pulses of constant shape, the extraction of the time corresponding to the centroid of the pulse can be done in several ways. The most common analogue approach uses a zero-crossing or a constant-fraction discriminator; such techniques deliver a logic pulse whose leading edge comes at a fixed time after the centroid of the analogue pulse. The time jitter of such circuits

clearly depends on the level of signal to noise in the analogue pulse but, for a good ratio, the timing can be extracted to better than 1 ns for pulses extending over many tens of nanoseconds. However, the sensitivity any such scheme is limited by noise, which reduces the efficiency for detection of small pulses.

The position of the pulse centroid can also be extracted digitally, as is proposed here. The analogue pulse is first digitized at a frequency which gives several samples over the duration of the pulse. These are fed into a digital pipeline and processed by a matched finite-impulse response (FIR) filter, which gives an output at a fixed time after the pulse centroid. Furthermore, the value of the digital output is related to the area, and not just the peak amplitude, of the original analogue pulse. Such an approach has four advantages over the analogue method:

- there is less sensitivity to noise-induced perturbations of the leading edge;
- the low-energy resolution of the energy measurement is improved;
- the two functions of time extraction and energy measurement are done simultaneously;
- the filter parameters can easily be made programmable, and hence matching to a wide range of pulse shapes is possible.

A requirement for the use of an FIR filter for pulse-centroid finding is that the pulse is of a known shape, which does not vary with the amplitude of the pulse (this is also partially true for edge discrimination schemes). The large dynamic range of the pulses from the ATLAS calorimeters makes this hard to achieve, while maintaining adequate resolution for low-energy signals. The solution is to set a maximum pulse amplitude which will be fully processed by the analogue electronics of the calorimeter; above this amplitude the pulse is truncated by the saturation of summing and shaping amplifiers. The maximum of the calorimeter trigger digitization scale is set below the maximum analogue amplitude, so that a large input pulse is cleanly truncated with one or more samples fixed to the maximum of the digital output. The resulting ‘digital saturation’ preserves the rising and falling edges of the input pulse provided there is no analogue distortion. This allows the position of the pulse maximum to be deduced from the number of saturated samples, by assuming both the pulse shape, and a fixed phase of the digitization clock with respect to the pulse maximum.

As a result of saturation, we anticipate two situations:

- unsaturated pulses, for which pulse shape is constant and independent of amplitude;
- saturated pulses, for which true signal amplitude exceeds some maximum and is truncated to a fixed value by the digitization of the calorimeter trigger.

The former describes the majority of ‘ordinary’ low-energy pulses over a dynamic range up to the maximum of the level-1 calorimeter trigger digitization scale, which is ~250 GeV. The latter is the case for pulses that exceed the maximum, which may arise from the high-E_T particles that of course are a feature of many interesting physics processes. It is essential that these largest pulses be visible to the calorimeter trigger, in order that the calorimeter signals resulting from these high-E_T particles can be used to trigger efficiently across the entire accessible E_T range. We note here that the largest pulses may be subject to a range of analogue distortions — these are discussed in more detail later. Separate BCID schemes are applied to these two cases, described below in Sections 5.3.1 and 5.3.2. In order to apply either of these algorithms successfully, the timing of the calorimeter trigger Preprocessor must be set up appropriately; strategies for achieving this are discussed in Section 5.3.3.

A property of both saturated and unsaturated BCID algorithms is that an identified pulse peak must always be preceded and followed by a zero output. This is exploited by the bunch-crossing multiplexing scheme described in Section 6.2.1.5, in order to double the transmission rate between Preprocessor and Cluster Processor for a given bandwidth.

5.3.1 BCID for unsaturated pulses

The BCID solution selected for unsaturated calorimeter pulses is an FIR filter followed by a peak-finder. A lookup table, which is not strictly part of the BCID logic, is also incorporated at this stage [5-5][5-6][5-7]. In the following, the results of performance and optimization studies are presented. In each study, simulated pulses following the form of LAr pulses described in [5-8] were generated for a fixed pulse peaking-time and scaled to a given E_T value, before simulated digitization and the addition of electronic noise and pile-up (see Section 5.3.1.4 for details). Such a pulse, of peaking-time 50 ns and corresponding to a particle of 2 GeV E_T, is shown in Figure 5-19, both as an analogue signal and digitized at 40 and 80 MHz. Tile Calorimeter pulses are expected to have both lower noise and be less affected by pile-up, as well as having shorter peaking-time. The energy resolution of the BCID algorithms acting on shorter pulses will be poorer, due to fewer digitizations, but this is offset by the lower intrinsic resolution of the Tile Calorimeter.

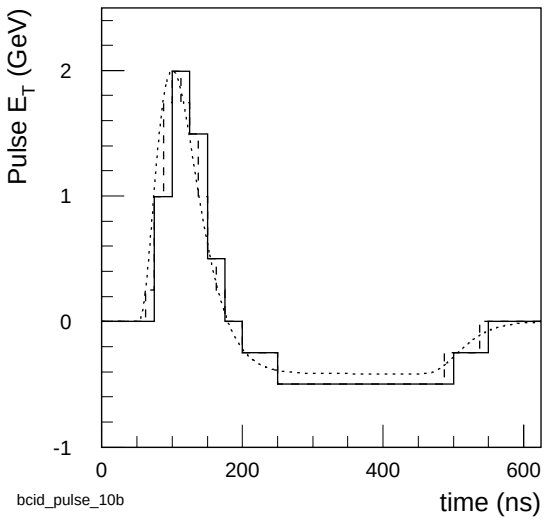

Figure 5-19 A typical LAr pulse of peaking-time 50 ns (dotted line), digitized with 10-bit precision at 40 MHz (solid line) and 80 MHz (dashed line).

5.3.1.1 Implementation

The FIR-filter BCID scheme is shown in Figure 5-20. A series of samples is held in a digital pipeline, which is advanced each time a new digitization of the pulse is produced. The filter acts on the contents of this pipeline in parallel, each sample being multiplied by a separate filter coefficient. To achieve optimum energy resolution and noise rejection, the set of coefficients is calculated with reference to pulse shape and noise spectrum. On each clock cycle, the weighted samples are summed to obtain the output of the filter. On the next cycle the pipeline advances, one new sample enters the pipeline, and each sample remaining in the pipeline is multiplied by the next coefficient in the filter. The sequence of filter outputs is processed by a peak-finder,

which vetoes any output which is not a local maximum in time. Any vetoed bunch-crossing is assigned an output value of zero. It should be noted that for some calibration procedures, the peak-finder will need to be disabled so that all data can pass through the BCID logic.

In parallel, the FIR filter output for each bunch-crossing is also sent to the lookup table mentioned above, which performs pedestal subtraction, final E_T calibration, and applies a noise threshold. The output of the lookup table will round the FIR-filter result to eight bits, which is sufficient precision for the trigger algorithms. The filter coefficients and lookup table are programmable for individual trigger towers, allowing the accommodation of tower-to-tower variation and the complete suppression of noisy towers.

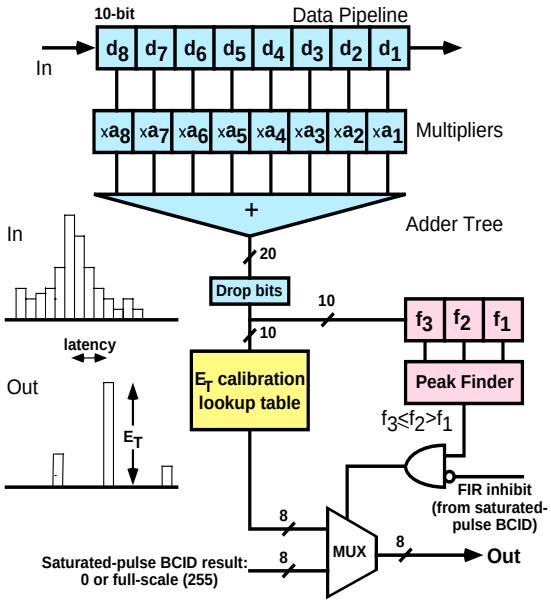

Figure 5-20 Finite-impulse response (FIR) filter for the implementation of BCID for unsaturated pulses. This example shows eight pipelined samples in the FIR filter.

The FIR-filter BCID logic is to be implemented within the Preprocessor ASIC, which processes two trigger towers independently. This ASIC, which is described in more detail in Section 6.2.3, performs a number of functions, and it is desirable that the size of the FIR-filter logic be kept small so that the ASIC complexity is minimized. The ‘scale’ of the logic depends on the number of bits per sample (A), the number of bits per filter coefficient (B), and the number of samples required to be pipelined. Each multiplier within the filter must handle $A + B$ bits, and the adder which sums the multiplier outputs must have adequate resolution. The appropriate values of these parameters are discussed in the following sections.

The number of samples used by the filter also determines the latency of BCID algorithm. If the filter requires X samples after that closest to the peak, then the BCID output is delivered with an irreducible latency of $X + 1$ bunch-crossings, allowing one bunch-crossing of latency for the peak-finder to take in the FIR-filter output from the bunch-crossing following a peak. There is of course additional latency, from the logic used to implement the filter, the peak-finder, and the lookup table (operating in parallel).

The successful implementation of an FIR-filter BCID scheme in both field-programmable gate arrays and ASIC devices was an important component of the level-1 calorimeter trigger demonstrator programme — see Section 7.4.

5.3.1.2 Digitization resolution and frequency

The two main parameters of the inputs to the filter are the digitization precision and frequency. FIR-filter BCID is required to operate in the regime up to $E_T \sim 250$ GeV, at which point the Preprocessor FADC saturates and an alternative BCID scheme must be used (Section 5.3.2). An 8-bit ADC spanning this dynamic range will therefore have a least-significant-bit (LSB) equivalent transverse energy of ~1 GeV, a 9-bit ADC will have an LSB E_T of ~0.5 GeV, and a 10-bit ADC will have an LSB E_T of ~0.25 GeV[1]. The digitization of an LAr pulse, with its characteristic bipolar shape, is illustrated in Figure 5-19, with a 2 GeV pulse digitized to 10-bit precision at 40 and 80 MHz. In principle, digitization at greater precision reduces the quantization noise of the ADC which, assuming ideal performance, is $1/\sqrt{12}$ of the LSB equivalent energy. In practice, the digitization noise is dominated by the ~450 MeV r.m.s. noise from the LAr calorimeter and its electronics in each trigger tower.

A comparison of simulated 8-bit and 10-bit digitization precision is shown in Figure 5-21 for a variety of digitization rates (see below). The influence of noise or variations in pulse properties can cause the filter to generate a maximum on either an earlier or later cycle, which will lead to an incorrect position in the trigger pipeline and hence both false and missed triggers. The efficiency of the algorithm at a given energy is defined as the proportion of pulses from which the FIR filter generates a non-zero output with the correct latency with respect to the original bunch-crossing. It is clear that FIR-filter BCID is efficient down to energies that are only two to three times the mean simulated noise, which is ~0.5 GeV. The algorithm is ~55% efficient at an E_T of 1 GeV, rising to ~90% efficient at an E_T of 2 GeV, >97% efficient at 3 GeV and effectively 100% efficient for pulses of 4 GeV and above (up to the saturation limit of ~250 GeV). Digitization with 10-bit resolution offers small efficiency gains at low energies compared to 8-bit resolution, ~5% at $E_T = 2$ GeV for example (in fact the ninth bit delivers most of this improvement). The use of a smaller step-size in the digitized energy values offers other advantages, by providing a finer resolution on thresholding of individual towers within the lookup table, thus enabling more effective and controllable suppression of noise. We should also remark that the apparent non-zero 'efficiency' at 0 GeV is due to digitized noise, most of which will be suppressed by the application of the noise threshold in the calibration lookup table.

Commercial devices for 10-bit digitization with acceptable latency and linearity performance are becoming available at reasonable cost, although they currently have greater latency and consume more power than 8-bit devices. We assume that these characteristics will improve further in the next few years. In our baseline design, 10-bit digitization will be employed, and so the resulting small gains in BCID efficiency will be realized. The additional low-order bits above 8-bit precision will be suppressed in the lookup table after thresholds are applied, resulting in 8-bit trigger-tower values transmitted to the Cluster Processor and used as the basis for the 2×2 sums made in the Preprocessor and transmitted to the Jet/Energy-sum Processor. The use of 10-bit digitization allows the fullest use of the 8-bit dynamic range even after pedestal subtraction and re-calibration of E_T have been done.

The preprocessing of each tower must produce an output for each 25 ns bunch-crossing, and so sampling must be at a multiple of 40 MHz. In principle, digitization at a higher frequency distributes quantization noise over a greater frequency range, reducing the contribution in the

1. The exact value of the step size depends on the limits of the dynamic range; the proportion of the dynamic range which will be dedicated to the negative-going part of the signals from the LAr calorimeter is not yet determined, but will probably be ~10%.

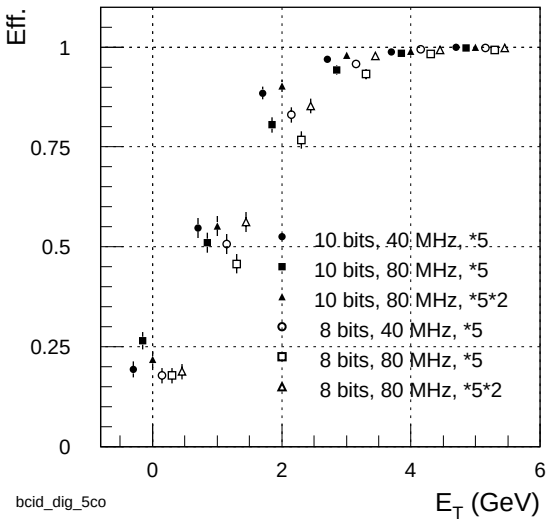

Figure 5-21 Efficiency of five bunch-crossing FIR-filter BCID for low-E_T pulses, with digitization at 40 or 80 MHz (single- or double-length FIR filter), and at 8- and 10-bit precision. Points at the same integer value of E_T have been offset for clarity.

signal region. In fact, digitization at 80 MHz samples a greater proportion of the noise spectrum than at 40 MHz, since the noise power spectral density peaks at 10 MHz, and the signal power spectral density peaks at 2 MHz with almost no component at frequencies higher than 10 MHz. The presence of noise components in the region above half the sampling frequency can generate aliasing noise at lower frequencies, degrading the signal-to-noise separation. This can be limited by the application of an anti-aliasing filter with a cut-off frequency at 20 MHz.

The gains for FIR-filter BCID by sampling at 80 MHz rather than 40 MHz are therefore expected to be limited to an effect similar to that achieved by increased digitization precision. This is shown in Figure 5-21: doubling of digitization frequency delivers improved efficiency with 8-bit digitization, but little or no improvement with 10-bit digitization. Overall, 80 MHz sampling with 8-bit precision offers a similar small efficiency improvement to 40 MHz sampling with 10-bit precision. There is negligible further gain from combining both improvements, i.e. sampling at 80 MHz with 0.25 GeV digitization precision.

The cost is that BCID processing at 80 MHz requires twice as many pipelined samples as at 40 MHz, and 80 MHz FADCs are significantly more expensive than 40 MHz equivalents. Note that 80 MHz digitization without doubling the number of samples used in the filter is in fact more noise sensitive than 40 MHz digitization, since the samples span less of the pulse. From the above considerations, it is therefore unnecessary to implement FIR-filter BCID for unsaturated pulses at 80 MHz if a 10-bit digitization at 40 MHz is available.

A further ingredient in the resolution for LAr signals is the proportion of the dynamic range required to digitize the negative-going part of the signal. This has two effects: firstly it determines how much of the negative ‘undershoot’ which follows the main pulse is digitized, and secondly it determines the equivalent energy of the least-significant bit, and hence the digitization resolution. The negative dynamic range is determined by a pedestal which is applied as an analogue voltage added to the incoming analogue signal. If the combined signal-plus-pedestal falls below zero (i.e. the negative signal is larger than the positive pedestal), the FADC gives an output of zero. Various pedestal values are illustrated for a five-coefficient 40 MHz FIR filter in Figure 5-22. The truncation of the negative part of the signal results in no

loss of sensitivity if the pile-up is ignored, and even coarsening the least-significant bit resolution by ~20% does not degrade the efficiency for small pulses.

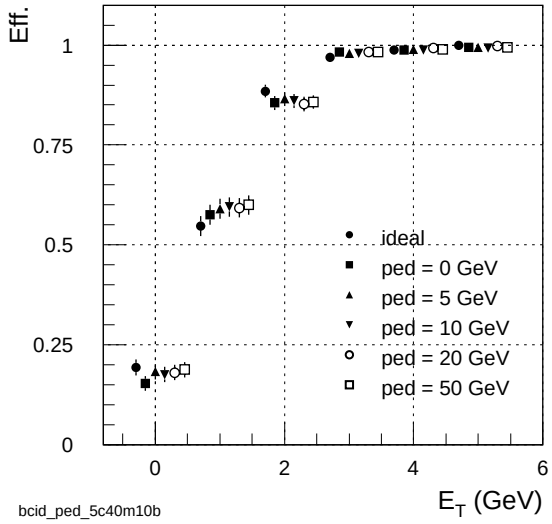

Figure 5-22 Comparison of FIR-filter BCID efficiency with pedestal values for a five bunch-crossing FIR filter and 10-bit digitization at 40 MHz.

5.3.1.3 FIR-filter optimization

In order to minimize the BCID implementation overhead, the parameters of the FIR filter — number of pipelined samples and filter coefficient precision — have been studied. In the approximation of white noise, an FIR filter with samplings covering the entire pulse-width is sufficient. Starting with eight 10-bit samples at 40 MHz, which completely span the width of the positive portion of a standard pulse of 50 ns peaking-time, the number of samples was progressively reduced, by removal of the smallest coefficient. The results are shown in Figure 5-23. The performance is almost indistinguishable down to four samples, but significantly degraded for fewer samplings. Five samples is the preferred filter length to give some latitude for variation of pulse width and timing.

The required precision of digital coefficients is illustrated in Figure 5-24, which shows that 6-bit coefficients in a 5-sample FIR filter achieve the same precision as real values, and performance is negligibly degraded down to 3-bit precision. In the study for the number of filter coefficients above, the full 8-sample filter included some negative coefficients corresponding to the tail of the pulse; these are not included in the 5-sample filter, which is of equivalent performance. Hence it appears that the use of 4-bit or 5-bit coefficients will give an adequate performance with some margin for safety, and that there is no need to allow for a sign bit in filter coefficients, although at very high luminosity there might be some reduction in pile-up noise if the first coefficient were to be negative.

In conclusion, the above studies suggest that adequate performance with a good safety margin can be obtained with 10-bit digitization at 40 MHz, five bunch-crossings sampled by the FIR filter, and 5-bit unsigned multiplier coefficients. The full precision of the output need not be retained, and 10 bits is sufficient for the address space of the lookup table.

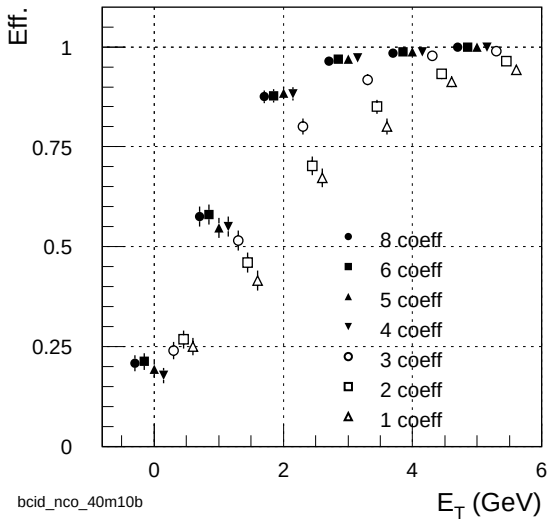

Figure 5-23 Comparison of FIR-filter BCID efficiency with filters of varying number of stages; in shortening the filter, the smallest coefficient was removed.

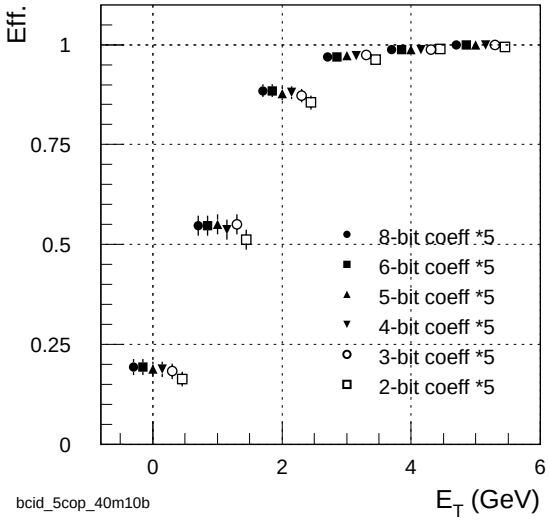

Figure 5-24 Comparison of FIR-filter BCID efficiency with number of bits used to encode coefficients of a five bunch-crossing filter.

5.3.1.4 FIR-filter BCID performance

The above studies have all been carried out with standard noise conditions. Electronics noise was generated from a time series taken from a Gaussian of width 450 MeV, which was then modified to reproduce the noise spectrum for the LAr electronics, followed by a fourth-order low-pass Butterworth filter with a cut-off frequency of 20 MHz, in order to reduce aliasing noise from the 40 MHz digitization of noise components in the >20 MHz range. Such a filter might be used in the Preprocessor. The resulting noise spectrum is shown in Figure 5-25a. Pile-up was generated as a single pulse for every bunch-crossing, of shape identical to the signal and with the E_T spectrum shown in Figure 5-25b (taken from ATRIG at luminosity of 10^{34} cm^{-2}s^{-1} and valid for $|\eta| < 3$). The two noise sources combined have an r.m.s. width of 525 MeV.

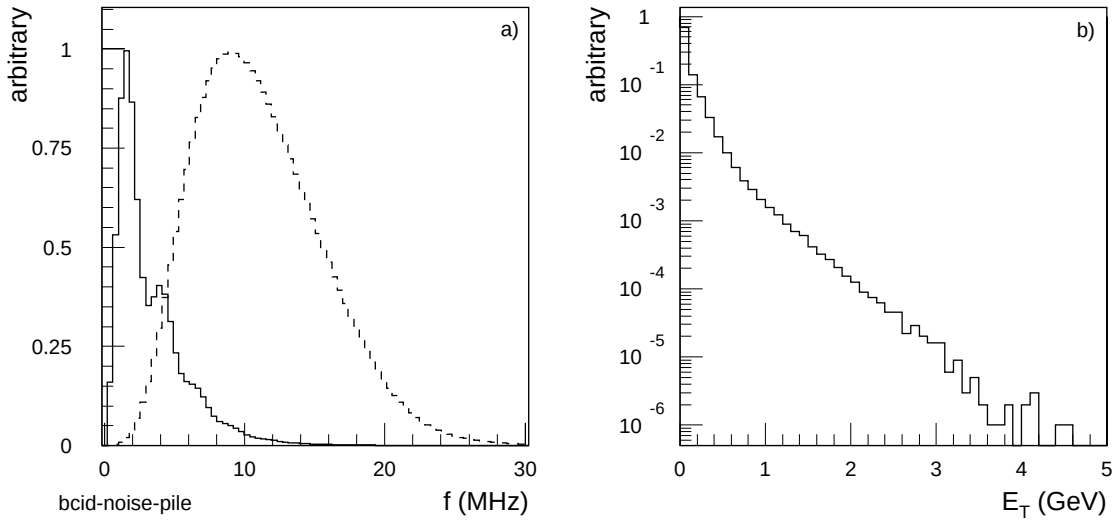

Figure 5-25 (a) The power spectral density used to modify white noise in simulating BCID (dashed); the signal power spectral density is superimposed as a solid line. (b) The pile-up spectrum — taken from ATRIG at high luminosity — used in simulating BCID.

Clearly it is important to know how the performance of a standard filter is affected by changes in these noise conditions. Figure 5-26 shows FIR-filter BCID efficiency for a range of values of noise amplitude. Small changes of ±10% affect the efficiency at 1–2 GeV by only a few per cent, and are negligible for higher-energy signals. Larger increases in noise do degrade efficiency significantly for $E_T < 5$ GeV. Note that the efficiency for fake BCID values (signal = 0) is largely independent of the noise amplitude; those noise fluctuations which pass the FIR filter and peak finder are simply any local maxima which remain after digitization effects. The probability for fake noise outputs therefore depends mainly on the combination of digitization and filter, plus the spectrum of the noise, and is to first order independent of the noise amplitude. The majority of those small amplitude fake pulses passed by the FIR filter will be removed by the threshold function applied within the calibration lookup table.

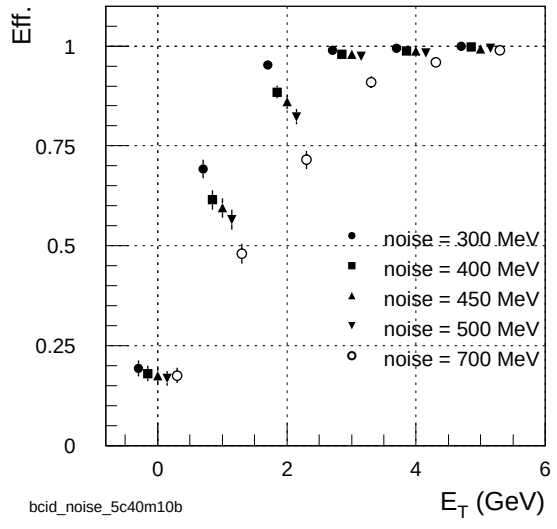

Figure 5-26 Comparison of five bunch-crossing FIR-filter BCID efficiency with r.m.s. of electronics noise.

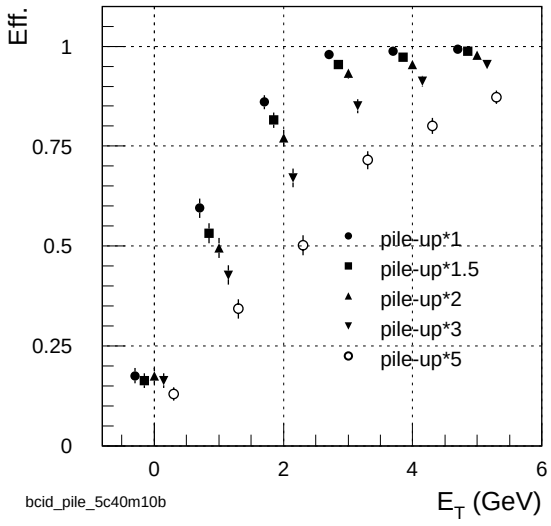

Figure 5-27 Comparison of five bunch-crossing FIR-filter BCID efficiency with scaling of pile-up spectrum.

Pile-up is a difficult noise source to reject effectively because it is composed of pulses of identical shape to the signal. A comparison of pile-up conditions is shown in Figure 5-27; in each case the level of pile-up was increased by the given factor. This rapidly degrades efficiency, with a doubling of the pile-up spectrum reducing BCID efficiency by ~10% at 1–2 GeV E_T, and ~5% at 4 GeV E_T. This will affect BCID performance in the 1–5 GeV range should the LHC exceed its design luminosity by a factor of two or more; conversely, at low luminosity, there will be a small improvement in BCID efficiency compared to that discussed here (although electronics noise is the dominant source).

Having established that FIR-filter BCID performs well for a standard pulse of known shape and phase, an obvious extension is to consider how well a given FIR-filter scheme will perform for a pulse which differs from the nominal values in one or both of these degrees of freedom. This models the robustness of the filter against uncertainty in either the pulse peaking-time or the synchronization of the digitization strobe with the analogue pulse.

For this study, a five-sample FIR filter optimized for a nominal 50 ns pulse width and a 40 MHz digitization strobe in phase with the pulse maximum was simulated over the two-dimensional space of pulse peaking-time, t_{pk}, and digitization strobe offset from the pulse maximum, Δt, determining the BCID efficiency at each point. The resulting efficiencies are shown in Figures 5-28a for pulses of 2 GeV E_T and 5-28b for 5 GeV E_T. It is clear that the good efficiency achieved for

'in-time' pulses is degraded when the synchronization of the pulse is altered beyond ±6 ns; within this region efficiency remains at a plateau of 100% for 5 GeV pulses, and is not degraded by more than 10% for pulses of 2 GeV. For digitizations offset from the peak of the analogue pulse by more than 6 ns, the efficiency rapidly falls by ~10% per nanosecond. Broader pulses are more sensitive to mistiming of digitizations.

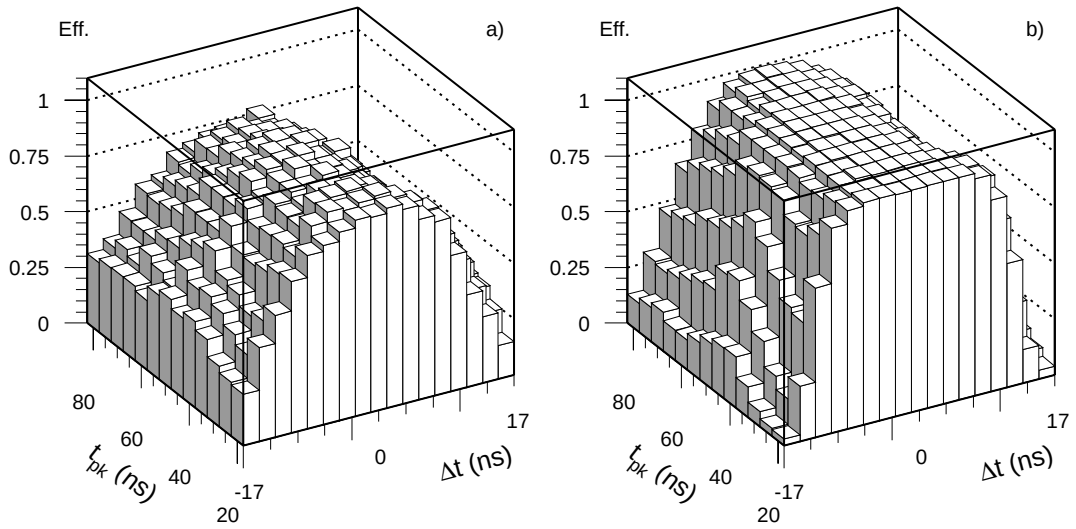

Figure 5-28 FIR-filter BCID efficiency for a nominal 50 ns pulse peaking-time over the two-dimensional space of actual pulse peaking-time, t_{pk}, and pulse-peak digitization offset, Δt, for pulses of (a) 2 GeV E_T and (b) 5 GeV E_T. The efficiency degrades as the pulse synchronization diverges from the zero offset for which the filter was designed.

The digitization strobe on each trigger tower can be adjusted in 1 ns steps, and the timing stability of analogue pulses is expected to be better than ±3 ns. Therefore the large phase differences between digitization and pulse maximum at which efficiency is significantly degraded should not occur in a well-calibrated system (how that calibration is achieved is addressed in Sections 5.3.3 and 8.7), and the FIR-filter BCID algorithm is robust for the variations of timing parameters anticipated.

A cautionary note should be added: the efficiency refers here to the correct identification of the peak of an analogue pulse, for values of the pulse peaking-time other than that for which the filter coefficients were designed. But a crucial parameter in relating the identified peak of a pulse back to the source bunch-crossing is the peaking-time itself; clearly a 70 ns pulse will peak more than one bunch-crossing later than a 40 ns pulse, in response to the same incident particle. Thus, even though a given filter finds the peak of a pulse of rather different width, one would need to have established the correct peaking-time to an accuracy of ~10 ns in order to assign the filter output to the correct bunch-crossing. And if that had been done, one would use an appropriate set of filter coefficients. The issue of timing set-up and tolerances is discussed more thoroughly in Section 5.3.3.

A further criterion by which to judge BCID performance is the energy resolution of unsaturated pulses. Particularly relevant are timing-dependent systematic shifts in the measured energy of pulses, which would clearly degrade the efficiency to trigger on objects passing fixed thresholds. The mean and r.m.s. energy for pulses over a range of E_T values were measured over the $t_{pk} - \Delta t$ space; the resulting mean energies are shown for pulses of 2 GeV and 5 GeV E_T in Figures 5-29a and 5-29b, respectively.

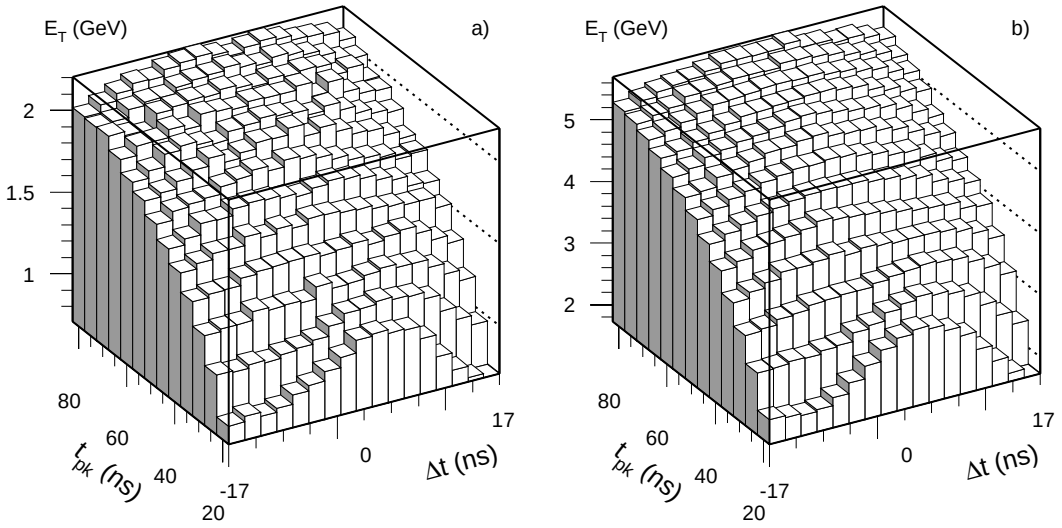

Figure 5-29 Mean measured energy for a nominal 50 ns pulse peaking-time over the two-dimensional space of actual pulse peaking-time, t_{pk}, and pulse-peak digitization offset, Δt, for pulses of (a) 2 GeV E_T and (b) 5 GeV E_T. The reconstructed energy is systematically reduced as the pulse becomes narrower than the one for which the filter was designed.

It is clear that pulses which are narrower than the nominal value for which the filter coefficients were chosen suffer a systematic shift in measured energy to lower values, since off-peak samples are at smaller fractions of the pulse maximum than for the broader pulses on which the filter was based. A smaller shift in measured energy to larger values occurs for pulses which are broader than the nominal width. It is this energy bias which in fact sets the limit on acceptable deviations from nominal pulse widths; in order that the measured energy be within 10% of the true E_T of the pulse, the FIR-filter coefficients must be designed for a pulse width within ±10 ns of the true value. The E_T resolution is not greatly affected by small excursions in pulse synchronization.

5.3.2 BCID for saturated pulses

Bunch-crossing identification must also be performed for pulses which exceed the maximum of the level-1 calorimeter trigger digitization scale and which also cause saturation in the analogue front end — such pulses are expected as signatures of many physics processes of interest. The truncation of the pulse by digital saturation of the FADC results in a series of saturated samples over the period for which the analogue signal was above the saturation limit. The FIR-filter algorithm of Section 5.3.1 can still act upon a truncated pulse and sometimes produce an output with correct timing[1]. However, the FIR filter fails when the samples in the pipeline are all saturated, which occurs for very high-E_T pulses when the pulse width is significantly different from that for which the filter was designed.

If the FIR-filter algorithm cannot be relied upon, other methods must be sought. Since the peak of a truncated pulse carries little information after digitization, more use must be made of the regions of the pulse which are below the saturation level, i.e. the leading and trailing edges of

1. Note that an FIR filter with a single maximum coefficient will produce a maximum output at some time for any saturated pulse of finite extent — whether this is in time for the true peak of the non-truncated pulse depends on the pulse parameters.

the pulse. However, a single sample on the leading edge of a large pulse does not provide much information on the timing of that pulse, and the large dynamic range of pulses makes it difficult to guarantee multiple samples on the rising edge of large pulses using a fixed digitization strobe. The time taken for a pulse of 50 ns peaking-time to reach a 250 GeV saturation limit is shown in Figure 5-30; for a pulse of 5 TeV E_T, the analogue pulse would in principle reach saturation within 8 ns, which would require a digitization rate of several hundred megahertz to provide two samples on the leading edge.

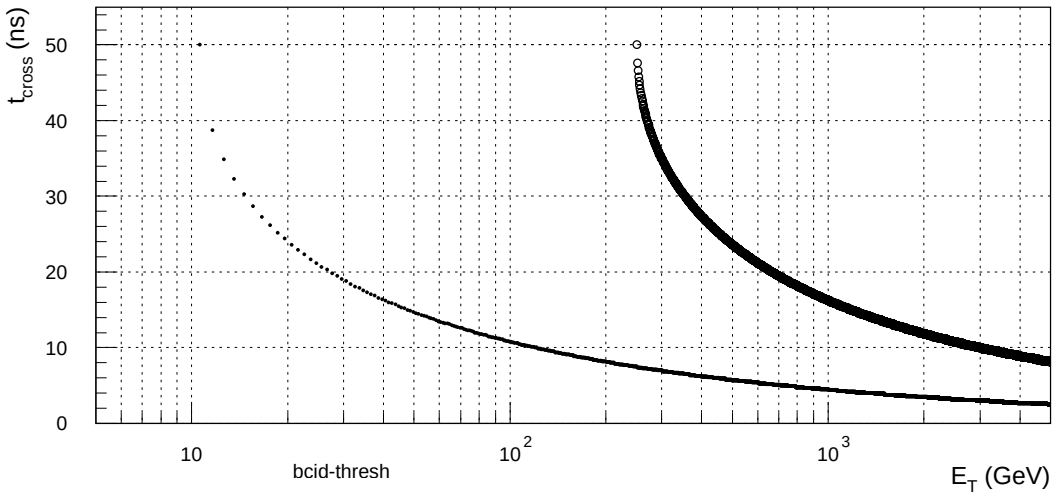

Figure 5-30 The time taken for an ideal LAr pulse of t_{pk} = 50 ns to cross thresholds equivalent to E_T = 10 GeV (small points) and E_T = 250 GeV (large points), as the amplitude of the pulse is varied over the range 1–5000 GeV. The threshold-crossing time scales linearly with the peaking-time of the pulse.

The position of the peak of a pulse of known shape can be determined from the number of saturated samples, provided that the peaking-time is unaffected by saturation, and that the maximum deviation of the pulse phase from a nominal synchronization is a small fraction of the sampling interval. The advantage of such an algorithm is that the logic can be integrated into the digital processing following the FADC, using samples taken at 40 MHz. There is no need to split the analogue signal, which depending on the implementation might be necessary if an analogue discrimination scheme were to be used. A sequence of samples can be digitally discriminated at a level close to the digitization maximum value and pipelined; the resulting binary sequence can then be interpolated for the pulse maximum using a lookup table. The proposed implementation of this lookup-table BCID scheme is described in more detail in Section 5.3.2.1, and its performance over a range of pulse peaking-times and digitization offsets is investigated in Section 5.3.2.2.

The drawback is that saturated pulses are required to be of known shape, and this can be hard to achieve for real analogue devices processing very large pulses. Although the digital saturation of the FADC provides a clean truncation of the pulse at a known amplitude, saturation of large pulses in the upstream analogue processing chain is not necessarily such a simple matter. Saturation of analogue amplifiers can result in prolongation of the saturated part of the pulse or other distortion of the falling edge, which becomes progressively worse as the amplitude of the original pulse is increased. The pulse shape is also altered when an ideally-truncated pulse is followed by a re-shaping element, such as the pole-zero cancellation of the LAr tower builder or the final integration in the LAr waveform monitor, or even the long cable from the detector to the trigger cavern. The pulse tends to be smeared and the ‘clipped’ nature of the pulse is lost. Finally, real amplifiers have finite slew rates, and the fast turn-on of the

largest pulses in the ideal case may not be exactly reproduced. Since the degree of saturation and subsequent distortion depends on pulse amplitude, then so does the effective pulse shape at the input to the calorimeter trigger Preprocessor. If the pulse shape *does* vary with amplitude, it is difficult to construct a lookup table which can correctly locate the true maximum for the entire range of amplitudes required.

The limitations of lookup-table BCID are under investigation, as improved modelling of the analogue pulse shape and saturation effects become available. Whilst the lookup table is the current baseline solution for large-pulse BCID, alternatives based on leading-edge discrimination are also under investigation and are described in Section 5.3.2.3.

5.3.2.1 Lookup-table BCID

The interpolation of a saturated pulse to find its maximum will be implemented in ‘saturation logic’ within the Preprocessor ASIC, processing the same stream of digitized samples as the FIR-filter BCID. Although the FADC provides an automatic truncation of large pulses, in practice discrimination will be applied by a separate digital comparator at a saturation level which can be adjusted as necessary, in order to have some control over analogue distortion effects for large pulses. Saturated samples are converted to single-bit values — 1 if above the comparator threshold, 0 if below — and each successive comparator result is stored in a shift register. The contents of the shift register can be rapidly decoded by a lookup table, which is addressed by the sequence of samples in the shift register. The table is set to produce a non-zero output only when a pattern matching one of the expected peak signatures is presented at the input. A schematic of the saturated-pulse BCID with lookup-table logic is shown in Figure 5-31.

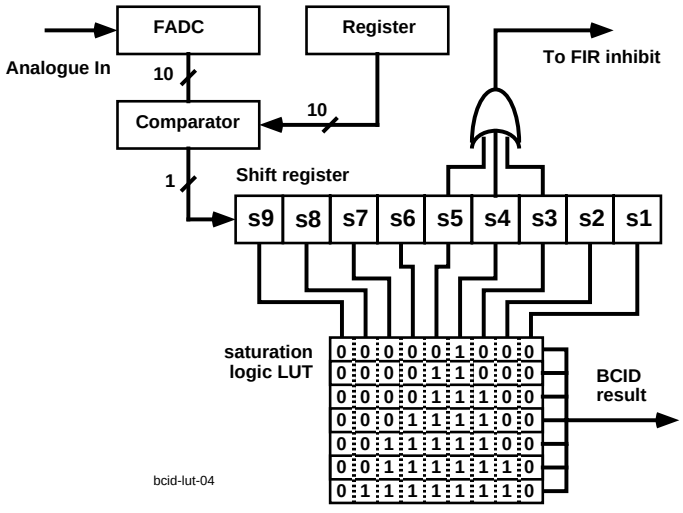

Figure 5-31 Digital comparator plus lookup-table logic for implementation of saturated-pulse BCID.

The most obvious version of such a lookup table is constructed by simply truncating a ‘template’ analogue pulse at the comparator threshold, adjusting the amplitude of the pulse until 1, 2, 3 etc. consecutive saturated samplings are observed. To evaluate the performance of the lookup-table BCID, such a table was constructed by discriminating the same LAr pulse shape as was used for the unsaturated BCID studies at 255 GeV, assuming zero pedestal on the analogue signal, t_{pk} = 50 ns and zero digitization offset. For a pulse of this peaking-time and phase, even a maximum-amplitude 7 TeV pulse can only produce up to four saturated samples.

To extend the table to broader pulses, additional rows were added to cope with large saturated pulses of peaking-time up to 80 ns. The patterns of 1s and 0s and E_T thresholds at which each additional saturated sample arises are given in Table 5-5.

Table 5-5 Lookup table for BCID of pulses up to 80 ns wide, valid up to 8 TeV, together with the minimum energies at which each pattern occurs (for a pulse with zero digitization offset from the maximum of the analogue pulse).

Bunch-crossing number											
1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**E_T threshold**	**t_{pk}**	**8-bit output (hex)**
0	0	0	1	0	0	0	0	0	256 GeV	50 ns	FF
0	0	0	1	1	0	0	0	0	352 GeV	50 ns	FF
0	0	1	1	1	0	0	0	0	464 GeV	50 ns	FF
0	0	1	1	1	1	0	0	0	822 GeV	50 ns	FF
0	0	1	1	1	1	1	0	0	2050 GeV	60 ns	FF
0	1	1	1	1	1	1	0	0	4772 GeV	60 ns	FF
0	1	1	1	1	1	1	1	0	3620 GeV	70 ns	FF

When a match to one of the above patterns is obtained on the lookup-table inputs, an output indicating a saturated pulse is generated, on the single bunch-crossing corresponding to the peak of the input pulse. The output from the lookup table must be set to the maximum of the digitization scale, which is FF (hex) with an 8-bit output from the calibration lookup table. This output, equivalent to an energy deposit of about 255 GeV, should assure that all saturated pulses can result in level-1 triggers. The latency of the lookup-table BCID scheme is determined by the number of samples after the peak required to make the identification; for the above table, a positive BCID cannot be made sooner than five bunch-crossings after the peak digitization. To this must be added the latencies of the comparator and the lookup table, perhaps another bunch-crossing together. The latency of the two BCID schemes — FIR-filter and lookup table — must be synchronized by adding a delay to whichever is the faster.

In some circumstances, both FIR-filter and lookup-table schemes may produce a BCID value from the same pulse — lookup-table BCID is given precedence whenever a saturated sample occurs. In order that the two schemes cannot possibly give non-zero output on consecutive bunch-crossings, the saturated samples held in the shift register are used to inhibit the FIR-filter BCID for one bunch-crossing before and after the identified peak.

5.3.2.2 Saturated-pulse performance

Lookup-table BCID has been tested against pulse parameters of peaking-time, t_{pk}, and digitization offset, Δt, in the same way as for FIR-filter BCID. Efficiency is defined in the same way, from the proportion of pulses for which the lookup table identifies the pulse peak with the correct latency. For large saturated pulses, noise from any anticipated source (other than disastrously high levels of crosstalk) is irrelevant to the pulse shape; all pulses of a given energy are effectively identical. The BCID scheme is therefore only affected by the timing of the pulse, and once the timing is sufficiently different from the nominal the lookup-table BCID essentially from switches 100% to zero efficiency.

The result is sharply-defined regions of parameter space for which failure occurs, dependent on the E_T of the pulse. For pulses of $E_T > 700$ GeV, this failure occurs for a digitization offset of as little as 5 ns, even close to the nominal pulse width for which the lookup table was designed. Since the region of good lookup-table performance depends on the pulse energy, the evaluation of integrated efficiency over the whole energy range is essential in determining the envelope of good performance. Integrated efficiencies are shown in Figure 5-32. The lookup table is suitable for pulses within ±10 ns of the nominal width, and with digitization synchronization within ±5 ns of the peak of the analogue pulse. This is comparable to the good performance region of FIR-filter BCID described in Section 5.3.1.

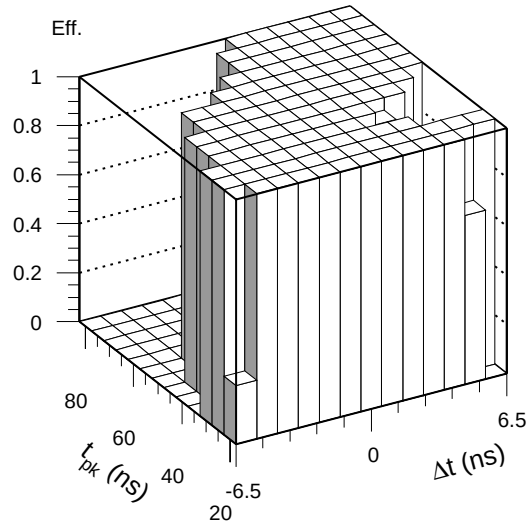

Figure 5-32 Lookup-table BCID efficiency for a nominal 50 ns pulse over the 2D space of actual pulse peaking-time, t_{pk}, and pulse peak digitization offset, Δt. The efficiency over the range 260–3000 GeV E_T is represented as a product, where an efficiency of zero at any energy for a t_{pk} and Δt combination results in a zero entry in the histogram.

It should be noted that these studies were made with an 'ideal' LAr pulse, which had no modelling of the distorting effects of amplifier latch-up, smearing of truncated pulses, or limited output slew rates. These potentially undermine the performance of the lookup table to a far greater extent than FIR-filter BCID; the lookup table relies on the binary classification of the pulse as above or below a threshold on each digitization, and a small change to the pulse amplitude can greatly affect the way the pulse is interpreted by such a binary scheme. It is essential that the understanding of the non-ideal nature of the analogue components will continue to improve, and that further simulation of lookup-table BCID algorithm in the near future will provide more evidence as to its ultimate suitability and be confirmed by measurements using prototypes of the calorimeter analogue electronics.

5.3.2.3 Alternative methods for saturated pulses

Given the current uncertainty in the analogue performance of large pulses, alternative methods have been sought for ensuring that large pulses will be identified in the correct bunch-crossing. One possibility is the use of leading-edge discrimination to detect a transition in the analogue input, and thus to determine the time at which the pulse 'started'. Note that the use of such a method only for large pulses ($E_T > 250$ GeV) greatly reduces the time skewing associated with the measurement of the pulse at a fixed amplitude over the full dynamic range. From Figure 5-

33 it can be seen that, for a peaking-time of 50 ns, the time of crossing a 10 GeV threshold is skewed by only about 5 ns in the amplitude range above 250 GeV. This contributes little uncertainty to the association of a leading-edge transition at a low level with the correct bunch-crossing for the pulse. Furthermore, the real-device effects mentioned above tend to slow the rise time of the pulse, and become stronger as the pulse amplitude increases — they therefore act in the opposite sense to the amplitude-induced time-skewing of the ideal pulse, and should reduce the actual time-skewing of the threshold transition. A threshold in the region 5–10 GeV is sufficiently high that noise will not contribute a significant rate of fake transitions, while keeping the time-skewing at a reasonable level.

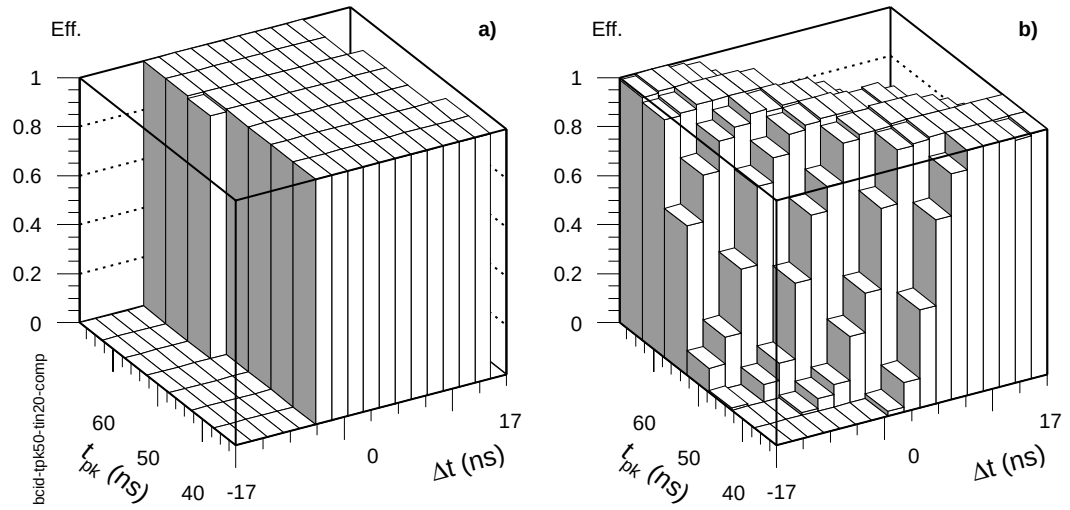

Figure 5-33 (a) Discriminator-BCID efficiency over the 2D space of actual pulse peaking-time, t_{pk}, and pulse-peak digitization offset, Δt. The efficiency over the range 260–3000 GeV E_T is represented as a product, where an efficiency of zero at any energy for a t_{pk} and Δt combination results in a zero entry in the histogram. (b) FIR-filter BCID efficiency for a 5 GeV E_T pulse of nominal 50 ns peaking-time, with strobe timing set on the leading edge of the pulse rather than the peak.

The transition of the discriminator output indicating the crossing of the threshold by an analogue pulse must be integrated into the digital processing stream. It is proposed to sample the discriminator output synchronously with the FADC digitization strobe, and to use this digital stream to determine the edge transition. This requires that the strobing of the discriminator follows the actual transition with sufficient delay that the time-skewing of the transition over the required amplitude range, or from small changes or uncertainties in the pulse parameters, does not shift the digital transition by ±1 bunch-crossing.

The sequence of sampled discriminator outputs determines the timing of the edge of the analogue pulse, but this timing can only be deduced with required uncertainty if the threshold is at a small fraction of the full-pulse amplitude. A second and much higher threshold must therefore be applied to confirm that a pulse is large enough to give small uncertainty in interpreting the low-threshold crossing time, and also that it is saturated and therefore the FIR-filter BCID result should not be used. This second threshold is implemented using a digital comparator acting on the FADC output, in order to limit the number of analogue components.

A suggested discriminator-BCID scheme is shown in Figure 5-34. The sequences of discriminator and comparator outputs — a low analogue threshold at 5–10 GeV and a high digital threshold at ~250 GeV — are pipelined in shift registers; from the first the edge timing is determined, and from the second the confirmation of the large pulse is made (using one or more

samplings subsequent to the low-threshold crossing). The logic shown to interpret these sequences could equivalently be implemented in a lookup table addressed by the contents of both shift registers. The output from this BCID scheme must be synchronized with the FIR-filter BCID result. If needed, the timing of the digitizations used by the FIR filter can be adjusted independently of the discriminator strobe timing by the use of an additional delay, as shown in Figure 5-34.

The discriminator BCID is faster than the FIR-filter algorithm, since the former does not require information after the peak of the pulse, and its result must therefore be delayed by an output FIFO. A further requirement is that any non-zero BCID result is followed by a zero; this arises automatically from the digital sampling of the comparator, which cannot make the 0–1 transition on two consecutive bunch-crossings. Finally, a non-zero output must also be used to override the FIR-filter BCID; this is done by the inhibit scheme shown in Figure 5-20. The required output value indicating a saturated deposit is generated from a register, and the multiplexer merges the FIR-filter and discriminator-BCID results according to the FIR-filter inhibit signal.

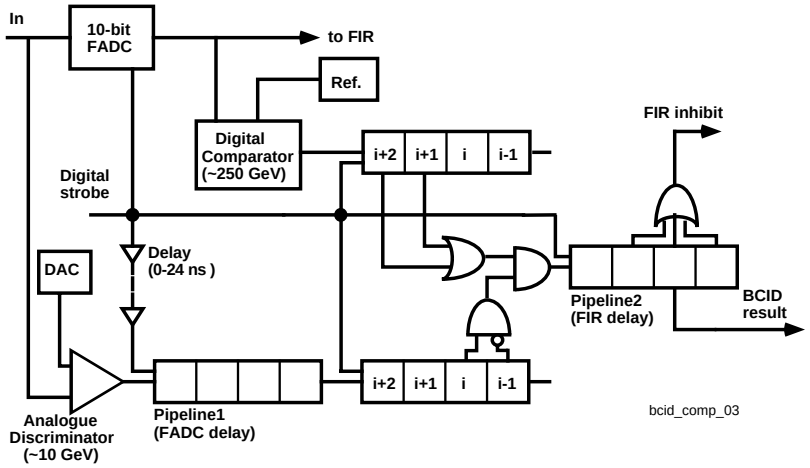

Figure 5-34 Discriminator logic for the implementation of saturated-pulse BCID.

Such a scheme has been simulated to determine the algorithm's response to uncertainty in the pulse peaking-time and the placement of the comparator sampling strobe, as for the FIR-filter and lookup-table BCID schemes above. For a pulse peaking-time of 50 ns and a low threshold of 10 GeV (which is crossed within 8 ns of the start of a pulse of E_T > 250 GeV), the discriminator was nominally sampled at 20 ns into the pulse. This would allow the timing of the pulse to be in error by several nanoseconds before the peak would be assigned to the incorrect bunch-crossing. From Figure 5-33a we see that the digitization strobe may be misplaced by up to 6 ns before the discriminator mis-identifies the bunch-crossing. Since the discriminator is relatively insensitive to the pulse peaking-time, it remains efficient when the peaking-time differs by up to 15 ns or more from the nominal value. Both of these margins are comparable to the degree of robustness offered by the FIR-filter BCID scheme. Because the discriminator samples only the first part of the rising edge of the pulse, this BCID algorithm is expected to be largely insensitive to degradation of the pulse shape by the real analogue effects discussed above, and so the illustrated performance is likely to be achieved in a real system.

It is useful to ask how lookup-table BCID and discriminator BCID differ in their sensitivity to the synchronization of the strobe with the analogue pulse. When the strobe was set up on the

peak of the pulse, the FIR-filter BCID algorithm was relatively robust to errors in the pulse timing (Section 5.3.1). However, when the timing of the strobe is set relative to the leading edge of the pulse, one can expect greater sensitivity of FIR-filter BCID to the pulse peaking-time, since a lengthening of the pulse shifts the pulse peak with respect to the digitization strobe, which rapidly degrades efficiency. This is illustrated in Figure 5-33b, which shows the FIR-filter BCID applied to pulses of 5 GeV E_T using the same timing as for the discriminator; the digitization strobe falls at 20 ns, 45 ns, etc. into the pulse, i.e. 5 ns from the peak of the nominal 50 ns analogue pulse, and the filter coefficients have been re-optimized for this. The FIR-filter BCID is now far more sensitive to pulse peaking-time, as shown by the strong correlation of maximal efficiency with t_{pk}. The sensitivity to the actual placement of the strobe is as before, efficiency being unaffected if the strobe is moved by ±5 ns.

A somewhat different solution, based on a clock-independent discriminator followed by a continuously-variable delay using monostables, is shown in Figure 5-35. The discriminator output can be fed as an additional bit into the data stream from FADC, before coarse alignment in time is performed in the clocked FIFO. The scheme ensures that the discriminator output bit can be used for BCID of saturated pulses in the downstream BCID logic.

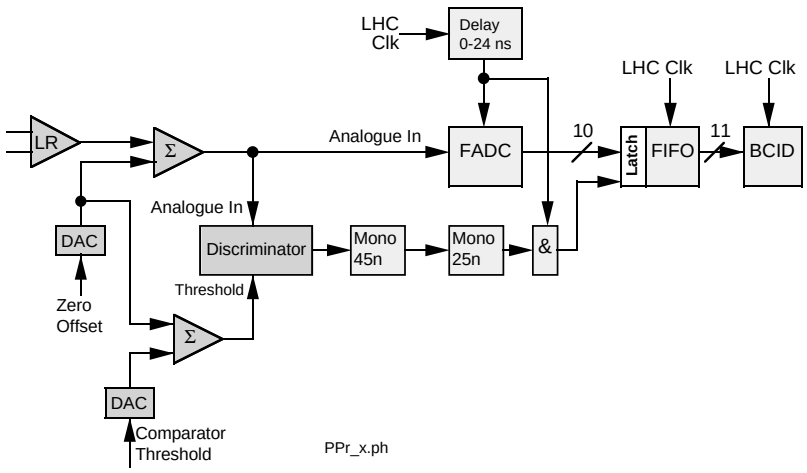

Figure 5-35 Variant of discriminator logic for the implementation of saturated-pulse BCID.

Assessment of the implementation and relative performance of the various saturated-pulse BCID schemes will continue, making use of more sophisticated simulations of the shapes of large pulses based on the properties of the analogue processing chain.

5.3.3 Timing strategy for BCID

In order to correctly interpret the physics interactions generated by a single bunch-crossing, all calorimeter deposits must be processed coherently, requiring that they be synchronized in time before their transmission to the trigger processors. This requirement dictates that the bunch-crossing results for all trigger towers must be synchronized to appear with the same latency with respect to the bunch-crossing which was the source of the particles generating the observed calorimeter deposits. As shown in the preceding sections, obtaining optimum performance from any of the proposed digital BCID algorithms depends on tailoring their parameters to the peaking-time and phase of the analogue pulses. The most important single parameter is the timing of the strobes which sample the pulses; for the FIR-filter scheme, the

filter coefficients must be matched to a known phase for the digitization of the pulse, and for the lookup-table and discriminator schemes the strobe must be placed away from a regime in which small changes in the pulse change the bunch-crossing of the BCID result.

The timing set-up must be performed individually for every tower in the system, since pulse parameters vary from tower to tower; peaking-time depends on the calorimeter parameters and pulse phase depends on flight time, both of which are a function of η.

To some extent, FIR-filter BCID and discriminator BCID have conflicting requirements in the placement of the strobe. The FIR-filter scheme requires that the peak of the pulse is of defined phase with respect to the set of digitized samples on which it operates, but the discriminator scheme requires that the strobe is accurately positioned with respect to the leading edge of the pulse, and the time at which the pulse peaks is largely unimportant.

If the FIR-filter timing criterion is met, the timing of the output is relatively unaffected by differences between the actual pulse peaking-time and that for which the filter was designed. However, the relative timing of the maximum FIR filter output and the source bunch-crossing is dependent on the pulse peaking-time. Since it is the latter which is required as the BCID result, the FIR-filter result can only be correctly used with a reasonably precise knowledge (to within 10 ns) of this peaking-time. This suggests that part of the timing calibration must be to determine the peaking-time for each tower with some accuracy.

For correct operation of the discriminator-BCID scheme, a digitization strobe must follow the crossing of the low-E_T threshold, sampling the consequent change of state of the discriminator output. The timing of the discriminator transition depends on the amplitude of the sampled pulse, occurring earlier as the pulse amplitude increases. The suitable reference point after which to place the strobe is therefore the time of the discriminator transition for the smallest pulse with which this algorithm is to operate, i.e. a pulse of peak amplitude equal to the high-E_T threshold. If the strobe is placed immediately after this reference point, then a small delay in the pulse can cause the following strobe to miss the discriminator transition, delaying the BCID result by 25 ns and throwing it into error by one bunch-crossing. Conversely, if it is placed too far after the reference point, then a small advance of the pulse, or even the amplitude-dependent time-skewing of the discriminator transition, will cause the detection of the comparator transition to occur one bunch-crossing too early, again throwing the result into error. The maximum tolerance of this scheme to shifts in the timing of the analogue pulse is clearly achieved when the strobe is placed ~0.5 bunch-crossing after the expected transition of a nominal pulse, of amplitude equal to the high-E_T threshold.

The above discussion shows that there are requirements to measure:

1. The peaking-time of a pulse of moderate E_T.
2. The timing of a low-E_T threshold transition for a pulse of ~250 GeV E_T.

The timing of the FADC digitization strobe can be shifted in steps of 1 ns over an entire bunch-crossing, and it is this degree of freedom which will be used to measure the relevant timing parameters, and to place the strobe in the optimum position. Requirement 1 can be met by using the LAr calorimeter calibration system to inject a pulse equivalent to a known E_T value into each tower, shifting the digitization strobe until a maximum digitization value is produced. Requirement 2 can be met by using a calibration pulse equivalent to 250 GeV E_T, and adjusting the strobe to determine when the low-E_T discriminator transition shifts by ±1 bunch-crossing from a ±1 ns change in strobe timing. The time difference between the first and second measurements can be used to determine the pulse peaking-time if the pulse shape is assumed.

5.4 References

5-1 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN / LHCC / 96–41, 1996.

5-2 *ATLAS Tile Calorimeter Technical Design Report*, CERN / LHCC / 96–42, 1996.

5-3 *Level-1 Calorimeter Trigger User Requirements Document*, version 1.1.0, ATLAS working Document, ATL–DA–ES–0001, April 1998.

5-4 U. Pfeiffer, Diplomarbeit Universität Heidelberg: *Analoge optische Signalübertragung vom Kalorimeter zum ATLAS Level-1-Trigger*, IHEP 96–16, 1996.
http://www.ihep.uni-heidelberg.de/Veroeffentlichungen/ps/pfeiffer.diplom.ps.gz

5-5 I. P. Brawn, *Bunch-crossing identification for the ATLAS level-1 calorimeter trigger*, Ph.D. Thesis, University of Birmingham, 1996.

5-6 J. Garvey and D. Rees, *Bunch-crossing identification for the ATLAS Level-1 Calorimeter Trigger*, ATLAS note DAQ–NO–051, 1996.

5-7 D. Rees, *Bunch-Crossing Identification Design Issues*, ATLAS note DAQ–NO–71, July 1997.

5-8 R. L. Chase et al., *A fast monolithic shaper for the ATLAS EM calorimeter*, ATLAS note LARG–NO–10, 1995.

6 Calorimeter trigger architecture and design

6.1 Architectural overview of the calorimeter trigger

In this chapter, we describe in some detail the proposed design of the Level-1 Calorimeter Trigger. Later, in Chapter 8, we expand on some of the technological and other issues involved in building this relatively compact but very complex system.

6.1.1 Interfaces to the calorimeters

The Level-1 Calorimeter Trigger receives analogue signals summed over multiple cells of the ATLAS electromagnetic (e.m.) and hadronic calorimeters as 'trigger towers' of moderate granularity. This was described in detail in Chapter 5. The two parts of the Level-1 Calorimeter Trigger — Cluster Processor (CP) and Jet/Energy-sum Processor — span different regions of the ATLAS calorimeters. The CP processes trigger tower inputs from the region $|\eta| < 2.5$; within this region all trigger towers are 0.1×0.1 from both electromagnetic and hadronic calorimeters. The CP trigger space is therefore 50×64 towers of e.m. and hadronic layers, for a total of 6400 towers to be processed. The Jet/Energy-sum Processor receives jet element inputs of 0.2×0.2 over $|\eta| < 3.2$; within the region $|\eta| < 2.4$ these are formed by pre-summing 2×2 groups of 0.1×0.1 towers, whereas beyond $|\eta| = 2.5$ only coarser granularity towers of 0.2×0.2 or 0.1×0.2 are available. The granularity change is handled by summing across the boundary to form one 0.3×0.2 element covering $2.4 < |\eta| < 2.7$, and another 0.3×0.2 element covering $2.9 < |\eta| < 3.2$. The resulting jet trigger space is 30×32 elements, with e.m. and hadronic layers combined. Finally, the missing-E_T and total-E_T calculations additionally use the signals from the forward calorimeters, which cover $\sim 3.2 < |\eta| < \sim 4.0$ and $\sim 4.0 < |\eta| < \sim 4.9$, with $\Delta\phi = 0.2$.

The calorimeter-trigger Preprocessor is the interface to these analogue signals. As well as performing the necessary signal processing to provide inputs to the digital trigger processors, the Preprocessor together with upstream patch-panels must also perform the reorganization of trigger towers as necessary for the architecture of the two processor systems. This includes regrouping towers for summing into jet elements, their allocation to serial links, the grouping of links for each processing module, and the duplication of serial links for towers or jet elements needed in more than one processor crate.

6.1.2 Division of tasks

The Level-1 Calorimeter Trigger is divided into two trigger processors, fed jointly by the Preprocessor over high-speed serial links. The Cluster Processor finds and counts electromagnetic clusters consistent with isolated electrons and photons, and similar isolated hadronic clusters coming mostly from tau decay. The Jet/Energy-sum Processor finds and counts collimated jets of particles, and measures the global quantities missing-E_T and total-E_T. Both of these systems receive their inputs from the Preprocessor, which for each trigger tower performs digitization, digital filtering to associate the broad pulses with a single LHC bunch-crossing (bunch-crossing identification, or BCID), pedestal subtraction, final E_T calibration, and noise thresholding. The Preprocessor also performs the summation of trigger towers to reduced-granularity jet elements required by the Jet/Energy-sum Processor. Trigger towers and

jet elements are transmitted from the Preprocessor to the Cluster and Jet/Energy-sum Processors, respectively, by high-speed serial links. The level-1 calorimeter trigger is shown schematically in Figure 6-1.

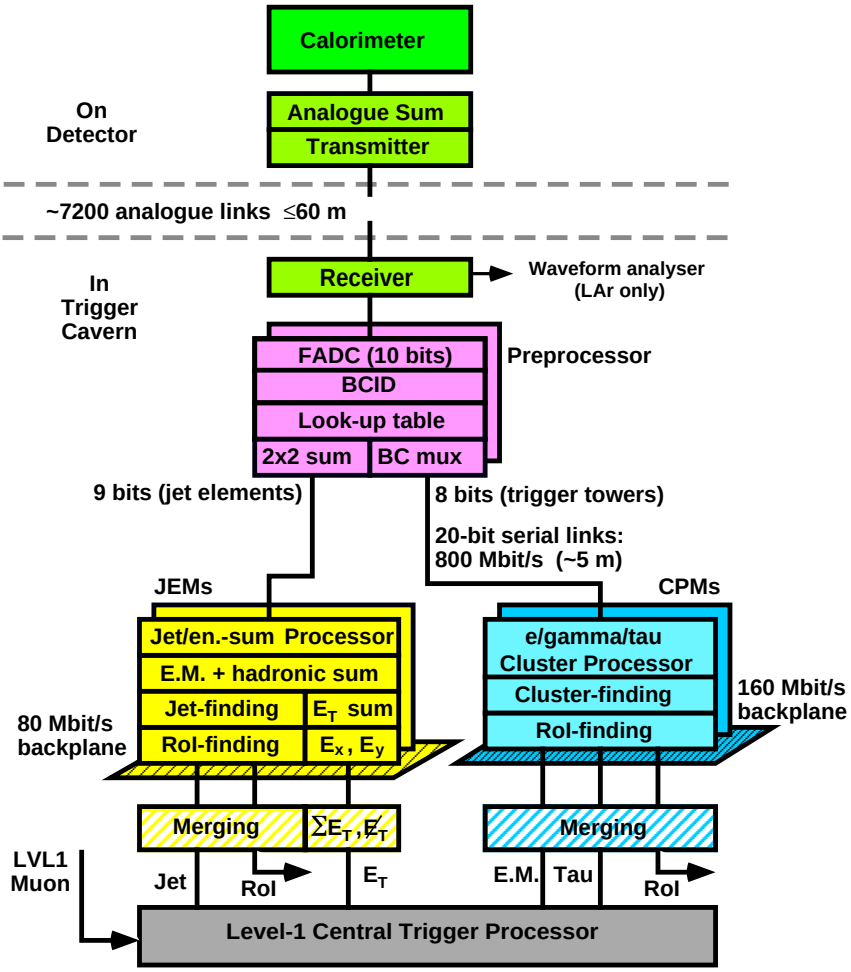

Figure 6-1 Overall architecture of the Level-1 Calorimeter Trigger.

6.1.3 Geometrical considerations

The trigger input space is divided into a number of overlapping, sliding η–ϕ subspaces, and signals from adjacent subspaces are shared between the devices which implement the trigger (ASICs, modules, crates etc.). The resulting fan-out of each trigger tower to a number of processor destinations is the defining characteristic of the calorimeter trigger, and the way in which it is handled is the predominant problem to be solved in achieving an optimal trigger solution.

The following convention will be used throughout to describe the regions processed by each kind of device in the trigger. Trigger subspaces are described as $m \times n$, where m is the number of towers in η and n the number of towers in ϕ; where it is necessary to refer to the number of layers, this is appended as a third dimension $m \times n \times p$ ($p = 1$ or 2). For describing groups of regions and the devices which process them the notation $a \times b \times (m \times n \times p)$ is used, meaning a assemblies of b devices, each $m \times n \times p$.

The basic window for the algorithms in the Cluster and Jet/Energy-sum Processors is composed of $4 \times 4 \times 2$ trigger towers or $4 \times 4 \times 1$ jet elements, respectively. Windows overlap, with neighbouring windows being displaced by one trigger tower or jet element in either the η or ϕ directions. In constructing a 4×4 window around any tower, one must examine two rows of towers on two adjacent sides, and one row of towers on the other two adjacent sides. Thus for any region of $m \times n$ overlapping 4×4 windows, the total environment required to perform the trigger algorithms on those windows is $(m + 3) \times (n + 3)$ towers, and likewise for jet elements. Two such neighbouring subspaces with m windows at the interface have a boundary $m + 3$ towers long and there are $3 \times (m + 3)$ towers which are required by both subspaces.

The CP and jet/energy-sum trigger spaces (see Section 6.1.1) are both divided according to a ϕ-quadrant architecture, whereby each $\pi/2$ azimuthal quadrant of the trigger space is processed by a number of identical modules within a single crate, the crate spanning the whole space in η. Each module within the crate covers a small part of the quadrant in η, but spans the whole quadrant in ϕ. The subspace processed by each module is an element in a one-dimensional array of subspaces within the region processed by the whole crate; the logical arrangement of processed regions is thus the same as the physical arrangement of processing modules, as indicated in the top half of Figure 6-2.

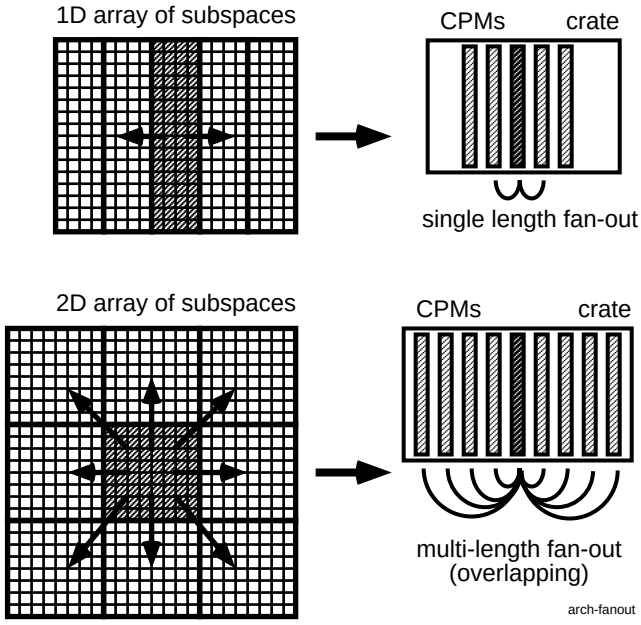

Figure 6-2 One-dimensional vs. two-dimensional fan-out on the backplane.

This yields the advantage that modules processing neighbouring regions in η are in adjacent slots in the same crate. Communication between modules within the same CP or jet/energy-sum crate is performed via a fast semi-serial backplane, and in the ϕ-quadrant architecture each backplane link is only one slot in length. This is a considerable simplification compared to a system which maps the trigger space as a two-dimensional array of subspaces within the same electronics crate, as indicated in the bottom half of Figure 6-2 and which was previously proposed [6-1], and evaluated in our demonstrator programme (see Chapter 7).

In order that the Preprocessor can provide inputs in parallel for both the Cluster and Jet/Energy-sum Processors, each quadrant is divided into the same subspaces for both. The modules in the Preprocessor cover the same subspace, although the e.m. and hadronic

calorimeters are handled separately. Thus two Preprocessor Modules (PPMs) — one electromagnetic and one hadronic — map onto each Cluster Processor Module (CPM) and Jet and Energy-sum Module (JEM). Note that because of the pre-summing to coarser granularity jet elements for the Jet/Energy-sum Processor, the JEM covers the same-sized region as an equivalent CPM, but has only 1/4 as many inputs. This universal mapping of modules within the calorimeter trigger is shown in Figure 6-3.

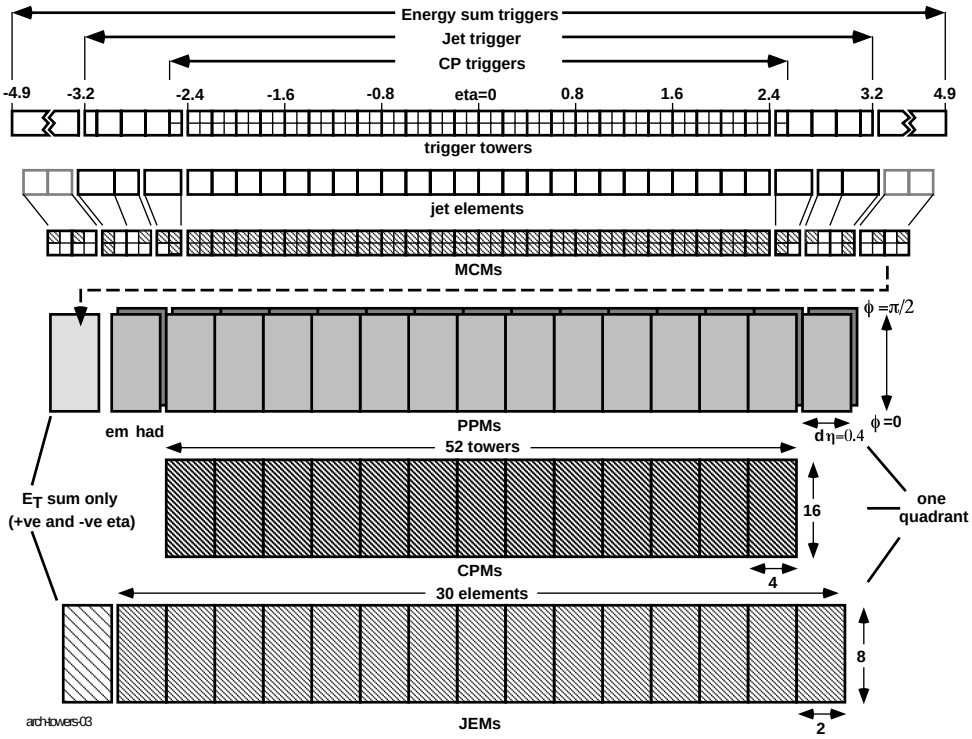

Figure 6-3 A possible layout of Preprocessor Modules, Cluster Processor Modules, and Jet and Energy-sum Modules.

6.1.4 Scalability of the design

The architectures of the Preprocessor, Cluster Processor and Jet/Energy-sum Processor are described in more detail in Sections 6.2, 6.3 and 6.4, respectively. Whilst the baseline architecture aims to divide the trigger space economically into quadrants, in order to minimize duplication of signals at the Preprocessor, some concerns remain about the resulting power density, cooling, available board space, and edge connections. These issues are to be addressed in the forthcoming design and prototyping phase, including the possibility of using lower-powered serial links, but a fall-back would be to lower the density of the system by re-partitioning. Such scalability is relatively simple in the ϕ-quadrant architecture, since it may be readily changed to ϕ octants by dividing each quadrant in half, assigning a smaller subspace to each module, and repeating the duplication of shared signals at the additional PPM/CPM/JEM edges.

This would lower the density of the Preprocessor by 50% and the trigger processors by about 40%, which should be more than sufficient to address the density concerns. An added advantage is the possibility of having all connectors at the rear of the modules, leaving the front panels free and making it possible to remove modules without disconnecting any cables. The disadvantage of this scaling is firstly the increase in cost of infrastructure — twice as many

modules, crates and power supplies — and secondly the doubling of the duplicated signals at the boundaries between crate subspaces. This increases the total number of serial links in the Cluster Processor by a further 25%, and hence the total power in the system. The comparative costs of these two architectures must be evaluated in detail, but the ϕ-quadrant remains the preferred option.

6.1.5 Results to CTP, level-2 and DAQ

The cluster and jet / energy-sum triggers do not actually make a yes / no decision on each event. Instead, the number of trigger objects passing each of the E_T threshold combinations for each class (electron / photon, hadron / tau or jet) is sent to the Central Trigger Processor (CTP) as the basis for a level-1 trigger decision. There are eight threshold combinations each for electron / photon triggers (a combination is a cluster threshold, an e.m. isolation threshold, and a hadronic isolation threshold), hadron / tau triggers (likewise), and jet triggers (a cluster threshold within a window of selectable size). There are also eight missing-E_T thresholds and four total-E_T thresholds. This is real-time data — a trigger multiplicity for each class of objects must be sent every 25 ns, with a fixed latency, from the bunch-crossing which generated those objects. The CTP requires one global multiplicity for each threshold combination within each object class, and so counts of objects from individual modules must be summed in merger modules before onward transmission to the CTP. The maximum multiplicity of objects at each threshold passed to the CTP is seven, requiring three bits per threshold.

Events which pass the level-1 decision are processed in greater detail by the level-2 trigger. Particularly at high luminosity, this system concentrates on regions-of-interest (RoIs) around each trigger object, in order to avoid reading out the full detector. The level-1 calorimeter and muon triggers must provide the coordinates of such RoIs to level-2 following a level-1 accept decision by the CTP, together with the thresholds which that object passed. All RoIs in an event must be passed to level-2, including those not used in the trigger decisions or those in excess of the seven-object maximum sent to the CTP.

Finally, level-1 data must be pipelined for readout to the ATLAS DAQ when a level-1 accept occurs, so that trigger information is included in the full event readout. Most of the DAQ load arises from the Preprocessor, which must at minimum read out the trigger-tower E_T values sent to the trigger processors so that it is possible to reconstruct where and how any trigger arose. For prescaled triggered events it will also capture the raw calorimeter trigger inputs over multiple bunch-crossings to allow signal timing and calibration to be done, as well as verification of the bunch-crossing identification (BCID). However, additional information is required from the processors themselves for a subsample of events in order to check the communication with the Preprocessor as well as the functioning of the processors.

All this requires considerable output bandwidth from even the smallest processing unit, although the overall sparsity of objects allows for a considerable reduction of data volume before communication with the level-2 trigger and the DAQ.

6.2 Preprocessor

6.2.1 Overview

The granularity of the Level-1 Calorimeter Trigger requires the collection of ~7200 analogue signals at the trigger electronics. Figure 6-4 shows the number of analogue input signals to the calorimeter trigger Preprocessor, and the processing steps before digital information is transferred to the trigger processing electronics. To accommodate our target of 64 channels on one module requires a high degree of integration, using ASIC technology. A more detailed view of the operations carried out on one tower is given in Figure 6-5. The stages of preprocessing are marked as A–K and will be referred to in the description that follows.

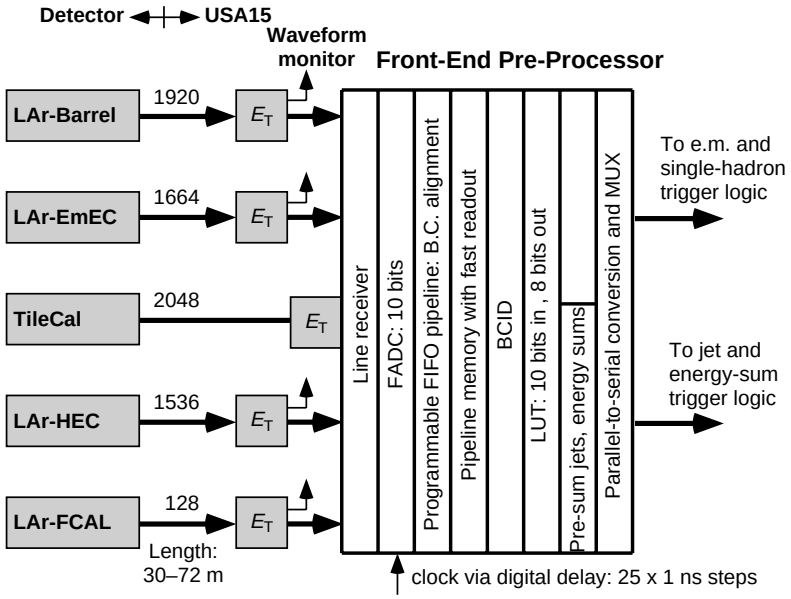

Figure 6-4 Overview of calorimeter signals and Preprocessor functionality.

6.2.1.1 Analogue signal input, digitization and timing

Each differential analogue signal is received and terminated in a line-receiver circuit (A). A signal from a trigger tower covers the range from 0 to 2.5 V linearly. This matches with transverse energy ranging from 0 to 256 GeV — larger energy deposits lead to saturation in the analogue stages upstream. All LAr signals have already been converted to E_T in the Receiver Stations, but Tile Calorimeter signals must be converted in these line receivers.

The analogue trigger-tower signal is digitized in an FADC (B) sampling at the LHC bunch-crossing frequency of 40 MHz. An FADC resolution of 10 bits is more than adequate for the purpose, giving an FADC step-size of 2.5 mV for one least-significant bit (LSB). It should be noted here that the expected noise on a liquid-argon (LAr) trigger-tower signal is almost double the step-size.

Signals from LAr calorimeters have a charge-cancelling bipolar shape. Tile calorimeter signals are unipolar, since they originate from photomultiplier devices. In either case the zero baseline should lie within the FADC digitization range. Shifting the zero level to a specified FADC value

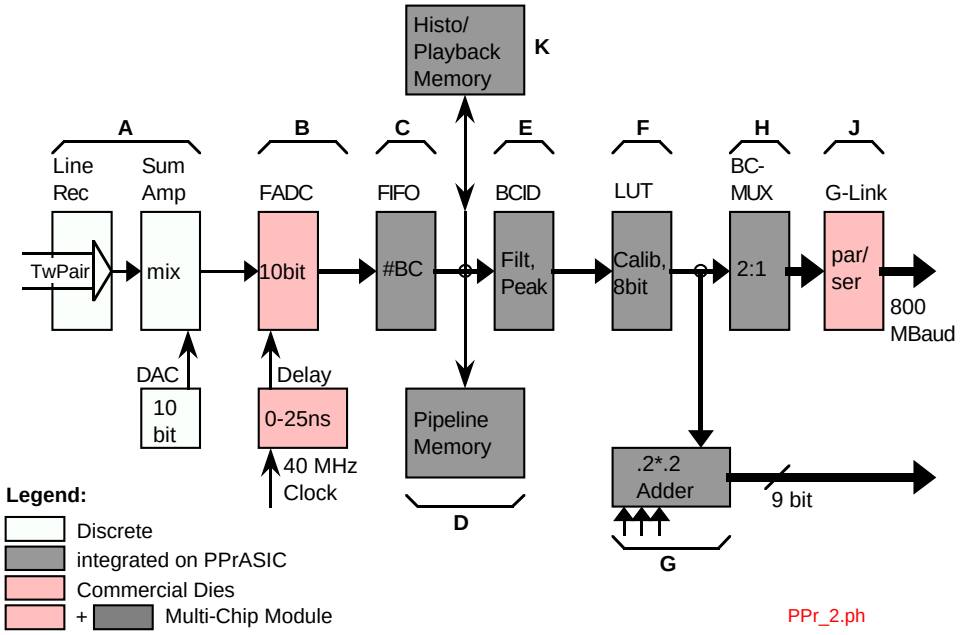

Figure 6-5 Functional diagram of the treatment of one trigger tower by the Preprocessor.

is achieved by a programmable DAC. Its resolution should be 10 bits for a voltage range of about 1.0 V (40% of the FADC range), yielding a step size of ~1 mV which is sufficiently smaller than the sensitivity of the FADC.

Time adjustment of the incoming analogue signal with respect to the sampling strobe is important to both the BCID efficiency and energy calibration. Optimum energy resolution is achieved by digitizing on the peak of the signal, so a programmable delay element is used for the phase adjustment between the analogue signal peak and the FADC clock.

The physical length of the cables transporting the analogue calorimeter signals to USA15 might differ by large amounts. Present installation plans foresee cable lengths from 40 m up to 60 m. Hence signal propagation times can differ by more than 100 ns, which corresponds to four LHC bunch-crossings. To align digitized data in time, a FIFO with a maximum depth of 32 locations and driven by the LHC clock is implemented in each channel (C). The FIFO depth is programmable, and is adjusted when setting up the experiment timing.

6.2.1.2 Readout of trigger input data

The raw data emanating from the FADC are read out on a prescaled sample of events, and therefore must be stored ‘on the fly’ in a fast scrolling memory (D), whose function is equivalent to a ‘level-1 pipeline’ on a subdetector. The memory stores data until readout is initiated by a level-1 accept signal. As the memory length must exceed the level-1 latency of 2.5 μs, a depth of 128 will be used, corresponding to 3.2 μs at 40 MHz. This comfortably exceeds the requirement. The memory is accessed upon receipt of a level-1 accept, when the relevant data for a given bunch-crossing (BC) are transferred to a derandomizing buffer memory for later readout to the DAQ system. Local readout in the Preprocessor is achieved by a high-bandwidth system called PipelineBus, the technical implementation of which is described in Section 6.2.5.2 below. The average level-1 trigger rate cannot exceed 100 kHz, but a peak rate of 8 MHz can be reached by an accepted event every fifth BC, since an artificial deadtime of four BCs is defined universally

for ATLAS. This peak rate allows a maximum of 125 ns to transfer level-1 data from the fast scrolling memory to the derandomizing buffer without introducing deadtime.

A second memory (K) of the same size serves two purposes:

1. Data playback: Recorded raw data or test patterns can be downloaded, allowing the system, including the full downstream trigger processing, to be run without input from detectors for setting-up and functional checking.
2. Histogramming: When the experiment is taking data, this memory could be used to accumulate statistical information for monitoring system performance, e.g. the frequency of FADC values for each channel would give direct feedback for fine adjustment of thresholds. The collection of statistics in hardware, instead of software, greatly reduces the amount of data that must be transferred to online computers.

6.2.1.3 Bunch-crossing identification

The bunch-crossing identification (BCID) circuit (E) comprises a finite-impulse response (FIR) filter, and a sliding peak-finder. It determines the maximum value by examining five consecutive FADC samplings, as described in Section 5.3.1. FADC samplings before and after the peak are set to zero. The output of the FIR filter is a measure of the transverse energy. Since the peak of the shaped pulse arrives at a fixed time after the event, the LHC bunch-crossing of the event is identified.

Calorimeter signals whose amplitude exceeds the upper limit of the FADC window result in a flat-top pulse. These saturated signals are handled in special logic, based on a lookup table (LUT) to derive their BCID. The result is a signal to the trigger at full scale, i.e. ~255 GeV, which will always produce some sort of trigger. If this purely digital solution proves to be viable for all shapes of saturated pulses, it will be implemented as described in Section 5.3.2.1.

Alternative solutions based on a discriminator, which uses the leading edge of the analogue pulse and the shaper-defined peaking time, are discussed in Section 5.3.2.3. This could allow BCID to be performed without exact knowledge of the full saturated pulse shape. The discriminator would be placed near the analogue line-receiver, and all other components could be implemented as dies on the multi-chip module (PPrMCM, see Section 6.2.4). A threshold for the discriminator would be defined by a DAC.

6.2.1.4 Lookup table

The 10-bit digital data are transformed by an LUT (F). Corrections for final transverse-energy calibration, pedestal subtraction, and application of a minimum threshold against noise are carried out in one step. It maps the 10-bit input to an 8-bit output of ~1 GeV LSB, as required by the trigger processors downstream. The use of an LUT is fast and allows arbitrary, nonlinear transformations, including the option to suppress noisy or faulty channels. The output data from the LUT are the basic ingredients used by the trigger algorithms, and *must* be read out to DAQ for every triggered event.

6.2.1.5 Transmission to Cluster Processor

Hewlett–Packard G-link serializers (J) are used in the present design to transport 8-bit data from the Preprocessor to the Cluster Processor. (Possible alternatives to G-link are discussed

Section 8.4.1.2.) This chip-set uses its own transmission protocol to reconstitute the clock and the corresponding data as a parallel word in the receiver chip. For distances of a few metres, electrical transmission on coaxial cable is feasible. The G-link provides two frame modes, allowing either 16-bit or 20-bit transmission. A further bit can be used either for the G-link's own internal error detection, or is available as a data bit. We have considered the possible use of two data rates for transmitting trigger-tower data, sending either two or four 8-bit trigger towers per serial link on each bunch-crossing. The resulting data rates are 640 Mbit/s of actual E_T data (800 MBd with protocol bits, 960 MBd if all possible bits are used, e.g. in the scheme described below), and 1280 Mbit/s (1600 MBd with protocol bits), respectively. Our choice at present is for the lower speed, since G-link operation at that speed is more robust and well inside the specification (see section Section 8.4.1.1), electrical links can be longer, and by using the properties of BCID as described below we can still transmit four trigger towers per link.

By the nature of the BCID algorithm (FIR filter and peak-finder), empty time-slices are created before and after a digitization identified as a peak value. The following time-slice may therefore be used to transmit the value of a neighbouring trigger tower. This scheme is called 'bunch-crossing multiplexing', or BC-MUX. Since 8-bit trigger-tower data are transmitted, use of the 20-bit frame mode allows the addition of the necessary flag bits to indicate which of the two possible bunch-crossings each byte of data is associated with. A format such as the one shown in Figure 6-6 can be used; the method for encoding the data and flag bits for the two successive trigger-tower bytes will be discussed below in Section 6.2.1.6.

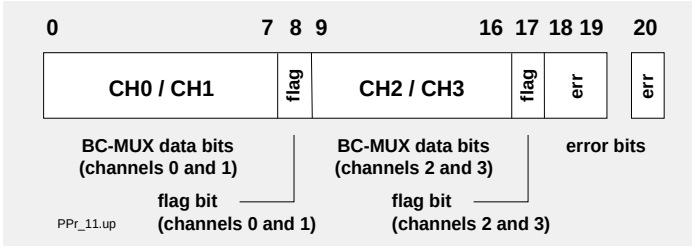

Figure 6-6 G-link data format for BC-MUX transmission.

In addition, two error (parity) bits are generated and appended to the data. Although the use of only one error bit per byte is not a powerful method for finding single errors, monitoring the error rate using these bits in the Cluster Processor will help to identify failing links, and thus to minimize the rate of false triggers. The BC-multiplexer (H) will perform the task of serialization, multiplexing, and generation of the flags and error bits. Thus, four trigger towers are effectively transmitted by one G-link to the Cluster Processor Module concerned — a factor-2 saving in cost and power.

6.2.1.6 Bunch-crossing multiplexing

We have considered two schemes for encoding two trigger towers, *A* and *B*, into a single data-stream two BCs long. These result in different conventions for the meaning of the flag bits. The schemes are:

- **Fixed**: Over two BCs, i and $i + 1$, always send the (non-zero) value from *A* first and the (non-zero) value from *B* second. The flag bit on each cycle indicates whether the relevant value came from BC i or $i + 1$. There is a fixed latency of 1 BC + d at the transmitting end, where d is the processing delay.

- **Variable:** The first non-zero tower out of the pair is sent out immediately, with a flag bit indicating whether it was *A* or *B*. If both are non-zero, send *A* as a default. On the second BC the other tower is sent, with the flag bit indicating whether it was in the same BC as the first or one BC later. Note that the two successive flag bits give different kinds of information. There is a fixed latency of *d* at the transmitting end.

The possible combinations of inputs and outputs over BCs i and $i+1$ are summarized in Table 6-1. *X* and *Y* indicate non-zero data in towers *A* and *B*, with the flag bits shown after the commas.

Table 6-1 Coding of BC-MUX for all possible combinations of data in two trigger towers.

	Inputs				Transmitted				
					Fixed		Variable		
Case	A(i)	A(i+1)	B(i)	B(i+1)	F(i)	F(i+1)	V(i)	V(i+1)	V(i+2)
1	0	0	0	0	0, 0	0, 0	0, 0	0, 0	
2	*X*	0	0	0	*X*, 0	0, 0	*X*, 0	0, 0	
3	0	0	*Y*	0	0, 0	*Y*, 0	*Y*, 1	0, 0	
4	0	*X*	0	0	*X*, 1	0, 0		*X*, 0	0, 0
5	0	0	0	*Y*	0, 0	*Y*, 1		*Y*, 1	0, 0
6	*X*	0	*Y*	0	*X*, 0	*Y*, 0	*X*, 0	*Y*, 0	
7	*X*	0	0	*Y*	*X*, 0	*Y*, 1	*X*, 0	*Y*, 1	
8	0	*X*	*Y*	0	*X*, 1	*Y*, 0	*Y*, 1	*X*, 1	
9	0	*X*	0	*Y*	*X*, 1	*Y*, 1		*X*, 0	*Y*, 0

Note that cases 6 and 9 are the same situation with a one BC delay. The fixed scheme sends two different patterns, whereas the variable scheme sends the same pattern 1 BC later. There is a fixed latency of 1 BC + *e* at the receiving end (*e* is the decoding time) in both schemes; one can never use the first value until it is known what the other tower of the pair was doing in the same BC, and this is only known on the following received BC.

In the fixed scheme, the transmitter and receiver must be synchronized, since the sequence of transmitted values determines which is interpreted as *A* and which is *B*. The variable scheme is self-synchronizing, since the first non-zero value after a transmitted zero is always explicit — the first flag bit says which tower was sent, and its value is always assigned to the next BC which the receiver is processing. The corollary is that the meaning of the flag bit changes between consecutive non-zero BCs. Each BC needs to be interpreted with reference to the preceding BC, but any transmitted zero data resets the link.

The variable scheme is preferred, because of its lower latency and self-synchronizing feature.

6.2.1.7 Summing and transmission to Jet/Energy-sum Processor

The 8-bit digitized values at this point represent transverse energy deposits within the standard trigger space ($\Delta\eta \times \Delta\phi = 0.1 \times 0.1$) for the identification of electromagnetic and hadronic clusters.

The trigger algorithms for objects spanning larger areas — jets, missing-E_T and total-E_T — use a coarser granularity (see Sections 4.4.1 and 4.5) of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. To simplify the architecture downstream and reduce the number of links, pre-summing is performed in the Preprocessor (G). These 10-bit sums of four trigger towers can be truncated to 9 bits for e.m. and hadronic jet elements separately without loss of trigger resolution, as shown in Sections 4.4.7 and 4.5.2.2. The 20-bit frame mode of the G-link is again used for the transmission of these 0.2×0.2 sums, but because BCID has already been performed independently on the four inputs to these sums, bunch-crossing multiplexing is not possible and only two pre-sums may be transmitted by one G-link to the Jet/Energy-sum Processor. However, the use of 9-bit sums leaves two bits available for error detection.

6.2.1.8 Integration of components

The design for ATLAS integrates blocks (C) to (H) into one ASIC, called the Preprocessor ASIC (PPrASIC). Furthermore, the integration of several channels in parallel into one Multi-Chip Module (MCM) is envisaged in order to achieve a density of 64 trigger towers on a single printed-circuit board (PCB). Suitable FADCs are commercially available in die form, as is the HP G-link. This Preprocessor MCM (PPrMCM) will have a relatively small pin-count, but thermal management will be an important engineering consideration as the power consumption of each G-link transmitter die is approximately 2.0 W.

The specification for a standard 9U VME crate is currently still open, but the number of crates needed to house the ~7200 channels will be at least eight.

6.2.1.9 Module layout

The η–ϕ space covered by the ATLAS calorimetry can be mapped onto crates, PCBs, PPMs, and ASICs as shown in Figure 6-7. The primary division is into 16 calorimeter sectors (4 endcap, 4 + 4 barrel, 4 endcap), two per crate. This requires 64 channels on a single-width Preprocessor Module (PPM). Thus a sector is mapped onto eight PPMs, occupying half a crate. A possible layout is given in Figure 6-8.

The current design requires 24 HP G-link transmitter chips per PPM (16 to the Cluster Processor and eight to the Jet/Energy-sum Processor), corresponding to a power consumption of 48 W from these ECL components. It is reasonable to assume that the remaining TTL-based electronics on a PPM will not require power in excess of 100 W. With 16 PPMs installed in a crate, the power dissipation may be as high as $\sim 150 \times 16 = 2400$ W per crate. Although high, this value lies within manageable limits for d.c. power supplies and cooling capabilities.

Although the aim of our design is to fit 64 trigger towers onto one module, there are still many technical issues that might prevent this. The crucial issue is PCB real-estate to accommodate Preprocessor electronics for 64 trigger towers in a single-width module. A crate height of 9U is definitely required. Allocation of one VME crate per sector would result in a more conservative configuration of 32 channels per PPM. However, 16 crates would then be needed to house the entire system. As crate space is a major cost factor, housing the Preprocessor in eight crates is a desirable goal, although rescaling to a 16-crate system is a possible fall-back option if the higher density cannot be achieved.

A possible mapping of the calorimeter sectors onto crates is shown in Figure 6-9. The massive cabling plant, due to the large number of signals, requires carefully organized routing. The

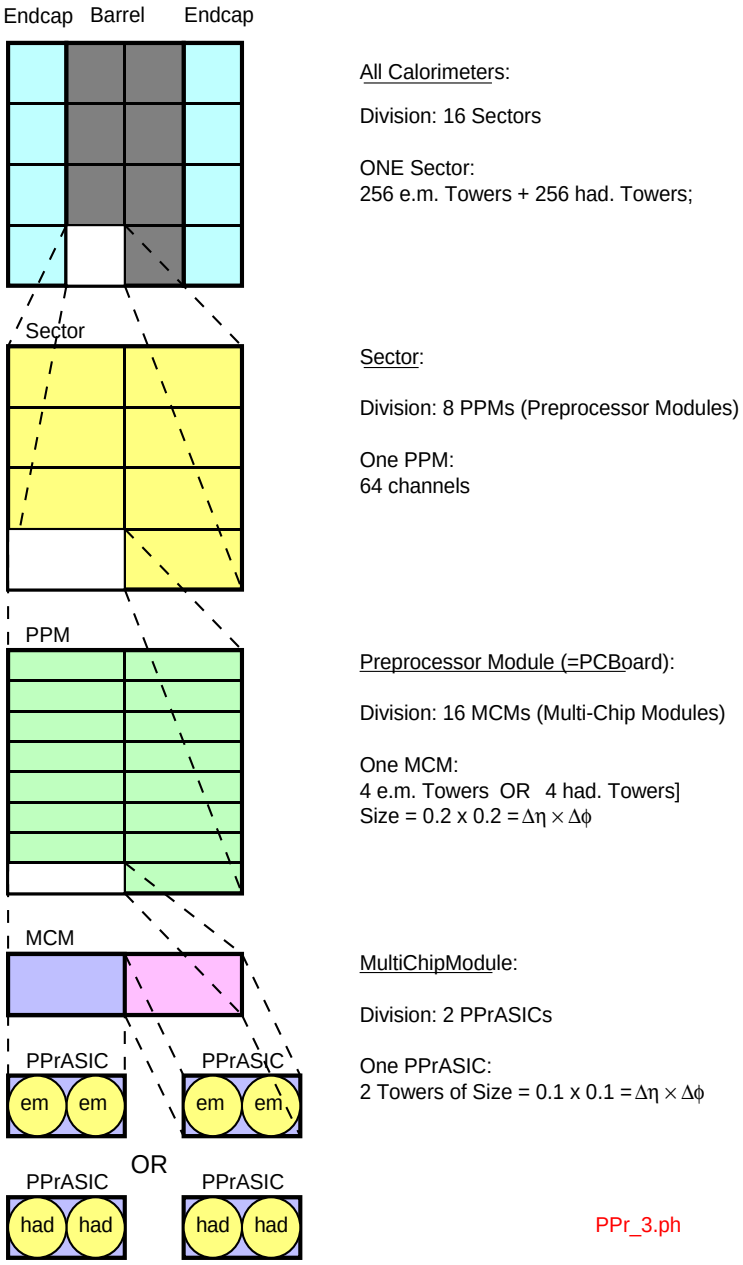

Figure 6-7 Mapping of ATLAS calorimetry space onto Preprocessor crates, PCBs, MCMs, and ASICs.

signal flow is directed into Receiver Stations (liquid argon) or patch-panels (tile) from the rear, emerging reordered on the front panels. Input cabling to the Preprocessor arrives at the front of the crates, and the high-speed G-link cables to the downstream trigger processor systems leave the system from the rear. In a possible rack layout for the entire Level-1 Calorimeter Trigger (Section 8.5), the downstream trigger processors could be placed in the same row between the two halves of the Preprocessor. Another scenario could use a second row of racks for the Cluster, Jet / Energy-sum and Central Trigger processors. The configuration shown in Figure 6-9 is still preliminary, and changes are likely in order to optimize the *total* length of cable, and hence latency, including both input cables and the output links to the CTP.

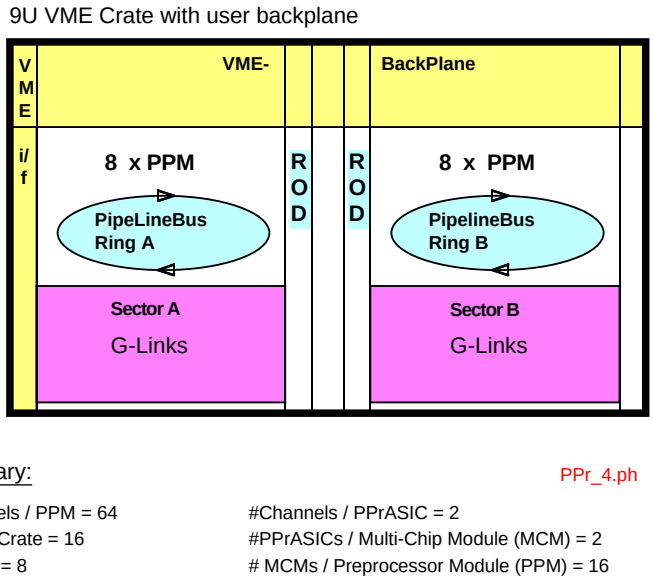

Figure 6-8 Layout of a Preprocessor crate.

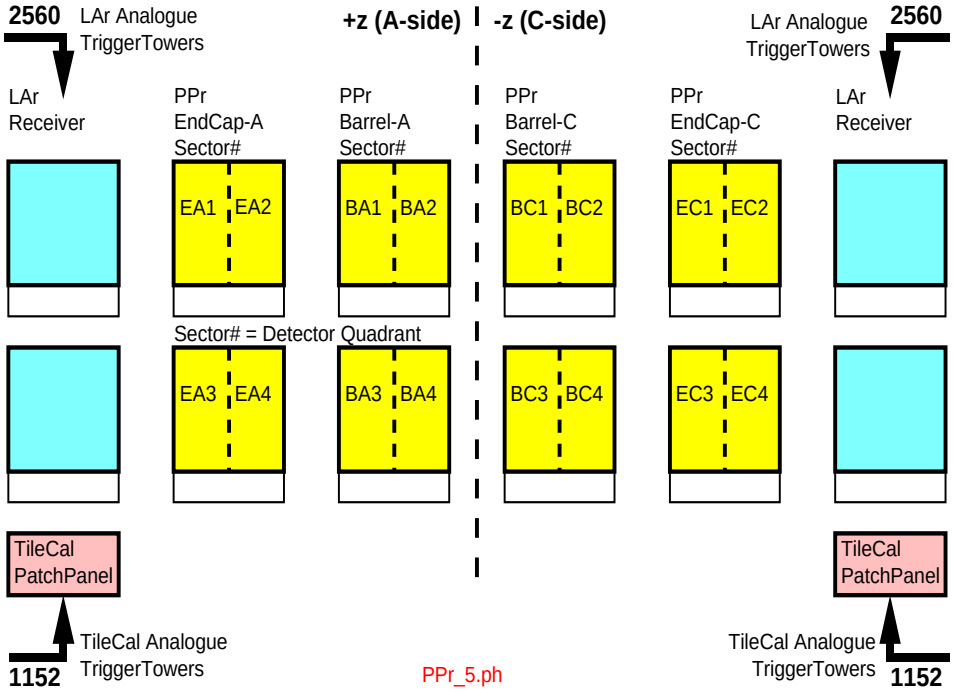

Figure 6-9 A possible mapping of calorimeter sectors onto Preprocessor crates.

It should be noted that maintenance work on the level-1 trigger will carry on through the full data-taking period of the ATLAS experiment. A clear and straightforward organization of the trigger electronics will make this easier, particularly in the later years of operation.

6.2.2 The Preprocessor Module (PPM)

The PPM contains all the signal processing components described earlier. Considerations of manpower resources for design, testing, and maintenance favour the development of only one module type. This is justified since each input channel is handled in the same way. Hence no special modules are considered. Differences stemming from detector-specific properties (e.g. adding overlapping trigger towers, etc.) are handled, if possible, upstream of the Preprocessor (see Sections 5.2.2.3, 5.2.3.2, and 5.2.3.4). Detector-specific parameters of the BCID and LUTs, however, are programmable and can be downloaded for each channel individually. This requires the design and construction of 128 PPMs (plus spares).

6.2.2.1 Functional description

A functional diagram projected onto a possible PCB is shown in Figure 6-10. The module receives 64 analogue signals. The Preprocessor imposes only a four-channel grouping at the input level, since it forms 0.2×0.2 sums in one kind of calorimeter for the Jet / Energy-sum Processor. After the differential line-receiver, a redistribution on-board is needed to allocate the signals in groups of four to the PPrMCMs. This is necessary because the 0.2×0.2 sums are built on each MCM. A preliminary layout study shows that the on-board space requirement for discrete line-receivers and signal redistribution is not negligible, and a vertical PCB slice of approximately 60 mm width will be needed for a design with minimized crosstalk.

The signal processing blocks contained in each PPrMCM are described later. Important at the module level are the outputs provided for the 0.1×0.1 data paths. Four channels are merged onto each serial G-link line (2×8 bits, plus flag and error bits, BC-multiplexed) leaving the PPrMCM, which must transport the data to a connector on the PCB edge. These signals are used exclusively in the Cluster Processor. The second data path brings 9-bit data (0.2×0.2 sums, truncated) out of the PPrMCM package, pairs of which are immediately serialized in an external (i.e. packaged) G-link transmitter to be transmitted to the Jet / Energy-sum Processor.

High-speed outputs are grouped at the edge of the PPM (lower section of backplane) according to the needs of the downstream systems. Architectural requirements demand the duplication (fanning-out) of a small number of signals — four to the Cluster Processor and three to the Jet / Energy-sum Processor. This could be accommodated, if there is not enough space on the PPM, on an auxiliary PCB, which is inserted from the rear into the crate.

The PPM must be supplied with external signals for synchronous operation. Most important is the LHC clock signal (40 MHz) from the Timing, and Trigger and Control (TTC) system. These signals are distributed optically and must be converted and decoded in a TTC receiver component on the PPM. Other control signals such as the level-1 accept —- used to initiate the readout of raw level-1 data — also arrive on this path.

Raw data, stored in the pipeline memories, are read out upon receipt of a level-1 accept. This is achieved by a token ring system called PipelineBus. The system is described later in detail, but its implementation on the PPM requires allocation of space for components and signal routes, and for connectors located in the middle section of the back edge.

A minimal VME connection is foreseen for parameter loading and basic control purposes. The VME interface and its connector for a 16-bit data bus occupy the top section of the backplane. Readout and loading of playback data, as well as set-up constants, are managed on the PPM by

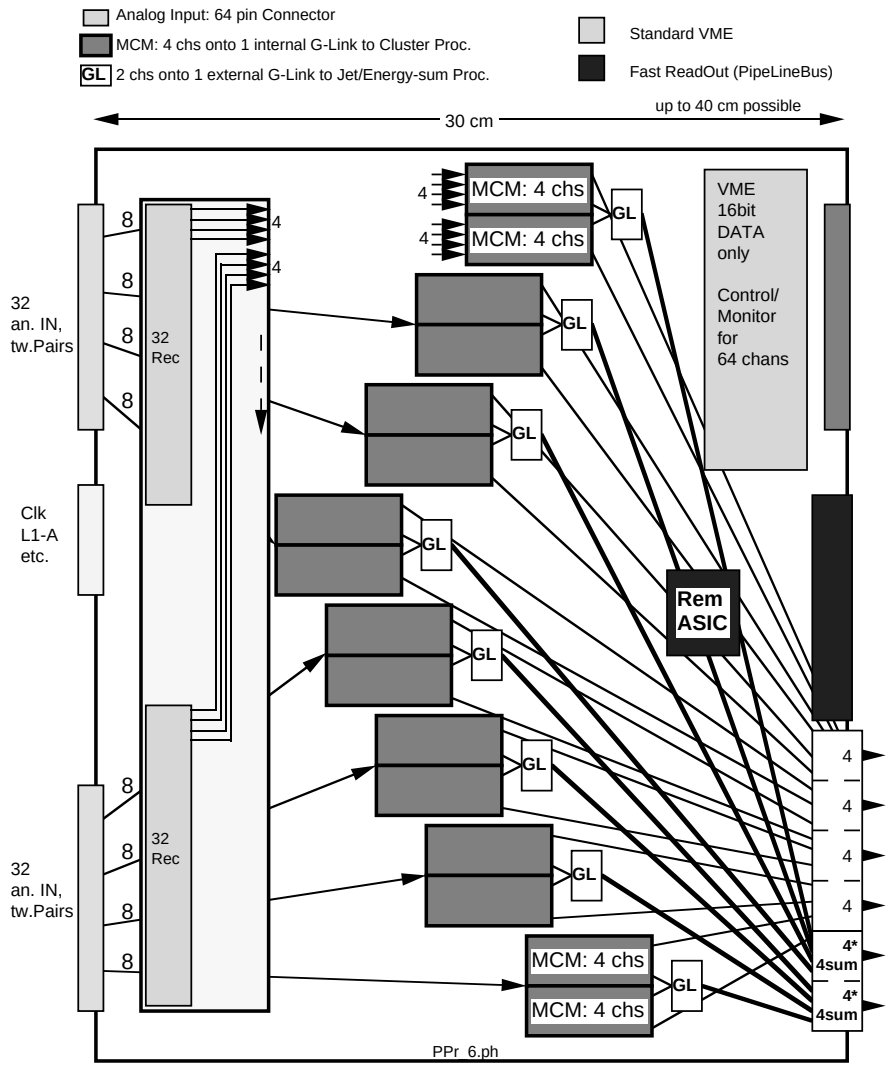

Figure 6-10 Layout of Preprocessor Module.

an ASIC called the REadout Merger ASIC (RemASIC), which links the local memories of the PPrASICs to VME or to the PipelineBus system.

Several sets of parameters must be pre-loaded into a PPM. The operations for each channel are:

- Load the digital-to-analogue converter (DAC) to set the FADC pedestal.
- Set the programmable fine delay within an LHC clock cycle to adjust the FADC strobe to the pulse-peak time.
- Set the FIFO depth to equalize signal propagation from the detectors in multiples of LHC clock cycles. Digital data emerging from the PPM must be aligned in time for downstream processing.
- Load coefficients for the BCID filter algorithm and for the LUT which is needed to do BCID for saturated pulses.
- Load the LUT for energy calibration (optimized for jets), a minimal threshold against noise, pedestal subtraction, and the conversion of a 10-bit FADC result to an 8-bit energy value needed in the Cluster Processor.

The VMEbus is used for loading parameters into the PPMs, and for technical check-out of the system. Stand-alone configuration and readout of a PPM, e.g. using JTAG for hardware diagnostics, must be implemented.

6.2.2.2 Components on the PPM

The front-panel connectors have to be of a rugged type, with clamps securing the input cables. It is not yet decided whether the line receivers and the signal redistribution should be implemented directly on the mother-board or on a daughter-board to save space.

The packaged PPrMCMs (Figure 6-11) will have considerable physical size. Their on-board placement is governed not only by signal routing requirements, but also by considerations of heat removal. The main heat source in the present design is given by the ECL-based G-links, either in die form on the MCM or in packaged form for the signals to the Jet/Energy-sum Processor. Special care has to be taken with the routing of the high-speed signals to the output connector. As bit-stream rates are approaching 1 Gbit/s, specially designed strip-line routing is probably needed to ensure proper functioning and minimal crosstalk. Fully synchronous operation is only possible with a clock distribution system compensated for propagation time across the entire module. There is a single interface to the external timing system on each PPM in the shape of the TTCrx receiver component.

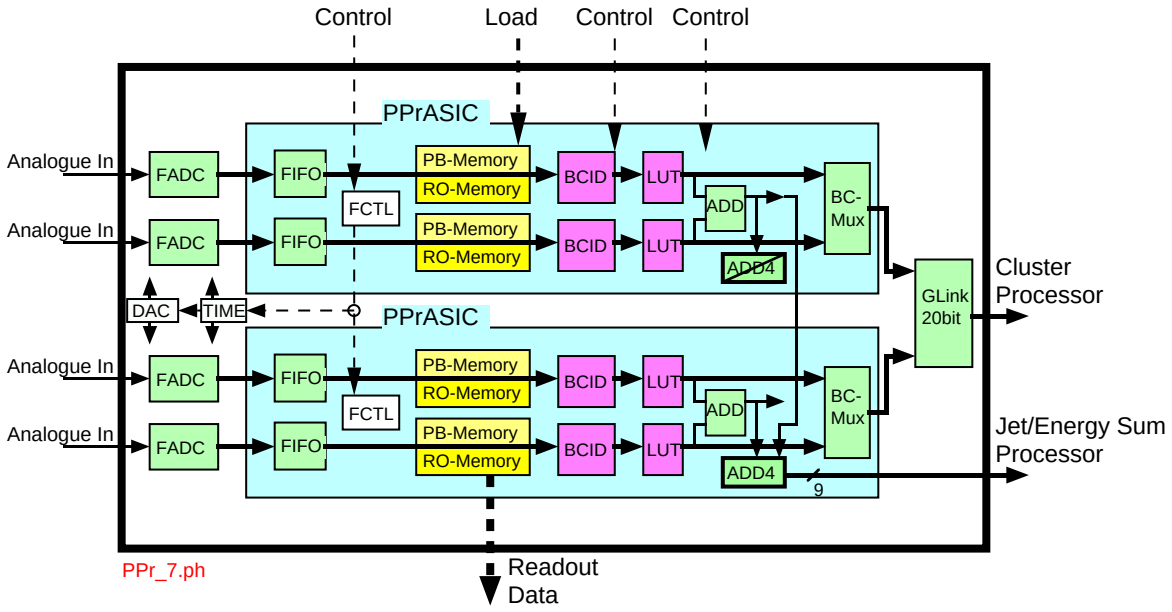

Figure 6-11 Block diagram of the Preprocessor MCM.

FADC
There are many 10-bit FADCs on the market, but there are also many constraints on our choice. For example, a wide analogue input range of 2.5 V (LSB of 2.5 mV), sampling at 40 MHz, TTL output signals, and low power dissipation are essential. No final choice has yet been made, but one suitable candidate is the Burr-Brown ADS 823 device. This device is a pipelined 10-bit ADC, needing five clock cycles for conversion. It has low power consumption (265 mW), TTL I/O, and a maximum input range of 2.5 V peak-to-peak. The maximum conversion rate is 60 MHz.

Timing, Trigger, and Control system (TTC)
Each PPM contains a TTC receiver chip (TTCrx), which is used to extract the LHC clock, the level-1 accept signal, and the bunch-crossing and event counter numbers (or the corresponding reset signals).

Programmable delay element (TIME)
This component, developed by the ECP/MIC group at CERN, provides four calibrated delays for digital signals. The delays can be programmed independently in 1 ns steps up to a maximum delay of 25 ns. The chip uses a 40 MHz clock as the timing reference and can be programmed using the I^2C standard. The maximum time jitter is 150 ps rms. The supply voltage is 3.3 V and the power consumption at 40 MHz is about 73 mW. Output signal levels are 3.3 V (low-voltage TTL).

DAC
The Analogue Devices AD7804 is a quad 10-bit digital-to-analogue converter with serial load capabilities, used for the base-line adjustment of the analogue trigger signal. The analogue outputs are possibly sent off the MCM to be added to the trigger signals before being digitized in the FADCs.

HP G-link
The Hewlett–Packard HDMP-D022 is a recently-introduced parallel-to-serial converter with TTL I/O; our tests to date have used the older ECL version. The transmitter device is part of a chip-set that performs all the tasks of encoding, multiplexing, clock extraction, demultiplexing and decoding. The 20-bit input data words are serialized at 40 MHz, and a 21st bit is used as an even or odd frame indicator for expanded error indication. The serial (user) data rate will be 800 Mbit/s and the total baud rate, including an additional four control bits, will be 960 MBd. One of the disadvantages of this device is its high power dissipation of 2.0 W, making it necessary to use a thermal cut-out on the PPrMCM substrate.

6.2.3 Preprocessor ASIC

The Preprocessor ASIC (PPrASIC) performs the main preprocessing functions on two trigger towers, as shown in Figure 6-12. It incorporates timing synchronization, BCID, calibration LUT, and bunch-crossing multiplexing. In addition, it includes readout components — pipeline memories, derandomizer buffers and a serial readout interface. It will be a semi-custom design, developed in the ASIC laboratory at the University of Heidelberg. It receives two digital inputs as 10-bit parallel data words from the FADCs every 25 ns. After processing, data from two trigger towers are BC-multiplexed into eight bits plus one flag-bit and one parity-bit each. These 10 bits fit into one half of a G-link data word. A JTAG interface offers the ability to test connectivity and bonding inside the MCM, and internal registers can also be read for test purposes. A detailed description of this essential component is given below.

At the input, 10-bit data words from each channel are stored in a FIFO with a programmable depth of up to 32 time slices (0.8 μs). To ensure stable data, storage will occur on the rising or falling edge of the LHC clock depending on the position of the programmable delay element for the FADC clock. The depth of the FIFO is more than sufficient to compensate for the cable length differences of up to 20 m between trigger towers. A BCID algorithm is then applied to data from each channel. A different algorithm is used for saturated and non-saturated pulses (see Section 5.3).

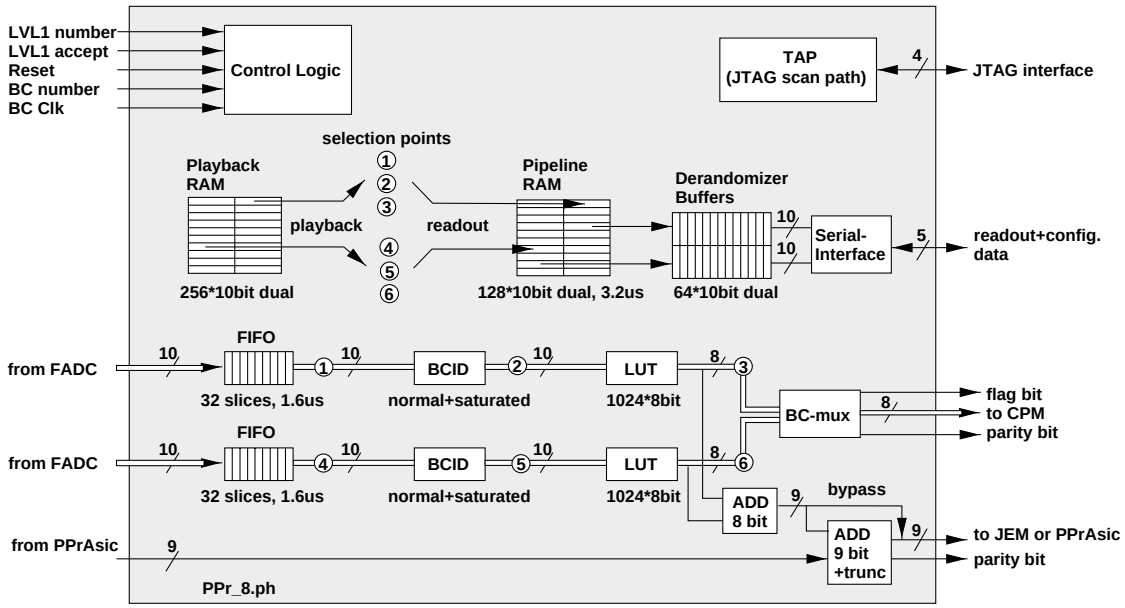

Figure 6-12 Block diagram of the Preprocessor ASIC.

After BCID a programmable 1024×8-bit lookup table (LUT), which is used for energy calibration, pedestal subtraction and application of a noise threshold, reduces the number of bits to eight. The data can bypass BCID for test purposes. To form 0.2×0.2 jet elements, pairs of towers are first combined in an 8-bit adder, and the results from the two PPrASICs on an MCM are then combined in a 9-bit adder (which is present but unused on half the PPrASICs).

Readout data can be captured at three points, which in Figure 6-12 are labelled 1, 2 or 3 for the upper and 4, 5 or 6 for the lower channel. Data from those points are copied with the full bunch-crossing frequency of 40 MHz into a pipeline memory, which has a length of 128 time slices (3.2 μs). These data are stored for up to 2.5 μs until the level-1 trigger has made a decision, when the corresponding data are copied from the pipeline memory to a 10-bit, 64 deep derandomizer buffer. With data from five time slices and a 10-bit header stored per event, the derandomizer capacity allows up to 10 events to be read out via the serial interface. The header contains four bits from each of the level-1 and bunch crossing numbers, which are used as a synchronization cross-check. These numbers are finally extended to their full size by the Readout Driver (ROD) module.

Data can be read out from the PPrASIC over a simple serial interface, which is also used to write configuration data. Since each PPrASIC provides independent input and output pins for the serial data, it is possible to operate several PPrASICs in a daisy-chain configuration. The serial interface is controlled by a clock and a control signal given as input to the PPrASIC, which operates as a slave. It uses a synchronous protocol for maximum throughput.

A very useful test feature for the trigger is a playback memory. The memory of size 256×10 bits can inject pre-loaded trigger data into the same points as mentioned above for readout purposes. For in-circuit testing at board level and at MCM level, a test access port (TAP) provides the possibility to check wire connections and to spy on internal registers.

The two custom ASICs required for the Preprocessor (PPrASIC, and the RemASIC described below in Section 6.2.5.1) are purely digital designs. They will be completely described with VHDL and synthesized in standard cell logic. In addition, memory blocks have to be integrated.

A sub-micron CMOS process will be used for manufacturing. This design methodology has been successfully employed for several prototype ASICs using a 0.7 µm process by ATMEL-ES2. Since ATMEL has stopped doing multi-project wafer runs, a different process must be chosen for the future. Candidates are the 0.6 µm process of AMS or the 0.5 µm process of Alcatel–Mietec. Both are accessible under EuroPractice for prototypes, but full production will require an engineering run.

6.2.4 Preprocessor Multi-Chip Module

6.2.4.1 General description

The driving force to use MCM technology for the Preprocessor is the large number of channels we wish to process in each PPM, and the large number of semiconductor devices needed for this processing. MCMs represent a technique where a very high internal pin count can be implemented on the substrate, with the package having a much reduced external pin-count. This section starts with a functional description of the MCM, followed by a more technical description of its production technique. More details of its development are given in Section 8.4.4.2.

The Preprocessor MCM (PPrMCM) shown in Figure 6-13 will receive and fully process analogue information from four trigger towers, sending the resulting data serially to the Cluster Processor. Data for the Jet/Energy-sum Processor are merged and serialized in an external G-link. This section identifies the stages of preprocessing with the corresponding dies on the PPrMCM. A channel-oriented description of the preprocessing was given in the system overview (Section 6.2.1).

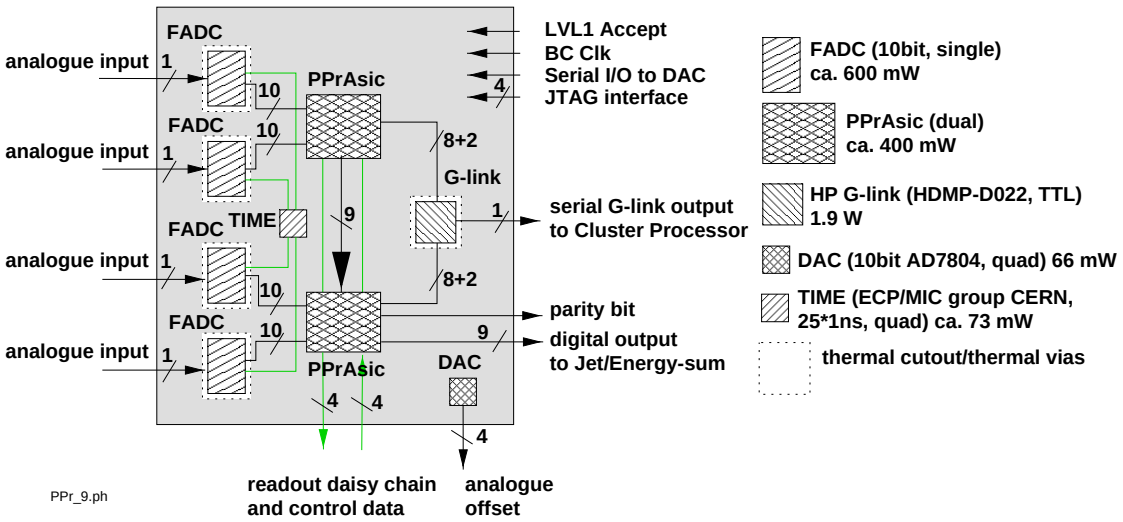

Figure 6-13 Layout of Preprocessor MCM.

The PPrMCM contains both analogue and digital devices. In total, there are nine dies: four FADCs, one DAC, one HP G-link, one programmable delay element and two PPrASICs. Their functional roles were described earlier.

Before an analogue trigger-tower signal arrives inside the PPrMCM, it is mixed with the analogue output of a DAC. This is done in the line receivers to adjust the pedestal and to shift the full signal into the digitization range of the FADC. The DAC itself is mounted on the PPrMCM to keep the routing of the digital inputs internal. Special attention will be paid to the analogue parts of the PPrMCM. Separate planes for ground and power should isolate the signals from the digital parts of the MCM to reduce the influence of electromagnetic interference (EMI) and keep crosstalk to a minimum. Potentials such as the top and bottom reference voltage of the FADC will be handled with care to ensure a good signal-to-noise performance.

Wire connections inside the PPrMCM are mainly channel-oriented. Only a few signals, such as the serial interfaces of the PPrASICs and the 9-bit data bus (the pre-summed half of a jet element), connect dies of different channels. The final truncated 9-bit sum for a full jet element (0.2×0.2), formed in every second PPrASIC, goes off the PPrMCM to an external G-link. The parity bit for the 9-bit jet element is provided by the summing PPrASIC.

The G-link transmitters are operated in the so-called ‘single frame’ mode, which serializes 20-bit words at a data rate of 800 Mbit/s. The high-speed signal from the internal G-link will be treated with special care, e.g. the wire length to the output pin of the PPrMCM will be kept as short as possible. Dedicated G-link power lines protect the rest of the MCM against spikes.

A large reduction in power consumption is achieved by the PPrASIC sending bunch-crossing multiplexed data to the internal G-link, thus requiring only one G-link to transmit data from four trigger towers to the Cluster Processor. Each PPrASIC data word contains 8-bit data, one parity bit and one flag bit. The flag bit is used to demultiplex the data at the receiving end of the link.

All dies inside the PPrMCM are identified in Figure 6-13. An estimate is given of the required area and the number of bits per data word transmitted between each die. The drawing is roughly to scale; the area estimation is based on the size of devices given in the legend. Control signals are not shown as they would fill up the drawing.

6.2.4.2 MCM production technique

The design process of the multi-layer structure is based on an industrially-available production technique for high-density PCBs. The process, called DYCOstrate, is characterized by its use of plasma-drilled, buried vias for interconnection between layers. It also combines low design constraints with low cost.

The multi-layer structure of the DYCOstrate MCM is fabricated starting from a polyimide core layer of 50 µm thickness. The core is surrounded by outer foils (35 µm) with copper tracks (25 µm) on top and underneath. This leads to a four-layer MCM structure of about 220 µm thickness. All layers are laminated at elevated temperature and pressure onto the primary base, which is copper or aluminium (1 mm).

The DYCOstrate design rules are: 100 µm track thickness, 100 µm line-to-line space, and 100 µm diameter for the plasma-drilled, buried vias. The via pad diameter is 300 µm for the core layer and 350 µm for the outer layers.

One of the most challenging tasks in the MCM design is the thermal management. The increased component density and the use of high power dies leads to growing failure rates with temperature. High-power dies on the MCM will need a low thermal resistance from chip to

case. This can be achieved using a cut-out for the G-link transmitter and thermally conducting vias for the FADCs. The total power consumption of the four-channel MCM will be about 5.0 W at full speed.

Figure 6-14 shows a side view of the MCM layer structure. The base is in thermal contact with a heat sink. The pins are arranged as a surface-mounted lead frame around the base. A high-density connector is also feasible to ensure a quick replacement following component failure. Each MCM substrate will be electrically tested (bare-board test) before bonding. All dies will be bonded onto the MCM using standard wire-bond techniques, and a plastic encapsulation (glop-top) will be used to protect the dies against the environment. The total height is about 14 mm including heat sink and lead frame. Lateral dimensions will be approximately 36×25 mm^2. The design will be done in the ASIC laboratory at the University of Heidelberg using Cadence Design Tools such as the Advanced Package Designer (APD). This is a physical layout system for creating and optimizing microelectronic packages such as MCMs. Other tools of this integrated software set will be used to check the layout for EMI rules, signal-to-noise performance and thermal behaviour.

6.2.5 DAQ and test facilities

Two categories of data have to be read out from the Preprocessor — raw calorimeter data after digitization, and the results of BCID. The PPrASICs are the source of these data. To transfer the data to the DAQ system they must be collected and sent to the ATLAS standard readout buffers (ROBs). The strategy is to use an ASIC on each PPM to collect the data from all 64 channels. This ASIC performs data compression and sends data over a custom bus system to a ROD, where data from several modules are merged to a data stream transmitted via S-link to the ROB. It is assumed that eight PPMs (corresponding to one calorimeter sector) are connected to one ROD, but if bandwidth requirements make it necessary one ROD could be used for four PPMs. This would still fit into the envisaged crate structure, but would double the number of RODs and ROBs.

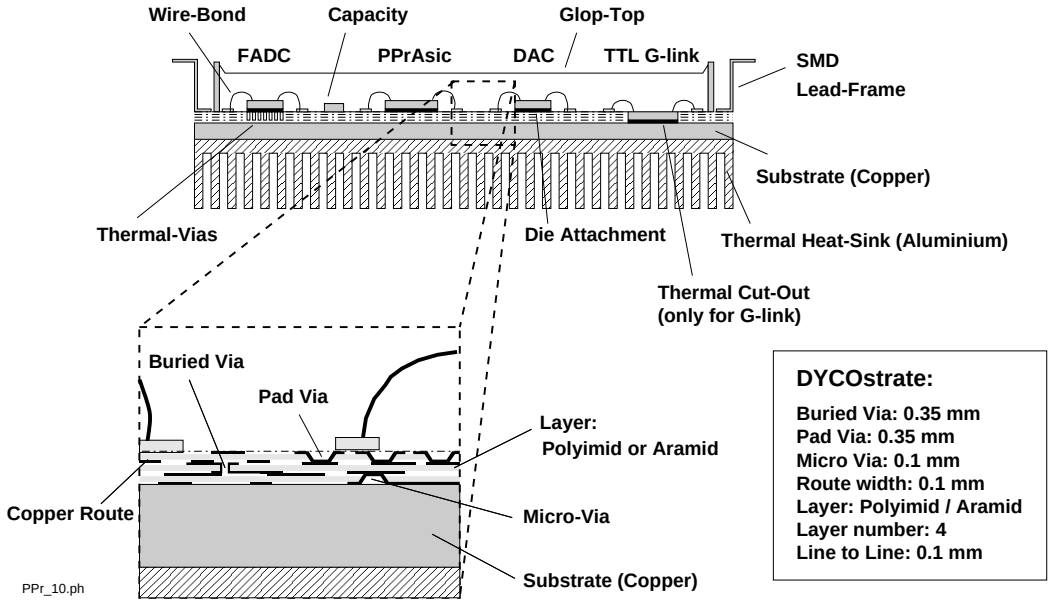

Figure 6-14 Structure of Preprocessor MCM (side view).

The essential parts of the system are shown in Figure 6-15. The Readout Merger ASIC (RemASIC) is used to collect the readout data. It is connected to the PPrASICs via a serial interface and provides direct I/O to the custom bus, the PipelineBus. A prototype of this ASIC has been developed and first functional tests have been performed successfully (see Section 7.6.2). A test system for detailed analysis of the performance of the prototypes is under construction.

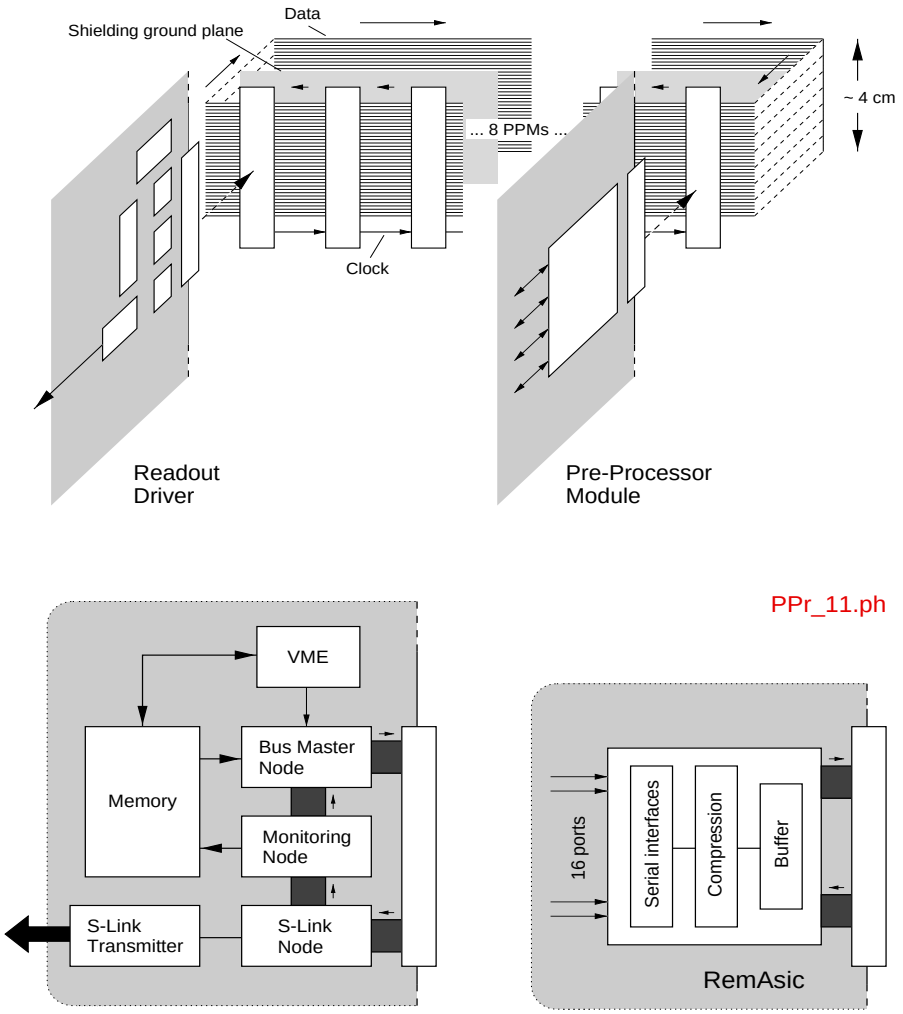

Figure 6-15 Preprocessor data-acquisition scheme.

6.2.5.1 Readout Merger ASIC

The RemASIC, shown in Figure 6-16, is used to concentrate readout data from many sources and to send them as a block over the PipelineBus to the ROD module. One RemASIC is used for each PPM, with 64 channels of data provided by 32 PPrASICs. The final RemASIC will have 16 ports to the serial interfaces of the PPrASIC, each connecting to two daisy-chained PPrASICs. One input and one output signal are needed per port, but the clock and control signals can be common to all ports. The RemASIC controls the PPrASICs.

A level-1 accept signal occurs every 10 µs on average at the maximum level-1 rate of 100 kHz. Fluctuations are smoothed by the derandomizer buffers, located on the PPrASICs, before serial

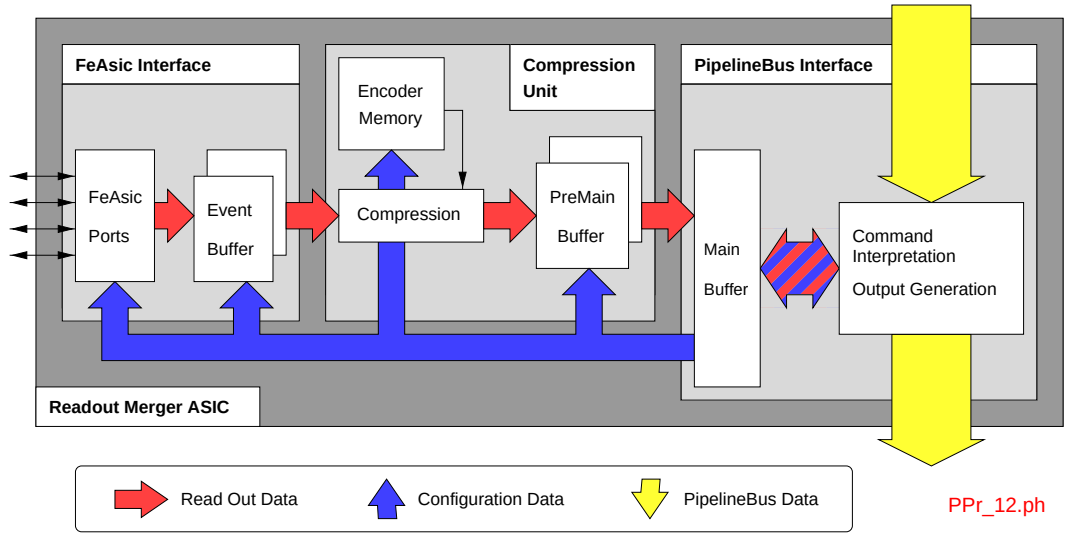

Figure 6-16 Readout Merger ASIC.

transmission. With a serial clock frequency of 40 MHz there are 400 clock cycles available to read out one event from one four-channel daisy-chain. This allows 100 bits per channel (100 clock cycles) to be transmitted. If the maximum amount of data to be read is assumed to be five 10-bit words (FADC output), plus one 8-bit word (LUT output), plus one additional 10-bit word to include part of the bunch-crossing and event numbers, there exists a safety margin of 40 clock cycles which can be used to transmit status and protocol information, and for processing time needed by the interfaces.

While reading out the PPrASICs, an event data block is constructed and stored in a buffer memory. The RemASIC incorporates a compression unit, which performs data compression algorithms on each block of data belonging to a specific event. After compression, data are stored in another memory, where they are made available to the PipelineBus by the integrated interface.

For configuration, the RemASIC receives data from the PipelineBus or the local VME interface and sends it via the serial interfaces to the PPrASICs. The RemASIC itself is configured by data received from the same sources. The RemASIC contains a boundary scan interface, which enables the backplane connections between different RemASICs to be tested.

6.2.5.2 PipelineBus

The PipelineBus connects a set of modules to the ROD, which in turn transmits the data to the next stage of data acquisition, the ATLAS-standard Readout Buffer (ROB). The ROD is also responsible for control of the bus.

The bus itself is made up of several nodes, connected by a 35-bit wide point-to-point connection in a ring-like fashion. All nodes receive a common bus clock signal, which is used to latch input data into the node and make output data available for the next node. Each node has to process the input data and to give out the same, or an appropriate, response to the data at the next clock cycle. In this way, data are shifted through the pipeline formed by the nodes and their connections.

The 35 bits on the bus consist of a 32-bit data word, two control bits identifying the type of user data and one parity bit for the data and control bits. Three types of data words are defined: command words, user data words, and empty slots.

Command words contain three parts: an address identifying the nodes to which the command is directed, a token identifying the command, and an argument space used for parameters or data produced by the node in response to the command. Commands toggle actions of the addressed nodes, send configuration information, read back status information, and are used to indicate error situations.

User data words contain configuration or readout data, sent to or from individual nodes. If no commands or data are written to the bus, the words containing no information are marked as empty slots.

Each PipelineBus ring contains one master node, which injects commands into the pipeline to control the other nodes. It receives the response to the commands and initiates the appropriate actions.

There are two major tasks, which the master has to cope with — configuration and readout. In the case of configuration, the master injects a command into the pipeline, switching the addressed nodes to a mode in which they are able to receive configuration data. It then sends blocks of data directed to individual nodes or groups of nodes. While the configuration process is executed, the master checks for error conditions and discards the configuration from the bus after use.

For a readout process the master injects a command into the pipeline, which switches the addressed node to a state where it is ready to send data into the pipeline. It then checks the nodes for available data. If they all have data available, it initiates a readout process by injecting a command into the pipeline indicating the start of a data block followed by empty slots. Each node appends its data to this command or to the following block by replacing empty slots with its own data. It terminates its data block by a command indicating the end of a data block, including the address of the node to which the data belongs. By this procedure a continuous block of data is constructed.

A node connected to an S-link transmitter reads this continuous block of data and feeds it into the S-link. Since the format of the data on the PipelineBus corresponds well to the format used by the S-link, there are only some header and trailer words to be generated by this node. After transfer to the S-link the readout data is deleted from the bus.

The maximum bandwidth of the bus depends on the frequency of the bus clock. A frequency of 40 MHz results in $40\,\text{MHz} \times 32\,\text{bit} = 152\,\text{Mbyte/s}$. This is the theoretical limit for the bandwidth, including protocol and idle times. By increasing the bus clock frequency, higher bandwidths can be realized.

6.2.5.3 Readout Driver (ROD)

The ROD module is the interface to the DAQ system. It has to transmit the readout data to the ROB by using the standard Readout Link (ROL), which is assumed to be an S-link.

On the other side, the ROD forms the interface to the modules which are the source of readout data. One ROD connects via a PipelineBus ring to eight Preprocessor Modules. It has to contain three PipelineBus nodes, a master controlling the bus activity, an S-link node transferring

readout data to the S-link interface, and a monitoring node used to control readout data independently of the main DAQ system. All nodes will be implemented using FPGA technology, and thus allowing for flexibility and evolution of algorithms and procedures.

In addition to the S-link and PipelineBus interfaces, the ROD contains a VME bus interface, which is used to configure and control the functioning of the ROD.

The master node receives configuration data required by the connected PPMs and sends it via PipelineBus to its destinations. It initiates and controls the readout process, and reacts to error conditions. The S-link node receives readout data and adds header and trailer information to form the standard readout format, and transmits the data via S-link to the ROBs. To double the throughput to the DAQ system a second S-link node could be added. The monitoring node copies PipelineBus data to a VME-accessible memory, which makes it possible to read out data independently from the DAQ system and to control correct bus operation.

6.2.5.4 Volume of readout data

The amount of data which has to be read out depends on the rate of level-1 triggers and the number of words to be read out per event. Table 6-2 shows the volume of data and the resulting data rates under several conditions, assuming a configuration of 64 channels per PPM and 8 PPMs per ROD.

Table 6-2 Volume of data to be read out and the resulting data rates.

	1 BC BCID result	**3 BCs raw input data**	**5 BCs raw input data**
1 channel	8 bits	3×10 bits = 30 bits	5×10 bits = 50 bits
PPM (64 channels)	512 bits	1920 bits	3200 bits
PipelineBus ring (8 PPMs)	4096 bits	15360 bits	25600 bits
Data rate at 75 kHz	36.6 Mbyte / s	137.3 Mbyte / s	228.9 Mbyte / s
Data rate at 100 kHz	48.8 Mbyte / s	183.1 Mbyte / s	305.2 Mbyte / s
Data rate after compression	19.5 Mbyte / s	73.2 Mbyte / s	122.1 Mbyte / s
Bandwidth of PipelineBus	40 MHz $\times$ 32 bits = 152.6 Mbyte / s		
Bandwidth of fibre-channel S-link	103 Mbyte / s		

These rates have to be compared with the bandwidth of the PipelineBus and the S-link. The theoretical maximum bandwidth of the PipelineBus at a clock frequency of 40 MHz is 152.6 Mbyte / s (protocol overheads reduce the effective bandwidth). By increasing the clock frequency to 60 MHz a bandwidth of 228.9 Mbyte / s results. For each PPM in the ring, one 32-bit word is added as a separator and one 32-bit word per event is added as a trailer. This results in a data rate of 3.4 Mbyte / s at a trigger rate of 100 kHz and a bus clock of 40 MHz, or 2.3% of the raw bandwidth.

The theoretical maximum bandwidth of the S-link is the same as for the PipelineBus at 40 MHz, i.e. 152.6 Mbyte / s, but available implementations do not yet reach this limit. By using two S-links per PipelineBus ring the S-link bandwidth could be doubled at the cost of twice as many links and ROBs.

The comparison of the data rates shows that the system with the described configuration is able to read out the results of BCID for all events under all conditions, as shown in Table 6-2. For reading out more data, such as five bunch-crossings per event, some form of data reduction is required. Simulations show that loss-free data compression by a factor of 2.5 is possible [6-2]. With the hardware compression unit integrated in the RemASIC it should be possible to read out all data at the maximum trigger rate.

In the proposed running scheme, the 8-bit BCID result will be recorded for all triggered events, and up to 5 BC of raw input data recorded only for a small, prescaled proportion of triggers. In conjunction with the results sent to the CTP, this allows the reconstruction offline of the mode and location of all triggers, which is useful both for physics analysis and for checking operation of the calorimeter trigger logic.

6.2.5.5 Physical implementation of the PipelineBus

The PipelineBus will be physically implemented as a custom backplane with three main layers. Point-to-point links between adjacent modules are made on one layer, and the connection over the whole width of the PipelineBus closing the ring is made on another layer. A ground plane between these two layers minimizes crosstalk.

The ROD drives the long connection over the full width of the bus. Then data propagate from one slot to the next and back to the ROD. The bus clock is distributed in the reverse direction to ensure correct timing.

The 64-bit expansion connector of CompactPCI can be used as a module backplane connector. It provides 110 signals arranged in five rows with a spacing of 2 mm, which is well matched to the 35 input and 35 output signals of the PipelineBus. The height of the connector is only 4 cm.

6.2.5.6 Timing, Trigger and Control system

The ROD also contains a TTC receiver chip, which is used to extract the LHC clock, the level-1 accept signal, and the bunch-crossing and event-counter numbers (or the corresponding reset signals). The ROD expands these numbers to their full length, checks for misalignment, and appends the numbers to the data block transmitted to the DAQ system.

On the PPM, the PPrASIC includes the low-order bits of the bunch-crossing and event numbers with the readout data, which the RemASIC uses to check synchronization between channels.

6.3 Cluster Processor

6.3.1 Organization of input signals

In this section we describe in more detail the organization of the Cluster Processor (CP) input signals. The processing of the CP trigger space — 50×64 trigger windows — is distributed over a number of Cluster Processor Modules (CPMs) hosted in four crates, each processing a quadrant in ϕ of the CP trigger space. In order to minimize the transmission bandwidth from the Preprocessor to the CP, each trigger tower is sent only once to any crate. Trigger towers shared by CPMs in the same crate are fanned out across a fast backplane, and trigger towers shared by CPMs in different crates are duplicated at the Preprocessor, i.e. the serial link output is put onto two cables going to corresponding CPMs in crates which process neighbouring quadrants in ϕ. This eliminates the need for data transmission between CP crates and minimizes the latency for distribution of digitized trigger towers. The ϕ-quadrant architecture creates four inter-crate boundaries of 50 towers, totalling $4 \times (50 \times 3 \times 2) = 1200$ duplicated towers, or a duplication overhead of 18.75% (of the 6400 towers in the $50 \times 64 \times 2$ CP trigger-tower space).

It is desirable that the dimension of a CPM in each direction be close to a factor of the trigger space dimension, whilst being efficiently divisible into smaller regions for processing at the ASIC level. Constraints of board space, power density and backplane connectivity set the basic CPM modularity at ~64 trigger windows per module. In the ϕ-quadrant architecture each CPM processes 4×16 trigger windows and each ϕ-quadrant crate has 13 such CPMs to cover the whole quadrant in η, as shown above in Figure 6-3. Note that the central CPM in each crate spans the $\eta = 0$ boundary and that the end modules in each crate are only supplied with inputs for 3/4 of the 4×4 windows they are capable of processing, the end rows of 1×16 windows being unused. If the $\eta = 0$ boundary were aligned between CPMs there would be 14 CPMs per crate instead of 13, which increases the component cost of the system by 7.6% and the amount of unused processing by a factor of three.

Each CPM needs an environment of $7 \times 19 \times 2$ trigger towers (e.m. and hadronic) to complete its 4×16 trigger windows; as explained in Section 6.1.3, any $m \times n$ group of 4×4 windows requires $(m + 3) \times (n + 3)$ towers to perform the trigger algorithms described in Chapter 4. The bunch-crossing multiplexing of pairs of towers in ϕ limits both serial communication from the Preprocessor and transmission across the CP backplane to an even number of towers in ϕ, and hence the basic environment is effectively extended to 20 towers in ϕ. The 20th row of towers received either via serial link from the Preprocessor, or fanned in from the backplane, is simply ignored after the demultiplexing of tower pairs.

The duplication in the Preprocessor of shared towers at inter-crate boundaries results in the mapping between PPMs and CPMs/JEMs being effectively one-to-three/three-to-one. A PPM covering a given quadrant sends all of its towers to a CPM covering the same quadrant (the ‘primary’ CPM), but also duplicates the two rows of towers at each edge in ϕ to send them to CPMs covering the same η position in adjacent quadrants (the ‘secondary’ CPMs). Conversely, each CPM receives the majority of its serial links from a primary PPM, with the additional links completing the environment coming from secondary PPMs in adjacent quadrants. The allocation of towers to serial links and the mapping of PPMs to CPMs (and also JEMs) is shown in Figure 6-17.

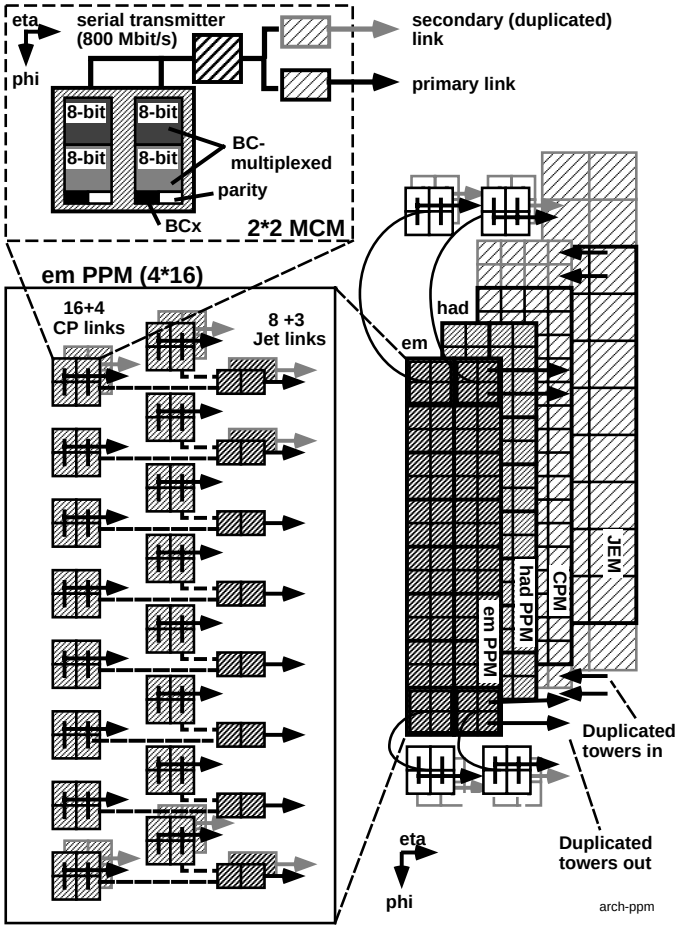

Figure 6-17 Mapping of CPMs (and JEMs) to PPMs.

Each CPM receives $4 \times 20 \times 2$ towers of its environment on 40 serial links (20 electromagnetic and 20 hadronic links), covering 4/7 of the full environment. Additional columns of towers needed to complete the environment in η must be fanned in from the backplane. These additional towers are supplied by neighbouring CPMs in the same quadrant, which receive them from the Preprocessor. Thus the directly received towers[1] of one CPM are fanned out via the backplane to become the fanned-in towers of the adjacent CPMs, and vice versa.

Within a CP crate, the fan-out of towers shared by separate CPMs processing neighbouring regions is accomplished by semi-serial distribution of towers on single-ended point-to-point links forming a high-density multi-layer backplane. In the ϕ-quadrant architecture, each section of the backplane behind a given CPM is occupied only by data lines carrying towers transmitted or received by that CPM; no other signals need pass through its section of the backplane. The backplane may in effect be subdivided into groups of short module-to-module links, which offers good signal grounding and high noise immunity.

The shape of the proposed CPMs is narrow in η, resulting in greater inter-CPM fan-in/fan-out than for squarer geometries; only $1 \times 20 \times 2$ towers on each CPM do not need to be shared with neighbouring CPMs. To complete the CPM environment, $2 \times 20 \times 2$ towers are received from the

1. Strictly speaking, only 3/4 of the directly received towers need to be fanned out, one column to one side in η and two columns to the other side.

CPM on one side and $1 \times 20 \times 2$ towers from the CPM on the other side; each CPM must also fan out the same volume of data to its neighbours. Thus for the $7 \times 20 \times 2 = 280$ towers of the CPM environment, $4 \times 20 \times 2 = 160$ are received directly from the Preprocessor and $3 \times 20 \times 2 = 120$ are received from neighbouring CPMs via the backplane. The source of towers within the environment of a CPM is shown in Figure 6-18.

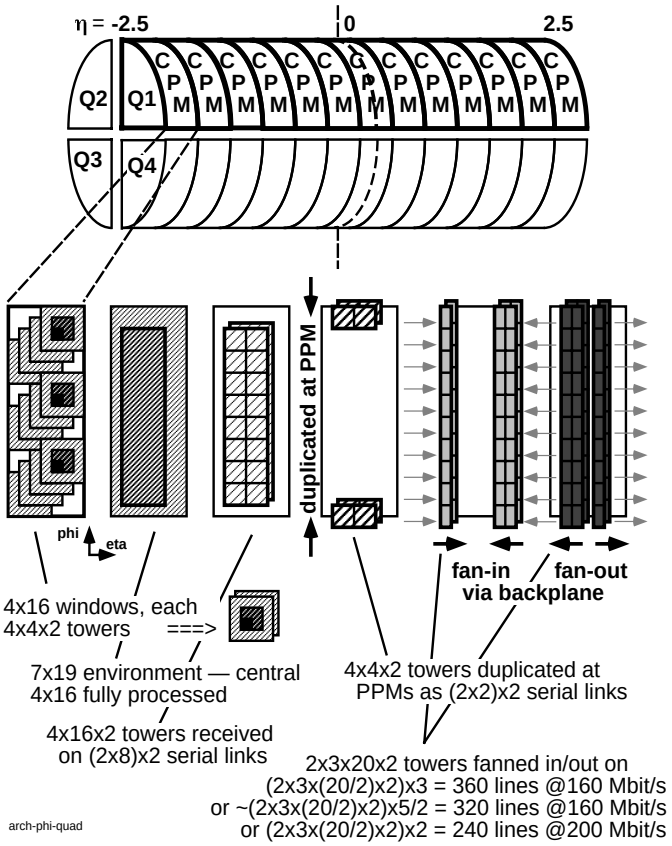

Figure 6-18 CPM layout, showing the ϕ-quadrant architecture.

One can compute the connectivity required for the CP backplane under a variety of assumptions about the speed of the backplane, and the way in which towers might be placed onto the backplane lines. A number of alternatives are shown in Table 6-3. Each row describes a different backplane data format, and the four columns show ways of grouping towers onto combinations of backplane lines, in order to enable bunch-crossing multiplexing plus parity (first column), multiplexing plus shared error-detection (second column), regrouping across layers with shared error-detection (third column) or transmission of an 8-bit value only (fourth column). The first version is the currently preferred option, since it economizes on bandwidth, whilst providing some simple error detection. The benefits of enhanced error-detection through jointly error-encoding both 8 + 1-bit words with two error bits are being evaluated against the cost of increased backplane fan-in / fan-out due to pairing towers in η as well as in ϕ. The final 'naive' scenario without bunch-crossing multiplexing or error detection is shown for comparison — such high connectivity would be difficult to realize in a reliable system.

The BC-multiplex and error-detection schemes require two extra bits in addition to the eight data bits per trigger tower. These could be accommodated on the backplane without increasing its speed by dedicating a fifth 160 Mbit / s line for each 2×2 group of trigger towers, carrying

Table 6-3 Numbers of CPM backplane connections for trigger tower fan-out/fan-in under a variety of backplane speed and grouping options — calculated as *bidirectional* × (η × φ × *layers*) × *lines-per-word*. The box indicates the currently-preferred option.

Backplane data format	1x2x1 tower pairs BC-MUX + parity error detection: 8+1+1 = 10 bits)	2x2x1 tower quads BC-MUX + shared error detection: 2x(8+1)+2 = 20 bits	1x2x2 tower stack BC-MUX + shared error detection: 2x(8+1)+2 = 20 bits	Single tower No BC-MUX or error detection: 8 bits
160 Mbit / s 2 lines per 8-bit word	n / a	n / a	n / a	2×(3×19×2)×2 = **456**
160 Mbit / s 3 lines per 10-bit word	2×(3×20 / 2×2)×3 = **360**	2×(4×20 / 2×2)×3 = **480**	2×(3×20 / 2×2)×3 = **360**	2×(3×19×2)×3 = **684**
160 Mbit / s 5 lines per 2×10-bit words	2×(2×20 / 2×2)×5 / 2 +2×(1×20 / 2×2)×3 = **320**	2×(4×20 / 2×2)×5 / 2 = **400**	2×(3×20 / 2×2)×5 / 2 = **300**	2×(3×19×2)×5 / 2 = **570**
200 Mbit / s 2 lines per 10-bit word	2×(3×20 / 2×2)×2 = **240**	2×(4×20 / 2×2)×2 = **320**	2×(3×20 / 2×2)×2 = **240**	2×(3×19×2)×2 = **456**

the extra bits for both BC-multiplexed tower pairs. The resulting backplane connectivity is 320 trigger-tower I / O lines per CPM.

Other backplane connections are needed for set-up and control (via a VME interface), power and grounding. Around 400 backplane connections is considered an acceptable figure, and so those options which require ~300 connections for trigger-tower traffic are preferable from an engineering standpoint.

6.3.2 Results for the CTP and level-2 trigger

Consolidation and compression of the CTP and RoI outputs are performed in 'merger' stages before onward transmission to the CTP and level-2. In this context, 'merging' means counting of trigger clusters. The CTP needs only the multiplicity of objects passing each set of trigger conditions; no information on the coordinates of those objects is required for the level-1 trigger decision. The local-maximum condition of the declustering algorithm (Section 4.2.2) allows only one trigger object of each class within any 2×2 set of windows. Each CPM may therefore find a maximum of 16 objects at any given threshold, although the natural sparsity of isolated clusters will mean that even in triggered events the majority of potential objects will not pass any thresholds. Furthermore, the CTP imposes a maximum count of seven objects for each threshold combination so that the multiplicity can be described with only three bits; any count exceeding the maximum will be represented as the maximum value. The Cluster Processor must therefore merge the 16 potential objects on each of 52 CPMs to a single global three-bit count, doing this in parallel for each of eight electron / photon threshold combinations and eight hadron / tau threshold combinations.

The counting of trigger objects proceeds in three stages:

1. Each Cluster Processor ASIC locates any electron / photon or hadron / tau objects within each 2×2 region and determines which thresholds were passed. The results from each ASIC are delivered in parallel as two 16-bit words (8 bits for electron / photon thresholds

and 8 bits for hadron / tau thresholds). If no object is found, all threshold bits are set to zero.

2. Each CPM must count all the possible objects independently for each threshold, truncating each count at a maximum of seven objects, and exporting its results as 16×3-bit multiplicities.
3. The Cluster Processor must globally count the 52 CPM counts to produce one 3-bit multiplicity per threshold — a task performed by the Cluster Merger Modules (CMMs).

Step 2 is relatively straightforward, and can be performed with a 16-bit sum on the CPM for each of the 16 thresholds, forcing each result to three bits. One method is to use FPGAs. Step 3 is more complex, since for each threshold the 3-bit sub-sums from 52 CPMs in four crates must be merged, which requires a very high degree of fan-in. The 16 thresholds will be handled independently, each on a separate CMM, with the sum implemented using LUTs. The summing of such a wide input must be done in a number of stages, increasing the latency of the addition — it is likely that two or three bunch-crossings will be required to produce the global count at each threshold. The 16 CMMs will be located in a separate crate close to the four CPM crates and the CTP, in order to minimize propagation delays. The hardware implementation of the merging of trigger sums for the CTP is described in more detail in Section 6.3.8.

Each trigger object found in an event accepted by the CTP must be described to the level-2 trigger by information on which threshold combinations it passed, as well as a global coordinate in the CP trigger space specifying its location with a resolution equal to the granularity of the trigger space, 0.1×0.1 in η and ϕ. A 16-bit word describing which electron / photon and hadron / tau thresholds were passed and a 14-bit co-ordinate[1] describing the location of the object together form a region-of-interest (RoI) transmitted to level-2. The RoI must also be tagged with the number of the bunch-crossing within which it occurred, in order that it can be correctly associated in time with the full-granularity information from the calorimeters.

Data on two potential RoI objects are produced in each CPASIC (see Section 6.3.7 below) for every 25 ns cycle. This must be held in a pipeline until the level-1 trigger decision is returned, initiating transfer to the level-2 trigger. A latency of at least 100 μs is acceptable before the RoI data reaches the level-2 trigger, but an average throughput of one event every 10 μs must be maintained. Simulation indicates that more than 99% of RoI data has no threshold bits set, and these data are eliminated before the remainder are merged into a single stream for delivery to level-2. This is a complex data reduction, but fortunately the low average rate of accepted events makes it possible to export all RoI objects from each CPM, including those with zero thresholds passed, shifting the compression task to a downstream RoI merger module. A single serial link provides sufficient bandwidth to export all 16 possible RoI objects from each CPM for every accepted event. There will be one RoI merger module for each CP crate, handling 13 serial links each. The four sets of retained RoIs must probably be further combined into a single frame per accepted event; this is to be agreed with the level-2 group in discussions over the detailed format of the level-1 / level-2 interface. The hardware implementation of this scheme is described in more detail in Section 6.3.10 below.

1. 12 bits can specify one location within the 50×64 CP space, but two additional bits are needed to allow the electron / photon and hadron / tau algorithms to decluster to independent locations within each 2×2 RoI region.

Additional readout must be provided for the full set of inputs and outputs for selected events, in order to test the transmission between the Preprocessor and CP, and the correct functioning of the CP itself. This will be similar to the above model for the RoIs, with serial transmission of multiple towers in parallel to a downstream ROD, which consolidates and compresses the data before communicating it to an ROB as an interface to the full ATLAS DAQ. This is described in more detail in Section 6.3.10 below.

6.3.3 Cluster Processor hardware overview

The Cluster Processor (CP) will process 6400 electromagnetic and hadronic trigger towers, each 0.1×0.1 in η–ϕ space, covering a total pseudorapidity region of ± 2.5. It will receive 8-bit trigger-tower signals from the Preprocessor, serialized and BC-multiplexed (see Section 6.2.1.5) at 960 MBd on ~2000 coaxial cables. The CP will provide electron / photon and hadron / tau trigger multiplicity information to the CTP, and RoI information to the level-2 trigger. The system will also provide intermediate and final results to the DAQ system for monitoring and diagnostic purposes.

The CP consists of four crates of Cluster Processing Modules (CPMs), each processing a quadrant of the calorimeter in ϕ over the full η space. This is performed by 13 CPMs per crate, each processing a 4×16 area of the calorimeter. The partial results from all the CPMs will be merged in a separate crate using Cluster Merger Modules (CMMs), and the results transferred to the CTP. RoI and DAQ data from all CPMs will be transferred to RODs in a separate crate. Figure 6-19 shows a block diagram of the system. Each of the six crates in the CP will include a module which will interface to the LHC timing distribution (TTC) system and the Detector Control System (DCS).

The real-time data flow through the system proceeds in several stages. Each 960 MBd data stream (800 Mbit / s of user data plus link protocol bits) received on a CPM is converted back to parallel data and then serialized into 160 Mbit / s data streams (two per 8-bit trigger tower) for input to the Cluster Processor ASICs and fan-out on the backplane for transmission to neighbouring CPMs. The cluster-processing algorithm and the transfer of results to the CTP will operate at 40 MHz. RoI and DAQ data are not used in real time, but are stored in dual-port memories and transferred out only on receipt of a level-1 accept signal. The backplane is discussed in more detail in Section 8.4.5.

The CP will include four module designs, two ASIC designs, and one MCM design.

6.3.4 Cluster Processor Module (CPM)

The CPM (Figure 6-20) will process a 0.4×1.6 η–ϕ region of the calorimeter, requiring signals from 280 electromagnetic and hadronic trigger towers. Of these, 160 will be transferred directly from the Preprocessor using HP G-link chip-sets (or similar devices) and 40 coaxial or twin-ax cables with compact connectors. The remaining 120 trigger towers will be imported to the module from the two nearest-neighbour CPMs via the backplane at 160 Mbit / s. The same number of trigger towers must be exported to these two neighbours, thus requiring a connector with 320 signal pins. Data will be transported at 160 Mbit / s, by using low-voltage CMOS signalling within the module and by using single-ended ECL signalling over the backplane.

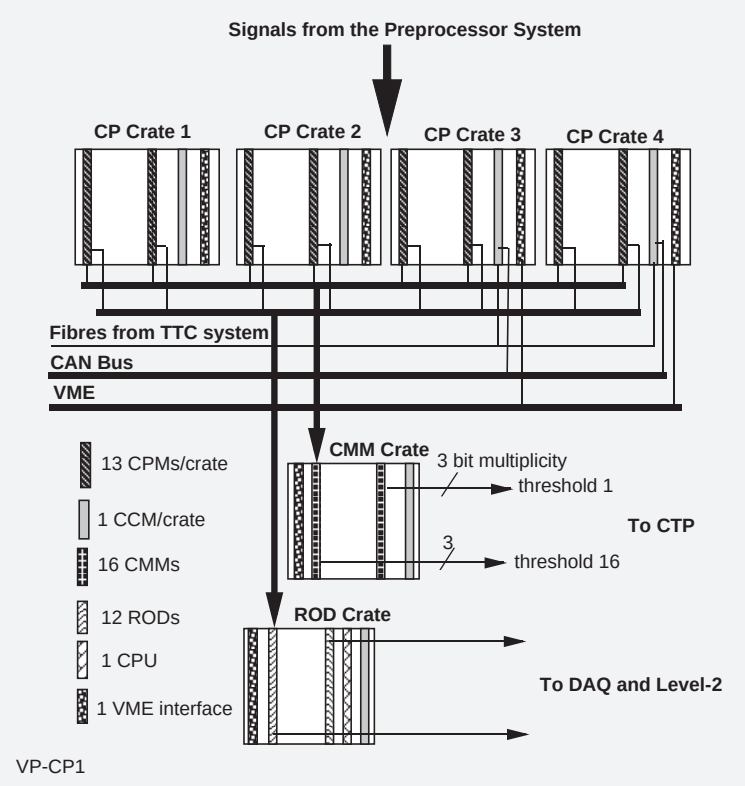

Figure 6-19 Block diagram of the Cluster Processor.

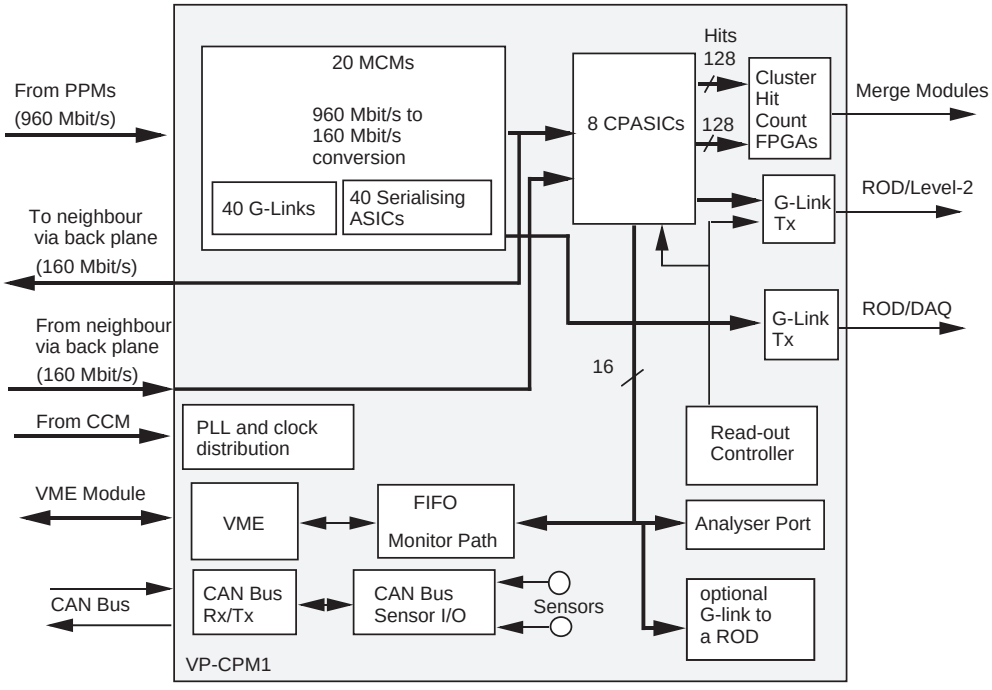

Figure 6-20 Block diagram of the Cluster Processor Module.

Processing the $4 \times 16 \times 2$ trigger-tower region will be performed by eight Cluster Processor ASICs (CPASICs) per CPM, each handling a $4 \times 2 \times 2$ region, to calculate RoIs and trigger cluster multiplicities for each electron / photon and hadron / tau trigger threshold combination.

Cluster counting will be performed by eight FPGAs on each CPM, four for each of the electron / photon and hadron / tau clusters. The partial counts from each CPM will then be transferred to the CMMs using low-voltage differential signalling (LVDS). See Section 6.3.8 for further details.

The readout control (ROC) function, implemented on an FPGA, will initiate and control the transfer of all ASIC data from each CPM to the ROD on receipt of a level-1 accept signal. These data have the crate ID, module ID, BCID number, and event number appended before transfer to the ROD. In addition, intermediate results from the CPASICs will be stored in a FIFO for diagnostic purposes. See Section 6.3.7 for further details.

The clock and control signals to drive all the synchronous processing within the CPMs will be received from the Clock and Control Module (CCM) via the backplane, and will be phase-locked, buffered and distributed on the CPM with minimum skew. Differential ECL will be used to distribute the CCM signals on the backplane to ensure stability of the 160 MHz CPM clocks.

A slow-control interface, such as VME, will provide access to the programmable registers on the module and on the ASICs. Module temperature will be monitored via the Control Area Network (CAN) bus. Live insertion of modules will be supported by multi-level connector pins. Modules will remain inactive until configured by software. Figure 6-21 shows the components on the CPM — this is not a layout drawing. Thermal modelling of the CPM using the appropriate environmental conditions (internal crate air flows, etc.) will be undertaken using Cadence ‘Thermax’ software to inform and guide board layout.

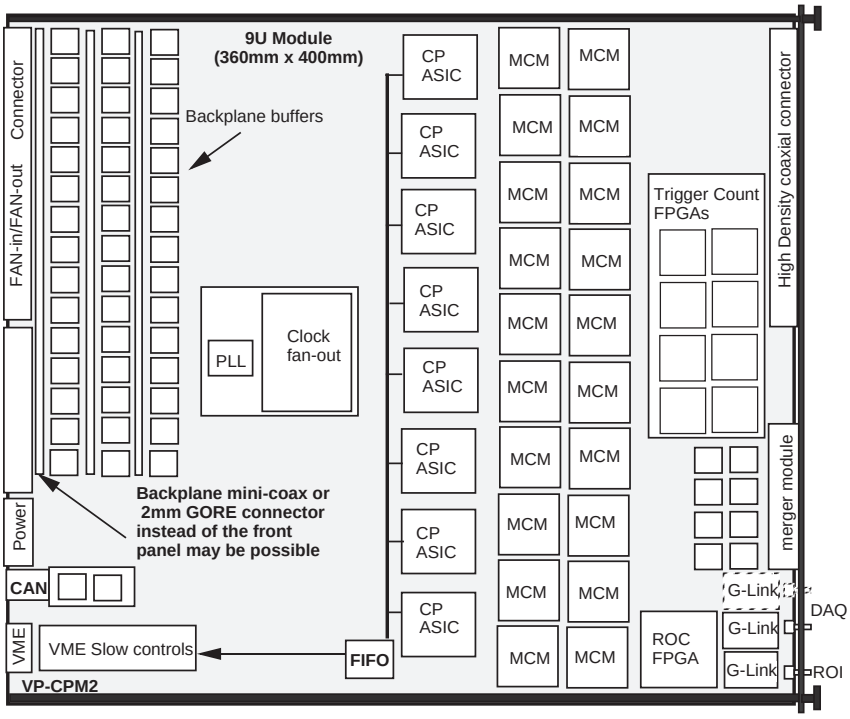

Figure 6-21 Components on the CPM (not a layout).

6.3.5 Cluster Processor Multi-Chip Module

6.3.5.1 Requirements

Multi-Chip Modules will be essential for the CP to achieve the desired channel processing density. Each Cluster-Processor Multi-Chip Module (CPMCM; Figure 6-22) on the CPMs will receive two serial 960 MBd data streams (20 data bits + 4 control bits every 25 ns) from the Preprocessor and will output four sets of five bit streams at 160 Mbit/s. Two types of die will be bonded on to the CPMCM substrate using standard wire-bond techniques:

- the high speed serial-to-parallel converter (HP G-link) to convert the 960 MBd serial data stream to a 20-bit word every 25 ns;
- the 160 Mbit/s Serializing ASIC to multiplex the 20-bit word on to five serial links operating at 160 Mbit/s.

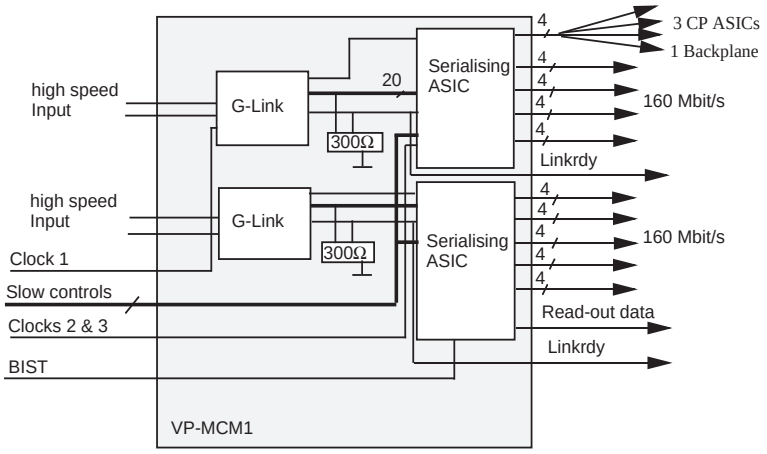

Figure 6-22 Block diagram of the CPMCM.

6.3.5.2 Specifications

Performance

- The CPMCM will receive two channels of 960 MBd NRZ (non-return to zero) format data via the package pins, and deliver 20 output signals (five bits with a four-way fan-out) at 160 Mbit/s per channel.
- Characteristic impedance control using micro-stripline design techniques will be applied to all fast (<200 ps edge speed) signal tracks.
- Crosstalk between the channels will be kept below –20 dB to achieve the bit-error rate (BER) requirement of better than 10^{-12}. Ground guard-traces will be used where necessary.
- The two channels will be fully independent, and require no synchronization on the CPMCM.
- All high-speed signals will be a.c.-coupled and differential, and all other I/O signals will be compatible with 3.3 V TTL.

Technology
The CPMCM technology will be multi-layer MCM-C (alumina 96%) with a lead frame attached to the substrate, which forms the package, and an encapsulating ceramic cover. Some of the important specifications are as follows:

- Substrate
    - four-layer construction (signal / power);
    - size 32 mm × 32 mm;
    - thickness 635 μm;
    - characteristic impedance 50 Ω for 800 Mbit / s data input.
- Package
    - the packaged dimensions will not exceed 32 mm × 32 mm, and the overall height (excluding heat sink) will be less than 5 mm;
    - the number of I / O and power connections will be ≤ 184;
    - the package will be a lead frame with a pitch of 650 μm which will be part of the MCM-C substrate, encapsulated using a ceramic lid and hermetically sealed at up to 55 °C and 95% relative humidity;
    - the thermal resistance of the dies to the package will not exceed 12 °C / W;
    - the package will be capable of dissipating at least 6 W;
    - the package pins will be capable of handling 960 MBd electrical signals.
- Bonding
    - the dies will be attached to the substrate using STAYSTIK 181 thermoplastic adhesive for reworkability, and wire-bonded to the substrate;
    - the passive components (50 Ω and 300 Ω resistors) will be printed on the substrate;
    - decoupling capacitors (100 nF) and decoupling networks (100 nF and 10 Ω) will be bonded close to the power pads of the chips.

Thermal management
Thermal management assumes the use of the current HP G-link receivers (HDMP-D024). Since this die dissipates approximately 2.5 W, thermal management of the CPMCM is crucial for successful operation. The technique will be to mount the dies directly on to the alumina substrate and avoid any layers directly underneath the substrate, as shown in Figure 6-23 (cross-section); a plan view is shown in Figure 6-24.

Testing, rework and ‘known-good’ dies
The HP G-link dies will be tested by the manufacturer for functionality, but not at full speed. However, dies that pass the low frequency tests and the process parameter tests should be ‘known-good’ dies to a high level of certainty.

The use of thermoplastic epoxy for die attachment will enable removal if necessary, by reheating the substrate to the appropriate temperature. The package lid will not be attached with a permanent hermetic seal until the acceptance tests are completed.

By designing suitable test facilities into the Serializing ASICs, possibly including boundary scan, their connectivity and functionality can be rapidly checked on the MCM.

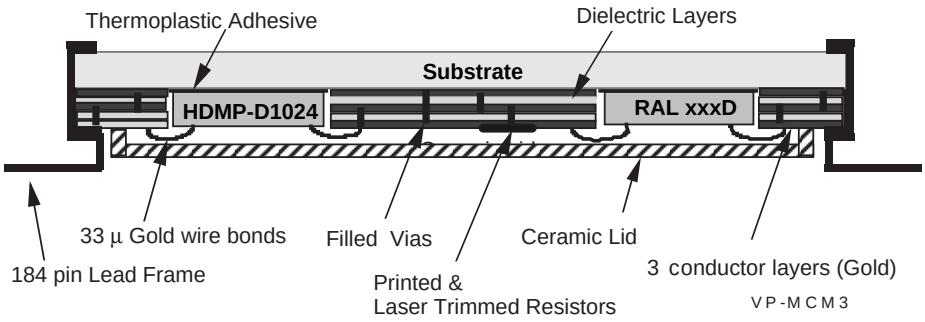

Figure 6-23 Cross-section of the CPMCM.

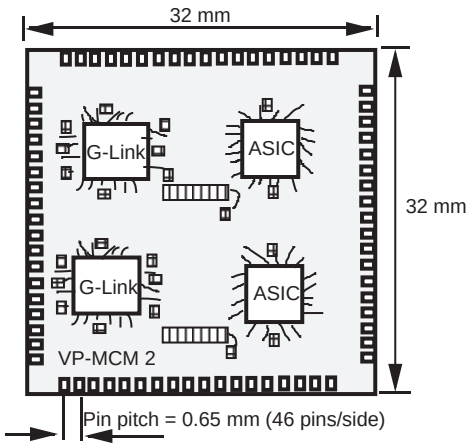

Figure 6-24 Plan view of the CPMCM.

6.3.6 160 Mbit/s Serializing ASIC

6.3.6.1 Requirements

Serializing ASICs are required in the CP to transform the data emerging from the HP G-link receivers (20-bit parallel data every 25 ns) into 160 Mbit/s data streams for input to the CPASICs and for fan-out on the backplane. The incoming data will also be captured in the ASIC memory and transferred to the DAQ (the readout scheme is outlined in Section 6.3.10). This ASIC will also provide test signals for the clock calibration logic required to set up the data and clocks on the CPASICs.

6.3.6.2 Specification

The Serializing ASIC will receive a 20-bit field carrying BC-multiplexed data from four trigger towers —16 data bits with two multiplexing flags and two error check bits — every 25 ns. It will transform the data bits into four 160 Mbit/s bit-streams and the remaining four bits into a fifth 160 Mbit/s bit-stream. (It would be very attractive to operate the ASICs at 200 Mbit/s, thereby requiring only four serial lines, but due to low-volume production not all commercial processes are available for this design at present.) A brief summary of the functionality of the Serializing ASIC is as follows (see Figure 6-25):

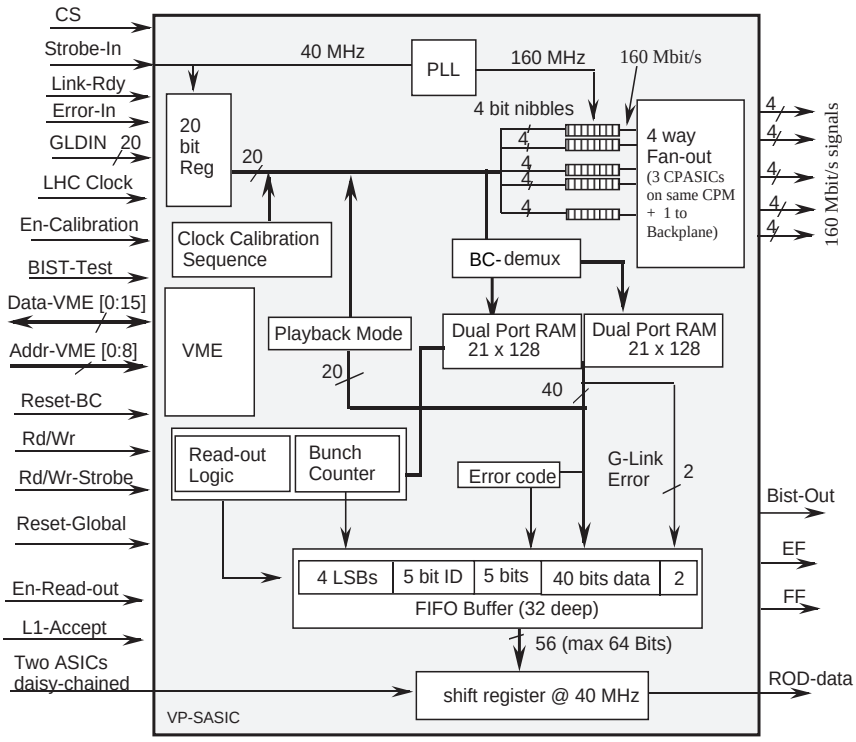

Figure 6-25 Block diagram of the Serializing ASIC.

- receive 20-bit parallel data every 25 ns (40 MHz);
- serialize a 4-bit nibble at 160 Mbit / s;
- provide four-way fan-out to feed up to four CPASICs;
- provide 128-deep dual-port memory for input data capture (and also for use as a test and diagnostic playback memory);
- provide 32-deep readout FIFO;
- transfer data from the dual-port memory to the FIFO on receipt of a level-1 accept signal;
- transfer data from the FIFO to the ROD on receipt of an Enable-Readout signal from the CPM;
- merge data with an ASIC ID and an error code (parity);
- shift the 40-bit readout data out of the ASIC at 40 MHz;
- capability of daisy-chaining two devices for data transfer;
- perform the calibration sequence for calibrating the clock and data on the CPASICs;
- provide a built-in self-test (BIST) facility;
- provide slow control access to all registers and memories.

6.3.6.3 Technology

The design will be implemented on a 0.6 µm (or smaller) geometry in CMOS technology. The estimated die size is 20 mm^2. A maximum of 120 pads is required for I/O plus power and

grounds. As the ASICs will be mounted on an MCM, they will be supplied and tested in bare die form. All inputs and outputs will be CMOS-compatible and capable of interfacing to 3.3 V TTL.

6.3.7 Cluster Processor ASIC

6.3.7.1 Requirements

The most complex part of the logic required for the electron / photon and hadron / tau trigger processor will be implemented on the CPASIC. It will fully process a $4 \times 2 \times 2$ trigger tower region and will provide cluster hit information and RoI information.

To fully process a region of this size, the CPASIC requires information from a total of $[(4 + 3) \times (2 + 3)] \times 2 = 70$ electromagnetic and hadronic trigger towers. With 8-bit trigger tower data in parallel, the CPASIC would require 560 input pins. By serializing (time-multiplexing) the data to 160 Mbit / s data streams this can be reduced to 140 pins. Using the BC-multiplexing scheme (Section 6.2.1.5) this number can be further reduced to 120 pins. A block diagram is shown in Figure 6-26.

6.3.7.2 Specifications

The functional specification of the CPASIC is as follows:

- receive digitized data for the electromagnetic and hadronic trigger towers from 120 bit-stream links operating at 160 Mbit / s;
- align the serial data to the 160 MHz clock, provide serial-to-parallel conversion and synchronize the parallel data with the 40 MHz system clock;
- provide BC-demultiplexing and error detection, and control the rate of the trigger output if errors are detected;
- perform the electron / photon and hadron / tau trigger algorithms on $4 \times 2 \times 2$ trigger towers at 40 MHz, using pipeline adders and comparator processing elements;
- provide eight sets of programmable threshold values for the isolated electron / photon cluster trigger (each set consists of an independently programmable cluster threshold, e.m. isolation threshold, and hadronic isolation threshold);
- provide eight sets of programmable threshold values for the hadron / tau trigger (as above);
- results and outputs:
    a. 16-bit hit result: two bits per threshold for the electron / photon trigger,
    b. 16-bit hit result: two bits per threshold for the hadron / tau trigger,
    c. 20-bit RoI for the electron / photon trigger,
    d. 20-bit RoI for the hadron / tau trigger;

  N.B. a and b are available on pins in real time, while c) and d) will only be available via the readout scheme;

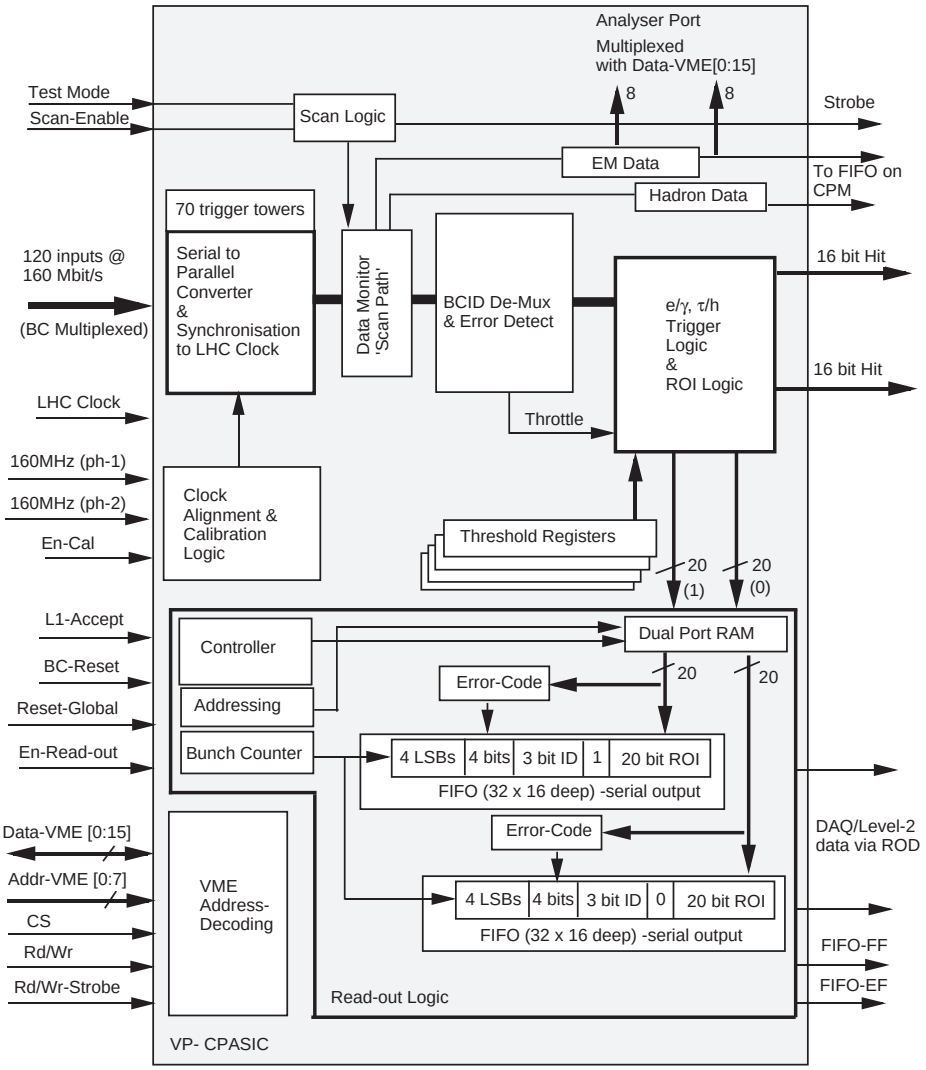

Figure 6-26 Block diagram of the CPASIC.

- provide readout logic to DAQ;
- include built-in self-test (BIST) logic;
- use a 160 MHz clock for the 160 Mbit/s links, and a 40 MHz clock for the core logic. The 160 MHz clock may be generated internally on the ASIC;
- provide a monitoring stage after the serial-to-parallel conversion.

6.3.7.3 Logic blocks

A brief description of the major blocks within the CPASIC is as follows:

Serial-to-parallel logic and clock-alignment logic
The 160 Mbit/s data streams do not have a clock recovery scheme encoded with the data. The function of the clock alignment logic is, therefore, firstly to select the appropriate 160 MHz clock phase to correctly capture the incoming serial data at 160 Mbit/s, and then to synchronize to the 25 ns clock period using 6.25 ns delay elements. This entire procedure takes approximately

1.5 μs. A test ASIC (RAL215) to evaluate this calibration procedure has been designed and tested successfully. Once the clock and data are aligned then each four bits within the 25 ns period will be converted to a parallel nibble for the next stage of processing.

BC-demultiplexing and error detection
The BC-demultiplexing logic will first look at the non-zero 20-bit parallel data (two eight-bit trigger towers with their associated flag bits and error codes) derived from the serial-to-parallel converter block. The BC-demultiplexing logic will use the flag bits to assign the data to the appropriate trigger tower and time slice. The data will then be latched and delayed until the next 20-bit word is decoded in a similar fashion. This process enables four trigger towers to be handled by a single 20-bit field. To avoid additional latency, error checking can be performed in parallel to the algorithm processing, but the procedure must be complete before the trigger logic can finish its task. If the error checking logic finds data from any of the trigger towers in error, then the corresponding part of the algorithm logic will be disabled to prevent the suspect trigger tower contributing to the triggering.

Trigger algorithms
Using 4×4 sliding windows within the calorimeter, the algorithm will search for isolated electron / photon and hadron / tau clusters and provide trigger objects and RoIs. The algorithms will be implemented using pipeline adders and comparators operating with 25 ns pipeline steps.

Regions-of-Interest
The RoI logic will select the 2×2 windows with the largest E_T from the fully processed $4 \times 2 \times 2$ trigger towers, and also indicate to the level-2 trigger which thresholds the RoIs passed. In an ASIC which fully processes $4 \times 2 \times 2$ towers, there can be two of these 2×2 windows. In a 2×2 window there can only be one RoI, and it can pass any one of eight thresholds for the electron / photon trigger or any one of eight thresholds for the hadron / tau trigger. Therefore we require 42 bits to flag the RoIs.

Readout logic
The results (RoIs and hits) must be transferred to two destinations: level-2, and the DAQ via the RODs. Since the RoI information includes the hit information, only the RoI records will be written to the dual-port memory. The results will be written to the dual-port RAM on each bunch crossing. On receipt of a level-1 accept signal the data corresponding to the appropriate bunch crossing will be transferred from the dual-port RAM to the FIFOs. A four-bit parity error code (calculated from the data), a three-bit ID (eight ASICs per CPM) and the four least significant bits of the BCID number will be appended to the data, providing a 32-bit word to the FIFO. The FIFO data will be shifted out serially at 40 MHz, requiring 800 ns to transfer each word to the ROD.

There will be eight CPASICs and 40 Serializing ASICs on each CPM. The ROD specification (Section 6.3.10) describes how data are transferred from the CPASICs to level-2 and the DAQ.

6.3.7.4 Design and technology

The designs will be carried out at RAL using the Cadence design suite. Design techniques including full custom, semi-custom cell-based or gate arrays, as well as synthesis from a hardware description language such as VHDL, may be used. The design will be implemented on 0.6 μm (or smaller) geometry in CMOS technology capable of handling 160 MHz clocks, phase-locked loops and the required I / O.

The estimated gate count is approximately 200 k gates (gate = 2-input NAND), requiring approximately 194 signal I/O pins plus supply pins, which could be packaged in a 224-pin pin-grid array (PGA) or preferably in a surface-mount compatible plastic quad flat-pack package. Approximately 564 fully-tested packages will be required (including spares).

Only a very limited number of vendors will be willing to manufacture ASICs in this low volume. However, at least two — Austrian Micro Systems and Alcatel-Mietec — will be able to supply the technology required for the processor ASIC. Other technologies such as gate arrays and laser-programmable gate arrays (LPGA) will also be considered before a final decision is made.

6.3.8 Cluster Merger Module (CMM)

The trigger-hit results from the CPASICs will be first counted on the CPM using eight FPGAs, and the results transferred to CMMs for further counting. Figure 6-27 shows a block diagram of the logic required on the CPM.

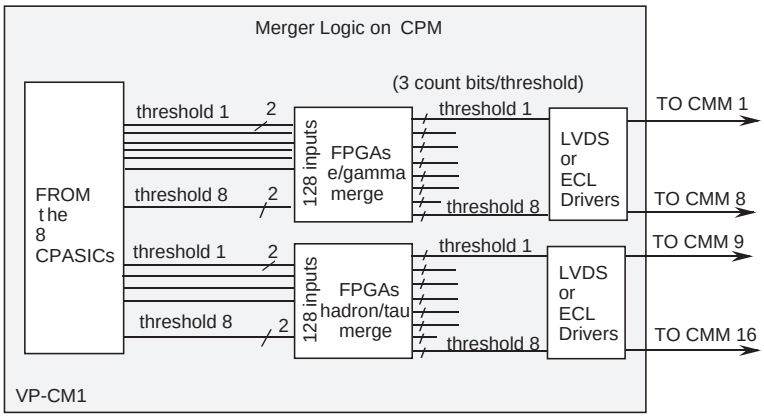

Figure 6-27 Merger logic on the CPM.

There will be 8 (thresholds) × 2 bits of information from each CPASIC and for each of the electron/photon and hadron/tau triggers which will require further processing (counting) before the final multiplicity result (three bits per threshold) is delivered to the CTP. This must be performed in real time and will contribute to the overall level-1 latency. Four FPGAs per trigger type will be used on the CPM to perform partial merging, yielding a maximum of three bits (seven hits) per threshold. A total of 48 (three bits × 16 thresholds) count bits will be transferred to the Cluster Merger crate and then distributed to the appropriate CMM. These signals will be transmitted to the CMMs using low-voltage differential signalling (LVDS).

Each CMM (see Figure 6-28) counts data for one threshold, so the CP requires a total of 16 CMMs. Each CMM will receive a total of 156 differential signals at 40 Mbyte/s from the 52 CPMs in the system. These signals will be transported to the CMMs through a suitable backplane connector (e.g. the 2 mm Metral connector system), and each CMM will then complete the counting for one threshold. The final three-bit multiplicities from all 16 CMMs will be transferred to the CTP using differential ECL or LVDS. The logic on the CMM can be implemented using fast LUTs to minimize the latency. For diagnostic purposes, the CMM input data and the results will be captured on spy FIFOs, which may also be used in playback mode for stand-alone testing.

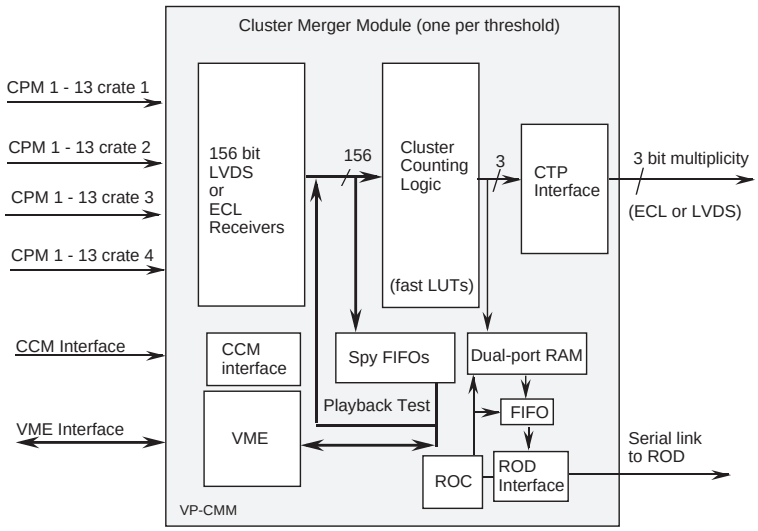

Figure 6-28 Block diagram of the Cluster Merger Module.

If the latency budget allows, the 3-bit count from the CPMs could be transferred to the CMMs at 160 Mbit/s using LVDS, which would reduce the number of signals and connector pins on both the CMMs and the CPMs and provide the additional bandwidth for a parity bit if required.

6.3.9 Clock and Control Module (CCM)

The CP consists of four crates of CPMs, one of CMMs and one of RODs. Each crate will have a CCM to provide the interface with the TTC system and the DCS (via the CAN bus). The timing and control signals are generated centrally (TTC system) and are distributed via fibre-optic cables to the experiment electronics. An ASIC (TTCrx) will receive and decode all signals distributed through this system. The signals will be brought directly from the TTC system using 18-fibre cables — one fibre per crate (6) plus one per ROD (12).

6.3.9.1 Interface to the TTC system — clock distribution

The clock distribution electronics will have the following functionality (see Figure 6-29):

- each CCM will receive TTC information on an optical fibre from the TTC system;
- the TTCrx ASIC will be used to decode the signal;
- the decoded TTC signals will be distributed to the processor modules on the backplane using differential ECL;
- the two-phase LHC clock, each phase of which can be individually delayed, will be distributed separately to the two halves of the crate to minimize clock skew.
- only clocks and control signals will be distributed, and distribution of parallel data avoided;
- all parallel data, such as bunch-crossing counter numbers, can be generated within the processor modules or within ASICs using the clock and control signals;

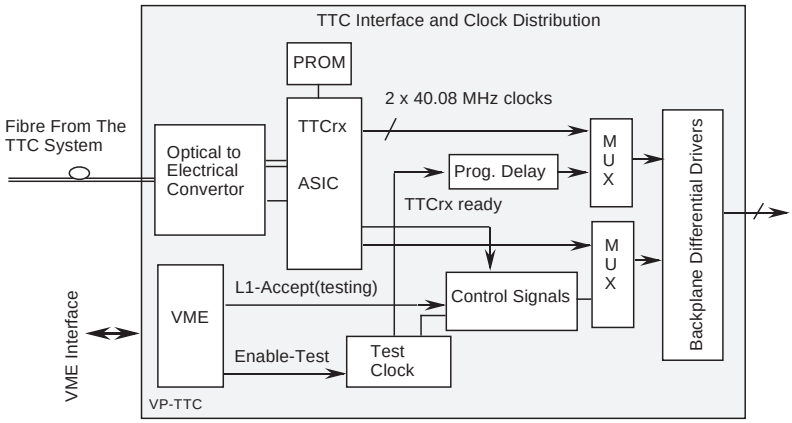

Figure 6-29 Block diagram of the TTC interface.

- a facility will be included to use local clocks and control signals when the TTC system is not available during testing of the processor modules.

6.3.9.2 Interface to DCS

All crates within the trigger system will provide access to the detector control system to monitor the power supplies, fans and the individual modules (where necessary) via the CAN bus. The micro-controller unit (MCU) on the CCM will provide the CAN bus interface to a network of nodes (‘crate nodes’), each interfacing to the sensors on the module plus the crate power supplies and fans. The ‘main bus’ node as shown in Figure 6-30 below will then interface the CCMs in the system to the DCS via a separate CAN bus. The MCUs and the dual-port RAM will provide the bridge between the ‘crate bus’ and the ‘main bus’. A block diagram of the CCM is given in Figure 6-31.

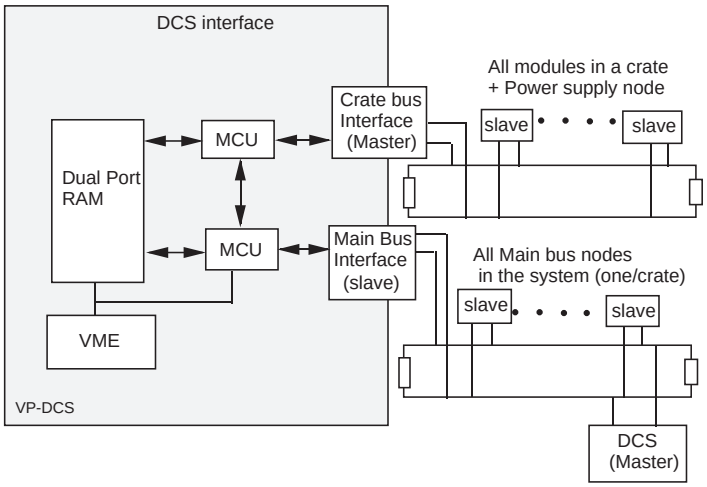

Figure 6-30 Block diagram of the DCS interface.

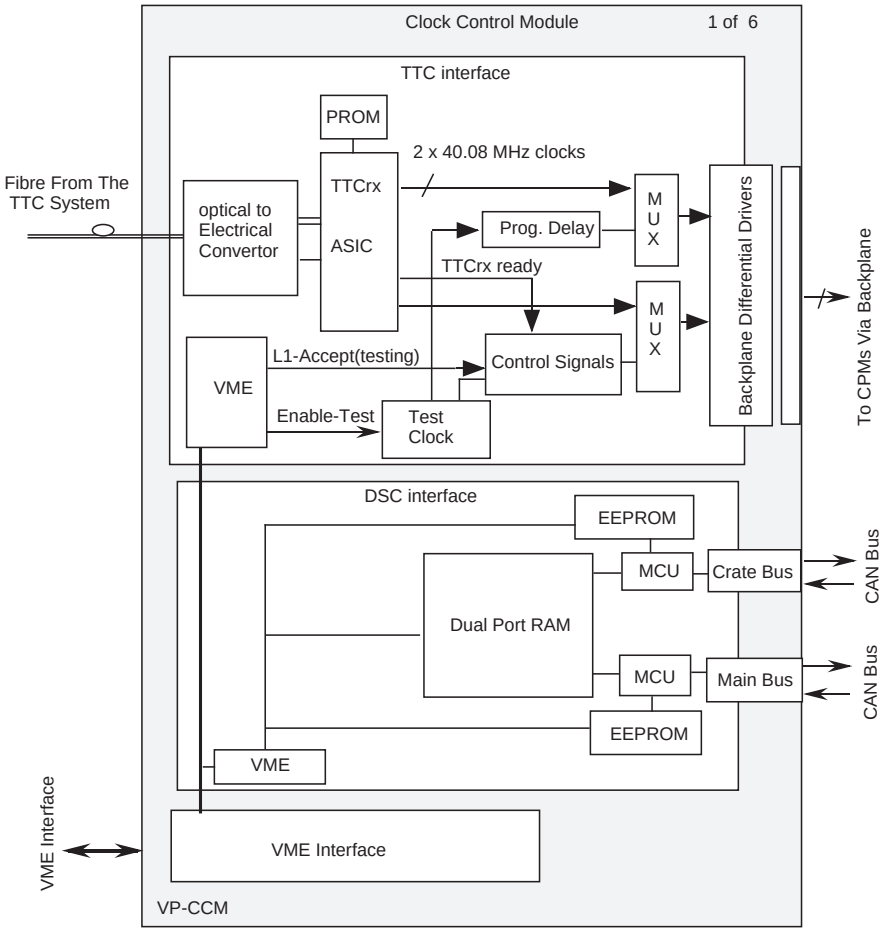

Figure 6-31 Block diagram of the Clock and Control Module.

6.3.10 Readout Driver (ROD) system

The following is a description of the ROD module which collects data from the Serializing ASICs on the CPM for the DAQ, and the RoI collector module for collecting RoI data from the CPASICs on the CPM for the level-2 processor. Both modules can be of the same design, and the methodology used for reading out data from the CPM makes the DAQ readout and the RoI transfer schemes the same. The Jet/Energy-sum Processor will use the same ROD design.

The data from the CPMs will be transferred to the ROD/RoI module (from now on referred to as the ROD) using a G-link (or equivalent) high-speed serial link. This provides flexibility and scalability, since the ROD need not be in the same crate as the processor modules. The number of RODs required will depend on the data volume and time available to transfer the data. Another advantage in having the RODs in a separate crate is that the addition of a CPU module will enable software cluster-finding on a sample of data to constantly monitor the performance of the system. S-links are the current ATLAS standard solution to carry the data from the RODs to the appropriate destinations.

6.3.10.1 Readout functions

For each crate of 13 CPMs, DAQ data from the Serializing ASICs could be handled by two RODs (so eight are required in the system), and RoI data from the CPASICs could be handled by one ROD (so four are required in the system). These numbers are based on the requirement that a data throughput of one event per 10 μs must be maintained to the ROB and to the level-2 trigger, and no form of data compression or zero-suppression (ZS) on this data is assumed.

The ROD will:

- receive data from the CPMs at 800 Mbit / s;
- compare the event numbers and bunch-crossing numbers with the on-board TTCrx generated numbers. If no match is found, it will use the TTC to reset the counters;
- perform ZS on the RoI data (and possibly also on the DAQ data);
- transfer the RoI data to level-2, and the DAQ data to the ROB in less than 10 μs.

6.3.10.2 Readout process

The process of collecting and transferring the data from the CPMs starts from the Serializing ASICs. Figure 6-32 shows a block diagram of the readout logic.

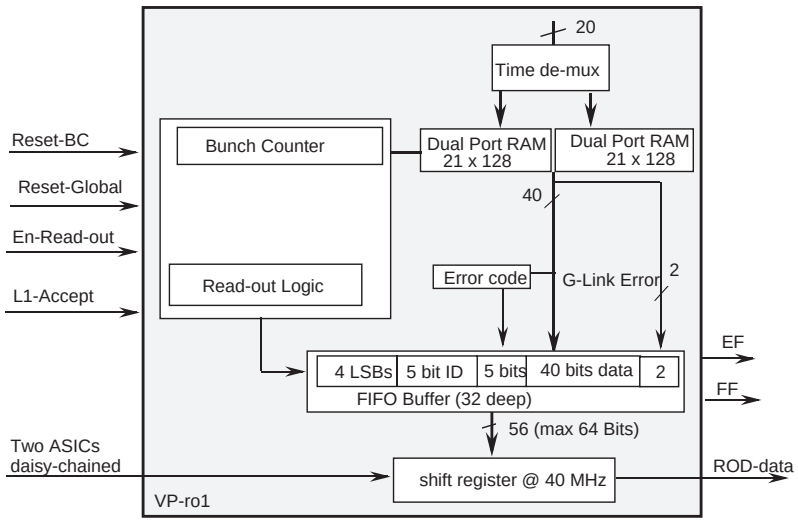

Figure 6-32 Readout logic on the Serializing ASIC.

For each bunch crossing, the data will be written to the dual-port memory. For ease of readout, two dual-port RAMs are cascaded to capture the BC-multiplexed data which arrive in two consecutive time slices. On receipt of a level-1 accept signal, the data associated with that particular bunch crossing will be copied to the FIFO together with information such as the ASIC number (40 ASICs per CPM) and an error code. The data will be shifted out of the FIFOs at 40 MHz. Two ASICs can be daisy-chained together to form 20 shift registers, so that all 40 ASICs on a CPM can be connected to a G-link, as shown in Figure 6-33. Another option would be to use two G-links without daisy-chaining, thereby halving the transfer time to the ROD from 2.8 μs to 1.4 μs.

Using a G-link transmitter chip, the readout shift registers on the ASICs can be connected to bits D0 to D19 of the 20-bit G-link word, while the flag bit can be used as an additional data bit for transferring the crate ID, module ID, etc. This serial data will be framed using the data available (DAV) signal of the G-link, as shown in Figure 6-34.

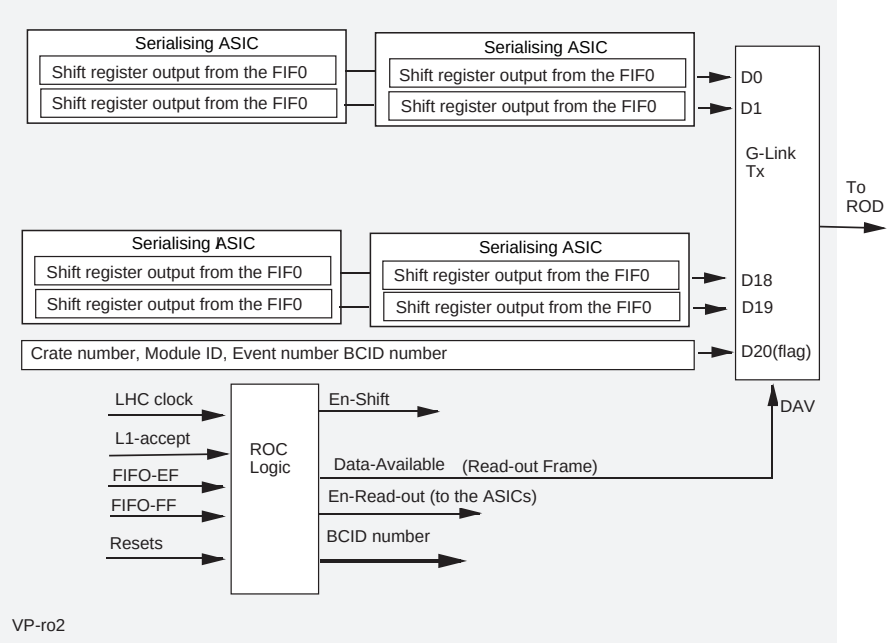

Figure 6-33 Readout logic on the CPM for the Serializing ASIC.

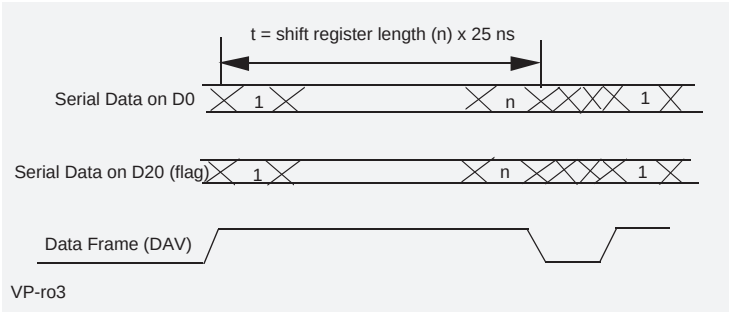

Figure 6-34 Framing serial data.

The readout control (ROC) logic on the CPM will merge the ASIC data with the crate ID (2 bits), module ID (4 bits), BCID number (12 bits) and event number (24 bits), and transfer it out of the CPM. The BCID number and the event number will be generated within the readout controller logic, and will be queued on the controller on every level-1 accept signal and transferred out with the readout data frames.

The ROC logic on the CPM will control the sequence of events by asserting the readout enable on the ASICs when it receives a level-1 accept, generating DAV for the G-link and terminating the sequence when all bits have been transferred. The entire process will operate at the 40 MHz LHC clock rate and will transfer data from the ASICs to the ROD without any data compression or ZS. The ROC will repeat this process until the FIFOs are all empty.The ROC logic will also

generate the BCID and event numbers, and will monitor the ‘full’ flag of the FIFO, indicating overflow via the Detector Control System. This condition should not occur if the CTP and the readout process of the CP are operating correctly.

The readout logic for the CPM RoI data in the eight CPASICs will be identical to that of the Serializing ASIC described above, except that two FIFOs will service the two windows. The data will again be shifted out serially, as shown in Figure 6-35. The output of these shift registers will be connected to a G-link (D0 to D15), with the crate ID, module ID, etc. connected to D20 (flag bit) as shown in Figure 6-36, and the data transfer to the ROD will be the same as described above.

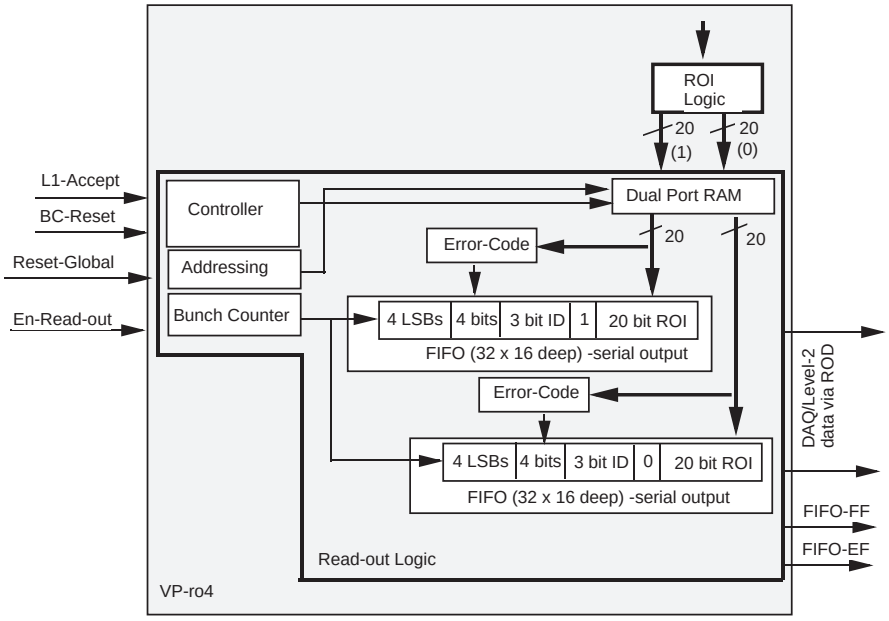

Figure 6-35 Readout logic on the CPASIC.

6.3.10.3 ROD module

Figure 6-37 shows the building blocks of the ROD. Using G-link receiver chips, the ROD will receive data and buffer them until requested by the DAQ, or by level- 2 with a RoI request. The entire process will be handled by the ROD control logic, which will monitor the DAV signal of the G-links and initiate the data receiving process. In parallel, zero suppression will be done (if required) and valid data written to the data buffers. Under the control of a single-bit token, these data buffers will be read out and the data transferred to the ROB and level-2. A selection of these data may be copied to another buffer for local monitoring. For example, the data could be transferred from all the RODs in the CP to a local computer unit to carry out ‘software’ cluster-finding and to report errors via the DCS.

On receipt of a level-1 accept signal, the control logic will place the event number and the BCID number generated by the TTC logic into the event / BCID buffer. When it receives the DAV from the CPMs, the ROD control logic will transfer the event and BCID numbers from the buffer to the FIFOs, and then enable all ZS logic. The control logic will also check the received BCID and event numbers against the locally derived numbers. If there is a mismatch then it must instruct the TTC system to reset all the event counters in the system. When ZS has been completed and the control logic receives a data request, a token will be sent to the first channel which is ready

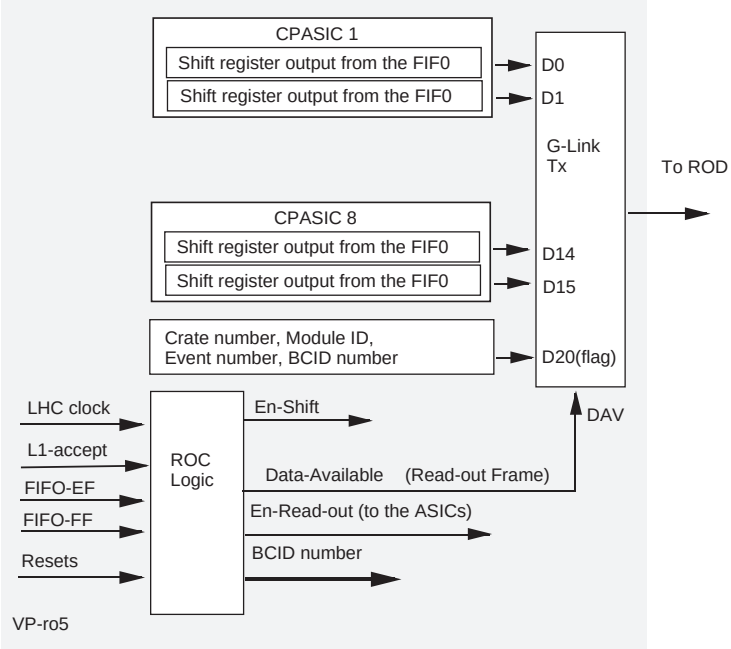

Figure 6-36 RoI-associated logic on the CPM for the CPASICs.

to transfer its data. All channels will then transfer their data under token control, until the last channel in the chain returns the token to the control logic to complete the readout sequence. The control logic will also monitor the XOFF signal from the ROBs to halt further data transfers if the ROBs become full.

The ZS logic will take control signals from the control logic, perform ZS and then write the data to the data buffers as 32-bit words, together with the number of valid words. If required, an error checking code such as a check-sum could be implemented on the same device as the ZS logic and written to the data FIFO as part of the complete data frame to be transferred out (see Figure 6-38 for the data format).

The ZS unit will inform the controller when it has completed its task (DONE). Since the time to perform ZS on each channel can vary, and also since the control logic is local, to save time in transferring data to the ROB or to level-2, a ‘first come, first served’ token-passing scheme could be used instead of a ‘sequential’ token-passing scheme. In such a scheme the control logic will monitor the DONE signals of the ZS units and pass the token to the first ZS unit which has completed its task, and others will be queued and given the token when it becomes free. This scheme is more fault-tolerant than a sequential token-passing scheme, since if a channel becomes faulty it may easily be bypassed. Features such as time-out can also be added. However, the order of the data might change from event to event, so this requires further discussion. Since all the control functions and ZS are performed by programmable logic, the ROD can evolve to accommodate other requirements, e.g. indicating RoI overflow to level-2.

6.3.10.4 Timing and scalability

The above scheme is fully scalable, so that the readout time to level-2 or DAQ may be reduced by increasing the number of RODs. Other combinations are possible, such as having more than one S-link per ROD. For example, the data volume from all the Serializing ASICs from a CPM is 2560 bits (64 bits $\times$ 40 ASICs). If the ASICs are daisy-chained and the data from all the ASICs are

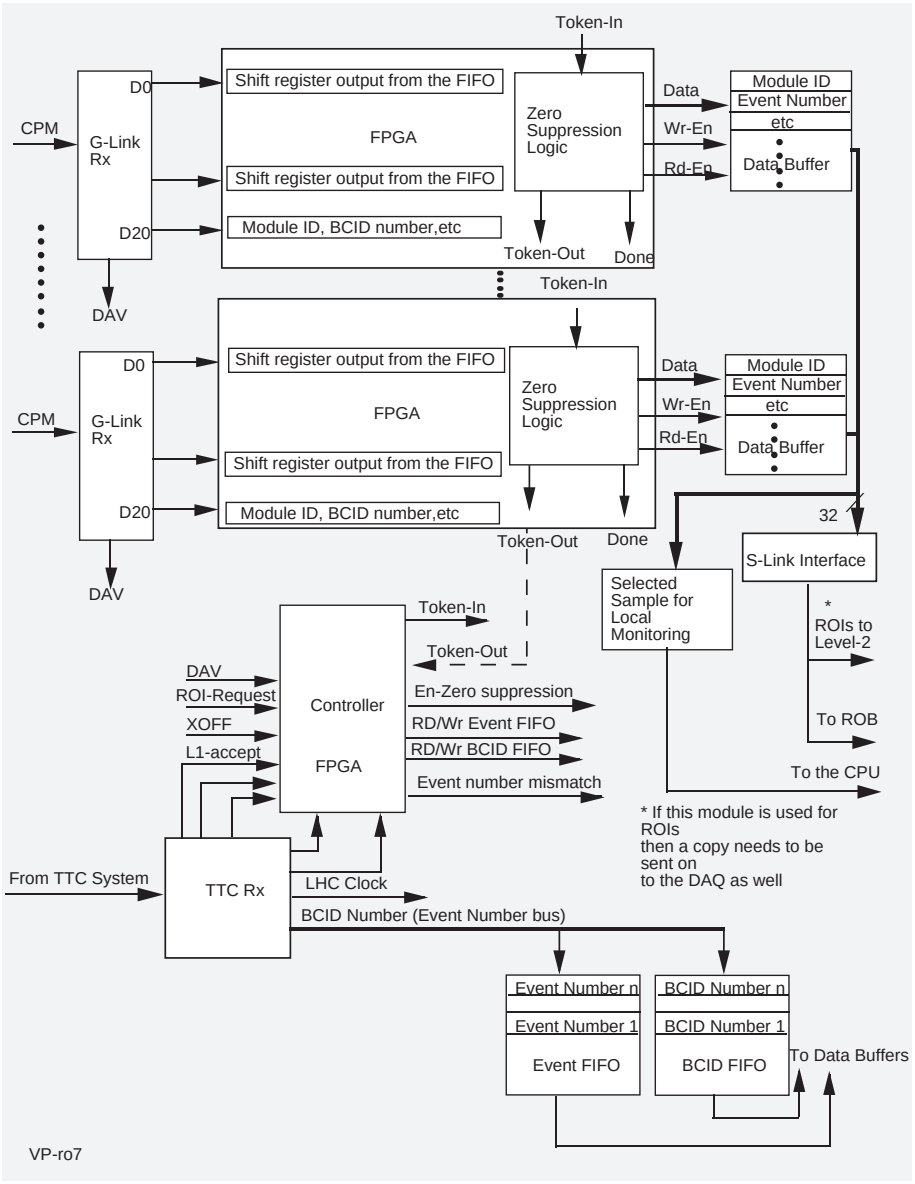

Figure 6-37 Block diagram of the ROD.

transferred in parallel using a G-link as described above, it will take 3.2 μs (128×25 ns) to get the data from the CPMs to the RODs. The data volume on the ROD will be 33,280 bits ($64 \times 40 \times 13$ CPMs / crate). The maximum rate at which the ROD can transfer data out is governed by the speed of the S-link. The S-link can operate at 32 bits every 25 ns, taking 26 μs to transfer 33,280 bits to the DAQ. Therefore, if one ROD per crate for the Serializing ASICs is used, then it will take approximately 26 μs (no zero-suppression applied) to transfer the data to the DAQ. Table 6-4 summarizes the time and the number of RODs required in the CP without any zero-suppression or compression applied.

It can be seen from Table 6-4 that to meet the 100 kHz level-1 rate, the first two options shown are not viable. Using the last two options, 12 RODs will be required in the Cluster Processor. Eight RODs, with two S-links each, are required for data transfer from the Serializing ASICS to the DAQ, and four more RODs for the collection of the RoIs and transfer to level-2. The RoI data

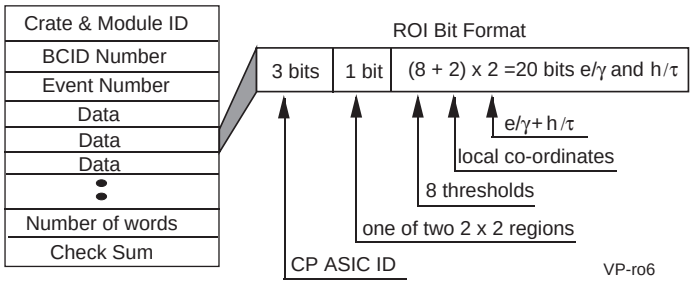

Figure 6-38 DAQ data format for RoIs.

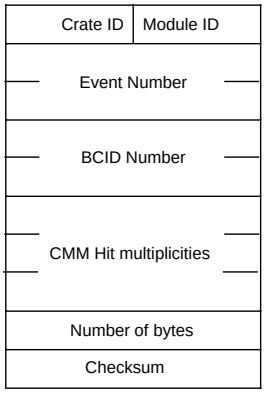

Figure 6-39 DAQ data format for CMMs.

Table 6-4 RODs required in the Cluster Processor for various options.

ASIC	RODs/crate	Total RODs	Time (approx.)
Serializing ASICs	1	4	26 μs
Serializing ASICs	2	8	13 μs
Serializing ASICs (2 S-links per ROD)	2	8	6.5 μs
Processing ASICs (RoIs)	1	4	5.2 μs

also need to be transferred to the DAQ, which can be easily achieved by duplicating the output signal and sending the same information to two destinations.

6.3.10.5 Testing

The dual-port memories on the Serializing ASICs can be loaded with test data, and with the appropriate signals supplied by the readout control logic the complete RoI and DAQ chain can be tested.

6.3.10.6 Cluster Processor DAQ formats and data volume

The data volume transferred from the Cluster Processor depends on the rate of non-zero RoI generation, which is a function of luminosity, physics, backgrounds, and threshold settings. It also depends on the frequency at which data for trigger monitoring and checking is read from the CPM inputs. Using the data format illustrated in Figure 6-38, and assuming a high average of 10 RoIs per event (including non-trigger secondary RoIs), the RoI DAQ traffic would be 120 bytes per event, or 12.0 Mbyte/s at a rate of 100 kHz. It is preferable to assemble the data into a single event packet before transmission to DAQ in the RoI builder. In this case, the overhead is reduced and the data rate becomes 6.0 Mbyte/s.

Data from each CPM Serializing ASIC occupies 8 bytes, as shown in Figure 6-32. Assuming that all 40 ASICs on all 52 CPMs are read concurrently, each time-slice will generate 16,640 bytes of data. As an example, if one time-slice is read out for each turn of the LHC (88 μs), the data rate will be 180 Mbyte/s. It is therefore reasonable to use a separate S-link and ROB from each CP crate to allow the capability to read at least 2–3 time-slices for data alignment checking. It

should be possible to compress this data using the same encoding as in the Preprocessor, remembering that the BC-multiplexing will have sent zero data for about half of the channels.

Using the DAQ format in Figure 6-39, the 16 CMMs provide 16 bytes per event, or 1.6 Mbyte/s at 100 kHz. This includes the information used to generate CTP bits. It is most convenient to carry this on a separate S-link to a separate ROB even though the data rate is not high. The volume of DAQ data from the CP is summarized in Table 6-5 for the conditions just described.

Table 6-5 DAQ data from CTP results and the Cluster Processor.

Data source	Volume (Mbytes/s)	No. of S-links and ROBs
Copy of RoI information	6	1
CPM Serializer data	~100	16
CMM hit data	1.6	1
Total	**~107.4**	**18**

6.4 Jet/Energy-sum processor

6.4.1 Architecture

The Jet/Energy-sum Processor performs the jet, missing-E_T and total-E_T trigger algorithms on 0.2×0.2 electromagnetic and hadronic jet elements provided by the Preprocessor. The goal of the design is to create a system based on off-the-shelf components, with conservative design parameters and a high degree of flexibility.

The Jet/Energy-sum Processor consists of four 9U crates processing quadrants in ϕ of the jet/energy-sum trigger space, and receives its data from the Preprocessor on Jet and Energy-sum Modules (JEMs) which perform trigger algorithms for a region of 2×8 elements in η–ϕ. The jet algorithms provide options for three jet-cluster sizes (2×2, 3×3, and 4×4 jet elements), which all operate within a 4×4 jet-element window, analogous to the CP algorithms. Therefore, each JEM requires an environment of 5×11 elements. As in the CP, fan-out of environment data in ϕ is performed by duplicated links from the Preprocessor, while fan-out in the η direction is performed via point-to-point links on a custom backplane. The similarity in layout to the CP offers the benefit of a simplified organization of the Preprocessor, which must provide data and duplicated links to both systems.

Fifteen JEMs in each crate perform the full jet algorithm and E_T summation for the range $|\eta| < 3.2$, and a sixteenth JEM receives additional FCAL sums which are added to the sum-E_T acceptance region. (If it is decided to implement a separate FCAL pseudo-jet trigger, it may be possible to perform the analysis for both FCAL ends on this sixteenth JEM.) Each JEM receives separate e.m. and hadronic data for 2×11 jet elements, which arrive on 22 G-links running at 800 Mbit/s.

To provide the required environment for 2×8 jet windows, each JEM also receives 22 jet elements via a backplane from its neighbour in –η, and 11 elements from its neighbour in +η. At the same time, each JEM transmits 22 jet elements to its neighbour in +η, and 11 to its neighbour in –η, via the backplane. All data transmission on the backplane between JEMs is performed via point-to-point links at 80 MHz, and each backplane link is only one slot in length. At this data rate, five backplane lines per jet element are sufficient to transfer 10-bit data every 25 ns. Thus 330 I/O pins per JEM on the backplane are required to fan-out data for the jet algorithm.

Results of the E_T summation and jet algorithms on each JEM are transmitted to a Sum Merger Module (SMM) and a Jet Merger Module (JMM) in each crate by 40 MHz point-to-point links. Each JEM transmits three 12-bit words to the SMM containing the E_T, E_x,and E_y sums for that JEM, and eight 3-bit sums to the JMM containing the number of jet clusters passing each threshold. Additional bits flag overflows in the input data or calculations, or transmission errors detected on the G-link data. The merger modules of each type from three of the four crates transmit their output data to the fourth module, which performs the final merging and transmits the final results to the CTP.

Input data and RoI bits are stored in pipelines on each JEM. Upon a level-1 accept, event information is transferred to derandomizer buffers to be read out to the DAQ system and RoI builder. Input data from the Preprocessor are read out via one G-link transmitter per JEM to data RODS for readout to the DAQ, while jet RoI information from the JEMs and energy RoI information from the SMM are transmitted via an 80 MHz daisy-chain to a G-link transmitter on the JMM, which sends the data to an RoI ROD for transmission to the RoI builder.

Crate-level control and configuration is performed over a reduced VME bus. A small 3U-height TTC card receives an optical fibre from the TTC subsystem and distributes the TTC signals electrically in the crate. A serial CAN-bus system is used to monitor temperatures, voltage levels, and other important slow-control information in each crate. The overall organization of the jet/energy-sum crate is shown in Figure 6-40.

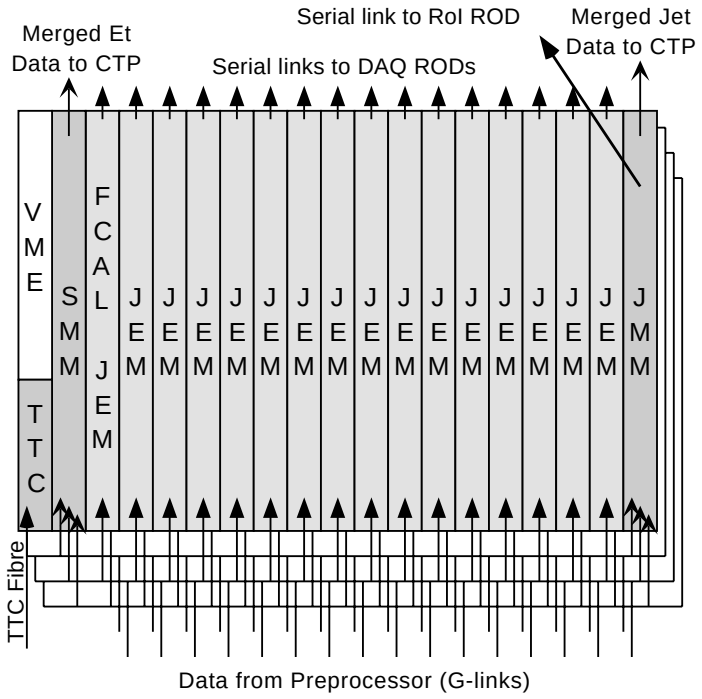

Figure 6-40 Proposed physical layout and data flow of the Jet/Energy-sum Processor.

6.4.2 Jet and Energy-Sum Module (JEM)

The heart of the Jet/Energy-sum Processor is the JEM, which receives data from the Preprocessor, calculates and forms sums of E_T, E_x and E_y, and identifies E_T clusters which pass the jet algorithm criteria. Fifteen JEMs in each crate cover the jet trigger acceptance region $|\eta| < 3.2$, while a sixteenth JEM accepts additional FCAL sums for E_T summing alone. A block diagram of the JEM is shown below in Figure 6-41.

Each JEM receives 22 jet elements on G-links. Sixteen of these jet elements correspond to the core 2×8 element environment covered by that JEM, while the other six are duplicated links from neighbouring regions across the ϕ-quadrant boundaries. The jet-element words have 9-bit resolution and a maximum value of 512 GeV, with a tenth parity bit used for error detection. If any input data word fails a parity check, its value is set to zero and an error flag is set.

6.4.2.1 E_T summation

The data from the G-links is parallelized to 40 MHz, and received by the Sum FPGAs. The 9-bit data words are checked for parity, and any input which fails a parity check is set to zero and an error flag is set. The e.m. and hadronic components of each jet element are then summed together, and E_x and E_y are calculated using LUTs.

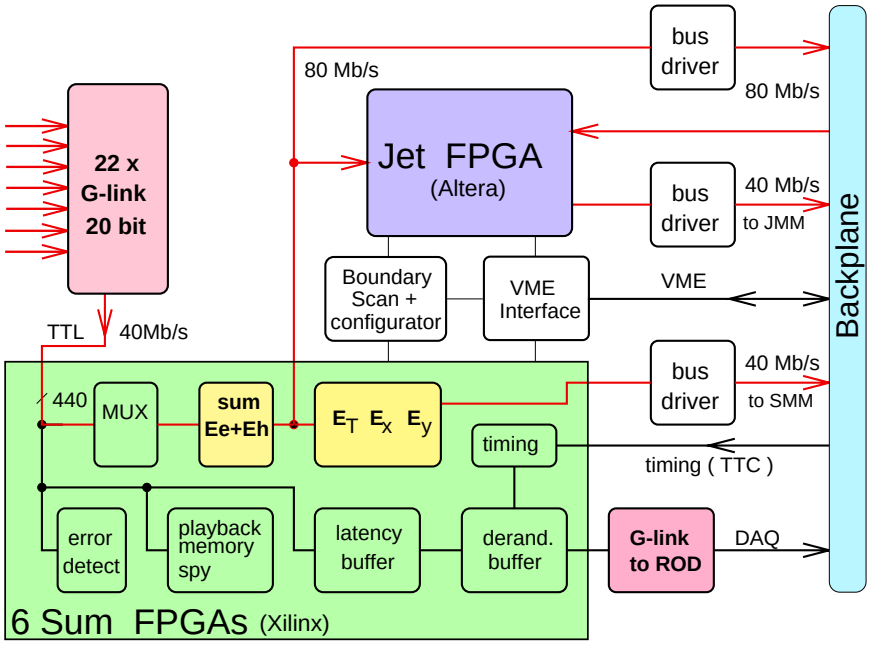

Figure 6-41 Block diagram of the JEM.

E_T, E_x, and E_y are summed for all core elements of the JEM. Because each jet/energy-sum crate corresponds to a single quadrant in ϕ, no sign bits are necessary for the E_x and E_y values. The initial pre-summing of the e.m. and hadronic energies reduces the number of LUTs necessary for the geometric calculations to convert to E_x and E_y, and reduces by half the amount of information which must be fanned out to the jet processor system. This reduction is made possible by the projective assignment of the jet elements, and the fact that the jet algorithms do not need to separate electromagnetic and hadronic energy deposits.

6.4.2.2 Jet element distribution

The 10-bit sum of electromagnetic and hadronic energies for each 0.2×0.2 jet element is multiplexed to 80 MHz and distributed to the Jet FPGAs. All 22 jet elements received directly from the PPMs by each JEM are sent to the Jet FPGA on that module. In addition, jet-element environment data is distributed to the neighbouring JEMs (22 elements in one direction, 11 in the other) over 80 MHz impedance-matched point-to-point links on the custom backplane. A unique Sum FPGA output pin will be assigned for each point-to-point link, either on-board or between modules, in order to simplify impedance matching and data synchronization. In the baseline design single-ended, series-terminated, low-voltage TTL drivers will be placed between the Sum and Jet FPGAs. An evaluation programme is being conducted to determine whether they are needed.

The choice of 80 MHz for jet-element distribution is made possible by the smaller number of shared signals between JEMs relative to the situation in the Cluster Processor. Since the Jet FPGAs internally process data at 80 MHz, an additional demultiplexing step which would have been introduced by a higher data rate is also avoided, reducing one additional source of latency.

The number and choice of FPGAs for the E_T summation is determined by the number of I/O pins necessary to distribute signals to all of the on- and off-board Jet FPGA inputs. The summation and missing-E_T calculations may be performed at 40 MHz, and only the

multiplexing and transmission of summed jet elements to the jet processor requires an 80 MHz clock. A baseline choice of devices would be six Xilinx XC4013XL [6-3] devices for initial data reception and processing. These devices would receive data from the 22 G-links, place them in buffers for later DAQ readout, perform the first two levels of energy summing, and transmit summed jet elements to the Jet FPGAs. The core jet element sums would then be passed to a larger Xilinx XC4062XL device for final processing and merging.

Transverse energy sums for each JEM are merged and transmitted by point-to-point links at 40 MHz to the Sum Merger Module (SMM) for crate- and system-level summing (see Figure 6-42). Additionally, input data and board-level energy sums are stored in pipelines for readout to DAQ, if requested, upon a level-1 accept.

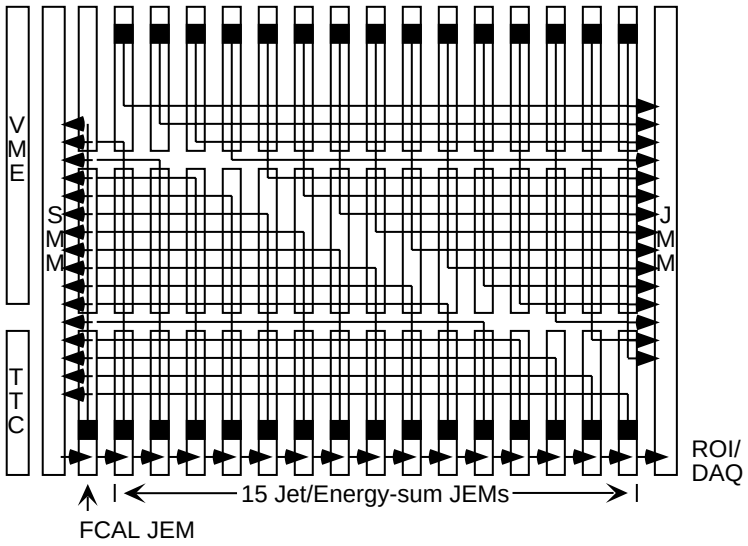

Figure 6-42 Proposed implementation of point-to-point links from the JEMs to the Jet Merger Module and the Sum Merger Module. Level-1 results are transmitted to each of the merger modules at 40 MHz, while jet and energy RoI information is transmitted via a daisy-chain at 80 MHz from the JEMs and the SMM to the JMM for readout to the RoI ROD.

6.4.2.3 Jet processing

Each JEM processes a core area of 2×8 jet elements within an 5×11 environment required for execution of all foreseen jet algorithms. These 55 jet elements constitute 550 data bits per 25 ns bunch-crossing interval, which are received at 80 MHz on 275 input pins of a large FPGA. The jet algorithm calculations are performed on the multiplexed data using 5-bit arithmetic functions between flip-flops clocked at 80 Mhz. This approach reduces the latency of the trigger algorithm, and also reduces the amount of routing and logic resources required in the FPGA.

As it is possible for up to four RoI objects to be identified as local maxima by each JEM, 3-bit sums for each of the eight thresholds are transmitted at 40 MHz to the JMM for crate- and system-level summing, as shown in Figure 6-42. Additionally, the RoI information is stored in a pipeline for readout to the DAQ and RoI builder upon a level-1 accept.

To execute the jet algorithm, an FPGA must be selected which has a large number of I/O resources, adequate logic to accommodate the full jet algorithm over all 16 jet windows per JEM, and sufficient speed to perform 80 MHz calculations. Initial studies indicate that the baseline jet algorithm can be implemented on larger SRAM-based FPGAs such as the Altera

Flex 10k250 [6-4], which has up to 250,000 usable logic gates and over 450 usable I/O pins. Pipelined calculations with 5-bit adders and comparators might be run at 100 MHz or more by a sufficiently high-speed device. This provides a comfortable margin for safely running an 80 MHz algorithm. A final choice of FPGAs will be made by evaluating a VHDL model of the full trigger algorithm on suitable devices from different manufacturers.

6.4.2.4 DAQ and RoI readout

Upon a level-1 accept, event information is transferred to derandomizer buffers to be read out to DAQ and the RoI builder. The JEM input data are transmitted directly to data RODs via a single G-link transmitter per JEM, while jet RoI information and energy-sum RoI numbers are read out via a daisy-chain to a G-link transmitter on the JMM.

Since each JEM receives its data over 22 G-links, input data for one bunch-crossing may be read out by a single G-link to the data ROD in 22 time slices, or 550 ns. This is much faster than the maximum level-1 accept rate, and makes it possible to read out data from multiple bunch-crossings if necessary.

The RoI daisy-chain consists of eight input and eight output pins on each JEM, and just eight output pins on the SMM. Additional point-to-point links between neighbouring modules signal when the previous module in the chain has emptied its current buffer, and a three-bit derandomizing bus ensures that all data from a single event have been transmitted before the next event is read out. RoI data are multiplexed on the daisy-chain so that each RoI data word occupies two 80 MHz cycles. Assuming that all possible RoIs are read out via an eight-bit wide daisy-chain, the time to read all RoIs from a single event to the RoI ROD is approximately 1.6 μs.

6.4.2.5 Timing and control

In the baseline design, a TTCrx chip on each JEM receives a 40 MHz clock signal, level-1 accepts, and control signals from the TTC system. The TTCrx chip is capable of providing two additional 40 MHz clocks with adjustable delays relative to the system clock, allowing different stages of the data path to be properly timed in. The board-level clock distribution is performed at 40 MHz, and clock signals are locally refreshed and multiplied to 80 MHz by PLL-based devices such as the ICS AV9170 clock synchronizer and multiplier [6-5].

6.4.2.6 JEM I/O connectors

Although the JEMs are 9U boards, space on the rear edge for backplane connections is an issue because of the large degree of fan-out necessary in the jet algorithm. The baseline choice of connectors is the AMP Z-pack HM 2 mm family, as shown in Figure 6-43. These connectors provide five signal rows plus two outside rows with ground-return shields, and have become a standard for CompactPCI backplanes. Equipping the backplane with five B-type connectors and one M-type connector provides up to 680 I/O pins plus three 40 A power pins. To reduce crosstalk and allow faster rise times, it is foreseen to populate the central row entirely with ground pins, in addition to the two dedicated ground return shields on either side. This leaves 544 I/O pins which can be used for data fan-out, merging of level-1 results, DAQ and RoI readout, the VME interface, and timing and control signals.

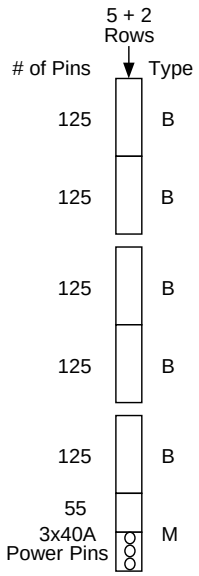

Figure 6-43 Configuration and pin counts of JEM card edge, using Z-Pack 2 mm-pitch connectors.

6.4.2.7 Backplane

The Jet/Energy-sum Processor backplane is a single 9U-high custom board of width 43 cm (21 slots). Most of the connectivity on the backplane consists of point-to-point links, although a reduced VMEbus and a serial bus for slow control are also provided. Slots for JEMs, the SMM, and the JMM are equipped with 2-mm-pitch connectors, as well as 40 A power pins for +5 V and +3.3 V distribution. A standard 6U VME slot is provided for the VME controller card, and appropriate connectors must also be chosen for the TTC card.

Short point-to-point connections between neighbouring modules for data fan-out to the jet FPGAs and for daisy-chain readout of DAQ and RoI information are run at 80 MHz, while longer point-to-point links from the JEMs to the JMM and SMM run at 40 MHz.

The backplane will be implemented in approximately 14 layers, with ground planes between each of the signal layers. All signal line-widths must be chosen to provide impedance-matched interconnections between modules.

6.4.2.8 Configuration, testing, and monitoring

The baseline control system for the JEM is a VME slave interface, which controls address and data lines for configuring and monitoring registers on all devices on the module. A JTAG boundary scan interface which can be controlled through the VME may be used to test and reconfigure the Sum and Jet FPGAs. A CAN-bus interface is used for monitoring local temperatures on each board.

Trigger parameters which change on a regular basis, such as energy thresholds and contents of LUTs, are implemented in read/write registers accessible by the VME. More fundamental changes in the trigger algorithms which require a reconfiguration of the FPGAs may be implemented in situ using the JTAG interface.

The JTAG interface may also be used to download special test configurations designed to test module and backplane interconnections in the system, and to provide in-circuit tests of functions within the FPGA.

6.4.3 Jet Merger Module (JMM)

The JMM receives counts of jet clusters identified by each JEM over each of eight thresholds. Twenty-four bits of results are received across 40 MHz point-to-point links from 15 JEMs (360 bits total), and sent to eight 80 MHz adder trees for crate-level summing. The outputs of the adder trees are truncated to three bits per threshold, giving 24 output bits per crate. The crate-level sums from three of the JMMs are sent as 40 MHz signals on 24 twisted-pairs each to the fourth JMM, which combines them with its own sums to produce the final level-1 jet multiplicities, which are sent to the CTP at 40 MHz on twisted-pair cable.

To reduce the numbers of different modules produced, all four JMMs are identical. In addition to receiving crate information through backplane connectors, each JMM also has one 30-pair ECL output connector and three 30-pair ECL input connectors. The extra two pairs are used for an accompanying ECL clock, an error detection bit, and overflow and error flags. The ECL output connector can be used either for crate-to-crate transmission of intermediate jet data, or output data to the CTP.

Crate-level merging of data is performed by a single, large FPGA, most likely identical to the jet processor FPGAs on the JMMs (Altera 10k250). The system-level merging is performed using high-speed lookup tables, implemented either using discrete 4 kbyte RAM chips or an FPGA solution. It is not necessary to install these components on JMMs which are not performing the system-level merging.

In addition to crate- and system-level merging of jet data, the JMM is responsible for collection of RoI information from the JEMs in the crate, as well as energy RoI information from the SMM. A G-link transmitter sends this information to the jet/energy-sum RoI ROD. Figure 6-44 illustrates the data flow in the JMM.

6.4.4 Sum Merger Module (SMM)

The function and overall structure of the SMM is similar to that of the JMM, although with different data requirements. As in the JMM design, the merging algorithms are implemented on FPGAs, and results are shared between crates via 80 MHz differential ECL signals. A diagram of the SMM data flow is shown in Figure 6-45.

The SMM receives three 12-bit sums (E_T, E_x, and E_y) at 40 MHz from 16 JEMs, making 576 bits of data in all. Additional bits may also be necessary to flag overflows or errors in the adder trees or input data. If such a large number of backplane connections proves unworkable, the data may either be transmitted to the SMM at a higher rate, or two SMMs per crate may be necessary.

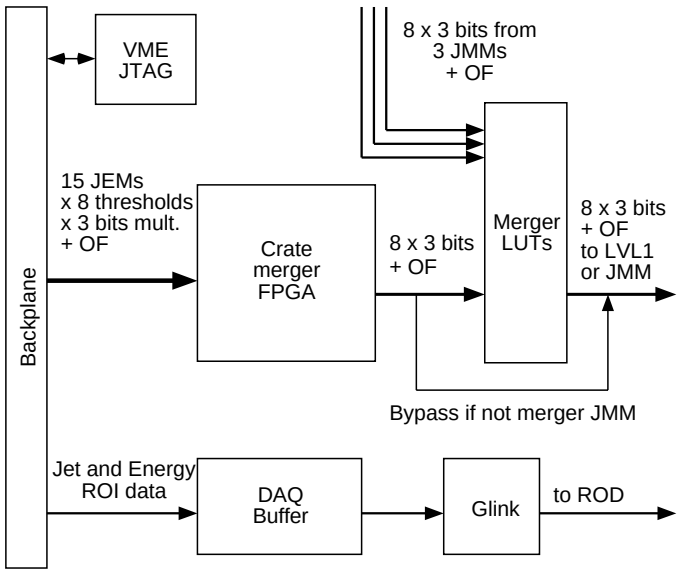

Figure 6-44 Data flow in the Jet Merger Module.

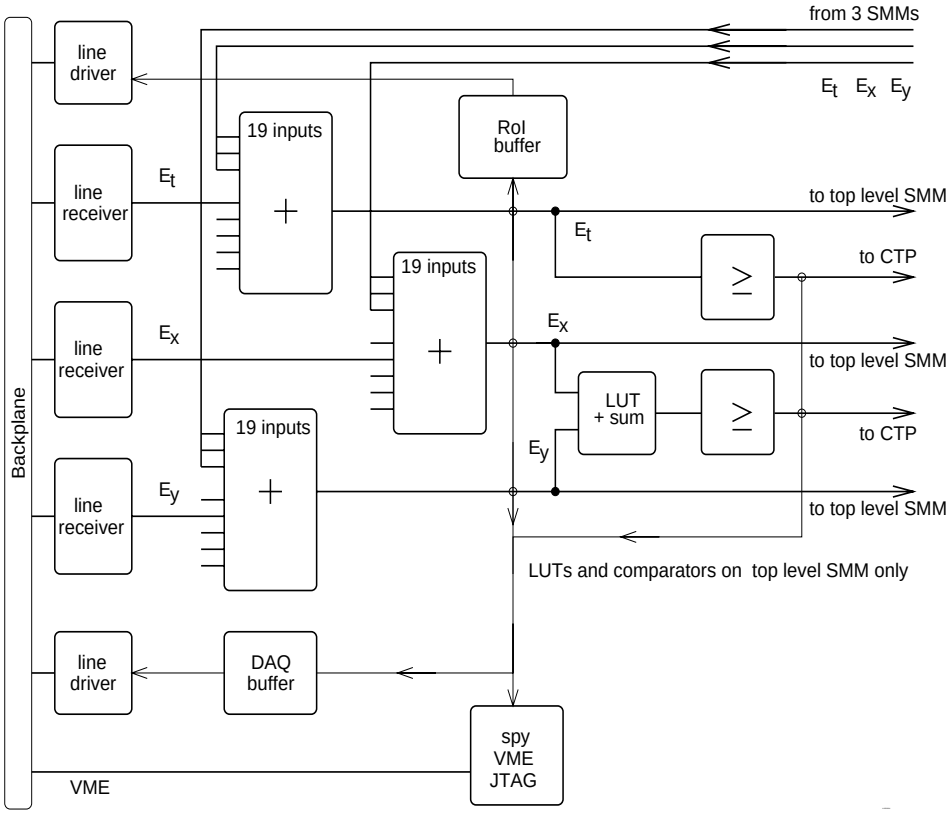

Figure 6-45 Data flow in the Sum Merger Module.

The crate-level sums are multiplexed to 80 MHz and transmitted over 20 twisted-pairs to the SMM assigned to merge the system-level data. The final results are sent over a 40 MHz twisted-pair cable to the CTP.

6.4.5 Results to Central Trigger Processor

Results of the jet and energy-sum algorithms are reported to the CTP on separate 40 MHz differential-ECL (or LVDS) ribbon cables. The energy-sum results are reported in the form of bits indicating which total-E_T and missing-E_T thresholds were passed, while the jet results consist of eight 3-bit multiplicities of jet clusters passing each threshold requirement. If a total jet-E_T calculation is also implemented (see Section 6.4.8 below), additional result bits from the jet trigger will indicate which total jet-E_T thresholds were exceeded.

6.4.6 Regions of interest

Upon a level-1 accept, the positions of jet clusters identified by the Jet/Energy-sum Processor are transmitted to a ROD that transmits RoIs to the level-2 trigger (see Section 6.3.10). Up to four RoIs may be identified by each JEM, and each RoI must contain a bunch-crossing ID, 10 bits to uniquely define the position of the RoI in the calorimeter, and eight bits corresponding to the eight possible thresholds which the RoI may have passed.

In addition to the jet RoI positions, the E_T, E_x, and E_y sums calculated for the level-1 trigger are also made available to the level-2 trigger, in order that the missing energy vector can be checked against any level-1 muons in the event, which are unseen by the calorimeter trigger.

6.4.7 Data acquisition, monitoring, and test facilities

The G-link input data, the JEM-level E_T, E_x, and E_y sums, and the JEM-level threshold results can also be read out to DAQ if desired. There are no plans to provide intermediate-level data for testing data transmission on the module and backplane, due to the prohibitive amount of data this represents. However, the JTAG interface will be used to download special test configurations to the FPGAs during periods with no beam, which will allow the connections at all stages of the data path to be thoroughly tested. The ROD modules used by the Jet/Energy-sum Processor for data readout to the RoI builder and the DAQ will be a common design shared with the Cluster Processor. See Section 6.3.10 for more information.

6.4.8 Total jet transverse-energy trigger

It is currently proposed that a simple trigger on the total E_T in jets be implemented in the CTP, as described in Section 15.2.3.1. However, we mention here an alternative hardware solution within the Jet/Energy-sum Processor. This would use the same simple multiplicity-based estimator discussed in Section 4.6.2, and is currently being investigated to see whether it brings added benefits and is feasible.

To execute this trigger algorithm ‘properly’, it would be necessary to alter the basic operational principle that there is no requirement to retain the E_T of identified trigger objects so that an adder tree could sum the total E_T of all jets passing some threshold. This would entail a huge extra I/O burden as well as unacceptable latency. However, simulation studies described in Section 4.6.2 indicate that an estimate of the total jet transverse energy of an event may be made simply by using the multiplicities of clusters passing the different E_T thresholds used for the jet algorithm. Studies of this seemingly-crude estimator indicate that it works surprisingly well, even for rather unevenly-spaced jet thresholds. Moreover, there is little difference between

using three-bit and unrestricted jet multiplicities, and adding the results. This estimate would then be passed through another lookup table and converted to result bits to be transmitted to the CTP. It is estimated that adding this functionality to the JMM would add at least another two BCs of latency to the jet trigger logic.

6.5 References

6-1 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, 1994.

6-2 B. Niemann, Diplomarbeit Universität Heidelberg: *Datenkompression für die Auslese des ATLAS Level-1 Triggers*, IHEP 98–02, December 1997.
http://www.ihep.uni-heidelberg.de/Veroeffentlichungen/ps/niemann.diplom.ps.gz

6-3 Xilinx Data Book, 1998.

6-4 Altera Flex 10k Embedded Programmable Logic Family Data Sheet, version 3, January 1998.

6-5 Integrated Circuit Systems, Inc.
http://www.icst.com/

7 Calorimeter trigger demonstrator programme

7.1 Introduction

The ATLAS level-1 calorimeter trigger group has used the ATLAS test beam every year from 1992 to 1998 as a test bed for processes and technologies which will be necessary in order to design and build the final calorimeter trigger for ATLAS. These test-beam operations form an overall programme of performance testing which demonstrates the feasibility of all the demanding parts of the system. Initially, design and test-beam work was performed as part of the R&D programme RD27, before this developed into the current ATLAS and CMS trigger programmes.

As already described, the ATLAS calorimeter trigger requires pipelined processing of a large number of trigger towers, and extensive data fan-out in order to perform the physics algorithms. All this has to be performed in a very short time (typically ~40 bunch-crossings from raw data to final result). In order to develop hardware which could achieve these goals, it was decided to test out ideas and electronics in the harsh reality of a test-beam environment, using prototypes of the final calorimeters as data sources.

The philosophy of the demonstrator programme [7-1] was to produce a small-scale version of the trigger system which contained all of the most important elements of the final design. In particular, each of the stages of the process which were identified as being critical should be performed in the demonstrator system. Alternative technologies were tried for various parts of the system throughout the programme in order to assess which would form the best basis for the final solution, so it was envisaged that some ideas would have to be abandoned, but that by the end we would have enough pointers to make the most reliable design, and sufficient evidence that this design would provide the performance needed for the final ATLAS level-1 trigger.

7.2 Brief history

7.2.1 Data source

For the test-beam runs from 1992 to 1995, data were mainly recorded using the RD3 liquid-argon (LAr) accordion calorimeter prototypes in the North Area at CERN. The signals were summed to form a convenient 6×6 array of detector towers, which were similar in arrangement to the final trigger towers in the ATLAS detector. This was a very useful geometry for testing the cluster-finding aspect of the trigger algorithms. The major difference from final ATLAS is that the area covered by these towers was smaller than the eventual tower sums of the LAr calorimeter, meaning that the noise seen in the test-beam set-up is about one-half of that which may be expected at ATLAS.

In 1996 and 1997 signals from the module-0 Tile Calorimeter were used. In this case, the geometry was not so convenient for the hard-wired trigger algorithms, but it still served as a useful insight into how the system responded to different calorimeter signal pulse shapes.

There was also significant testing performed in the laboratory at RAL during 1998 using pseudo-random input data.

7.2.2 Data acquisition technique

The means to investigate the performance of the hardware was based on a fairly simple data acquisition setup. Data flowed through a sequence of modules, and wherever appropriate, were clocked into circular memories on each 'bunch-crossing' tick. When an interesting event occurred (usually triggered by the standard beam-line scintillators), the whole system was stopped in order to be read out. This meant that the progress of a particular sequence of data could be followed through the various elements of the processing chain. Clearly this is not the way readout will be performed in the final system, but it gave a simple way of testing the integrity of data processing without developing buffered readout.

Online analysis could quickly check for some errors, but the data were also recorded to tape for later offline analysis. Thus the data integrity throughout the system could be determined, allowing measurement of bit-error rates and latency stability. Unfortunately, in terms of statistics for very low bit-error rates, this comparison of data offline only provides a fairly limited verification. Most of the results quoted come from this sort of analysis, and so are statistics-limited. Other techniques, as detailed later, were developed to test some aspects of the system in real time and obtain a better estimate of the system performance.

7.2.3 Test-beam summary

The first RD27 test-beam occurred in November 1992 [7-2], although at that stage very little custom-built hardware was available. LAr calorimeter pulses were digitized and recorded in order to study pulse shapes for BCID. In April/May 1993 [7-3] custom-built FADCs were used and cluster-finding was performed on the 36 channels using a module containing nine cluster-finding (CF) ASICs. This system was further enhanced in November 1993 [7-4], when BCID was performed on a single channel via a new BCID module. This system was used again in September 1994 [7-5] when data from prototype tile-calorimeter modules (RD34) were also taken. Most of the important results from the earlier years of test-beam running were summarized in a NIM paper [7-6]. A typical system set-up is illustrated in Figure 7-1.

In 1995, use of BCID was extended to all channels in a newly designed module using FPGAs, and this was used to study the performance of different algorithms. In the same year new purpose-built FADCs were used to produce the digitized data. Again the 36 channels were passed to the cluster-finding modules [7-7]. This was the end of the phase-1 demonstrator programme.

The phase-2 demonstrator was designed to closely emulate the final system, and in particular simulate the various data transmission and connectivity problems that will eventually be encountered. Figure 7-2 shows the full phase-2 system that was built by the end of the test-beam work. Completely new BCID modules and electromagnetic (e.m) cluster-finding modules had to be designed for this phase. The first tests of the new system were made in June 1996, but at that stage the most useful results came from investigations of analogue data transmission systems. More data were taken in September 1996 [7-8], when problems were discovered in the backplane technology used. This led to some redesign work before the October 1997 run and subsequent laboratory tests. These runs included several other modules to exercise new aspects

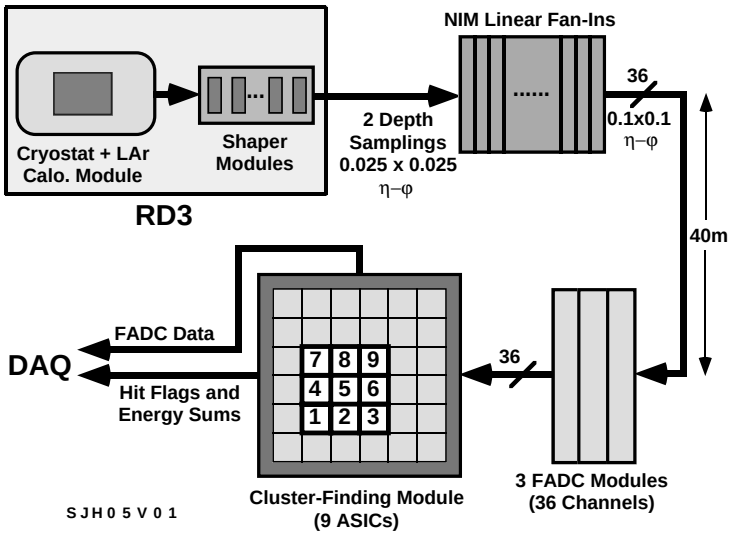

Figure 7-1 Block diagram of the liquid-argon calorimeter and demonstrator trigger system in the test-beam runs of early 1993.

of the system, such as jet processing and integration with the Central Trigger Processor demonstrator.

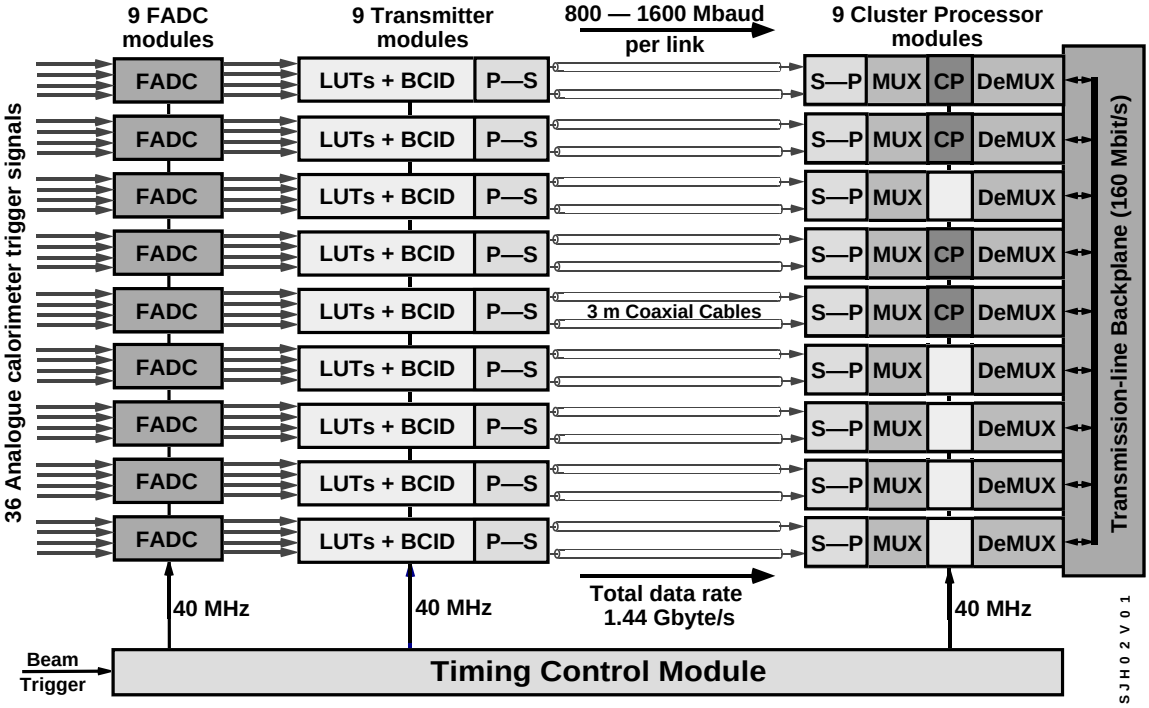

Figure 7-2 Schematic diagram of the full 36-channel phase-2 demonstrator.

7.3 Pipelined trigger algorithms

The trigger algorithms described in Chapter 4 need to be applied to all the data at 40 MHz to generate trigger flags for each individual bunch-crossing. The two most intensive processing tasks are the e.m. cluster-finding algorithm and the jet algorithm. Both of these require information from many neighbouring trigger towers and combine the data in a relatively complex way to produce several triggers. It was necessary to show that this can be done in parallel on all channels reliably at 40 MHz with a small latency. This was done in pipelined logic with either custom-designed ASICs or FPGAs, as described below.

7.3.1 Cluster-finding ASIC

One of the first hardware tasks was to implement the e.m. cluster algorithm in a CF ASIC. A cut-down algorithm, using only 16 inputs (one layer of towers) and only two thresholds for each energy requirement, was implemented in ASIC RAL114 [7-9], and was ready for testing in the 1993 test-beam periods. At the time that the ASIC was designed, the LHC bunch-crossing time was foreseen as being 15 ns, so the ASIC was designed to work up to at least 67 MHz. In fact, it was shown to work correctly at speeds above 70 MHz, limited only by other on-board components. The latency of the process was six ticks for central cell energy summing, and seven ticks for the final trigger decision.

7.3.2 Phase-1 cluster module

The original test module for the CF ASICs was a 9U VME board which housed nine ASICs in order to fully process the nine 4×4 arrays contained within the full 6×6 trigger-tower environment of the early test-beam work. Data were brought into the module across the backplane, but with only one CF module, backplane traffic was not an issue for this stage of the demonstrator. In Figure 7-3a, the energy sum computed by the CF ASIC is plotted against the expected sum (using the LAr data) for a typical run. As can be seen, the two agree perfectly. Note the bunching of data around three typical values, corresponding to pedestal energy, pions, and electrons. Figure 7-3b shows the sharpness of the cluster-finding algorithm (with the isolation threshold switched off), as should be expected with a digital threshold.

Analysis concentrated on the correct performance of the ASICs, which operated at speeds in excess of 70 MHz [7-6]. So with the real LHC clock frequency of 40 MHz, and the progress in ASIC technology since the design of this ASIC in 1992, it should be possible to implement the full algorithm with a shorter tick latency than with this prototype ASIC by performing more complex operations at each step. The CF ASIC was also used in the phase-2 demonstrator, but the focus was by then on backplane issues, as the correct functioning of the ASIC had already been established. Its hit outputs were however used in the full-slice tests (see Section 7.6).

7.3.3 Jet processor module demonstrator

In 1997, work began at Stockholm on the jet trigger. This differs from the cluster-finding architecture in that it is proposed to use programmable logic chips rather than ASICs. The JPMD was designed to test FPGA implementation of the jet logic and to integrate with the existing demonstrator modules so that the full processing chain could be demonstrated as a complete ‘slice’ of the final system.

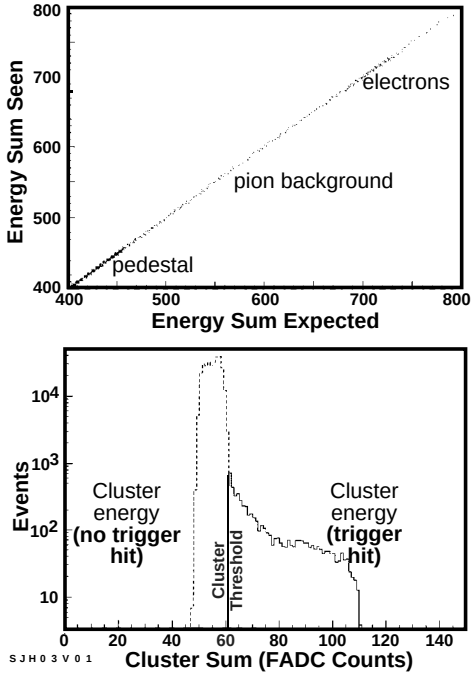

Figure 7-3 Analysis of the cluster-finding ASIC performance: (a) observed vs. expected energy sum, (b) sum for events failing and passing threshold.

The JPMD was a 6U VME board which received 16 8-bit tower sums. It analysed these sums using a pipelined processing algorithm implemented on a commercial FPGA. The JPMD used a 411-pin Xilinx XC4000EX FPGA to process the 128 bits of input data and produced an 8-bit result to be sent to the CTP demonstrator (CTPD). In addition to the input and output data bits, eight address and eight data pins are also assigned to accommodate up to 256 internal registers for setting thresholds, configurations, etc. The processor FPGA could be configured by a download cable from the front panel, by an on-board EPROM, or remotely through the VME interface.

The baseline processor algorithm used at the test beam was implemented using four sets of adder trees, whose results were each compared with three thresholds which are individually programmable by the VME interface. The main elements of the design (adders and comparators) will also be the main elements of the final jet processor algorithm implementation. The input and output data lines of the JPMD were shared by 12 ns RAMs, which are used to capture and read out event information for diagnostic purposes. The RAMs could also be used to program and play back simulated data for testing various elements of the trigger.

The JPM demonstrator board was successfully integrated into the level-1 slice test during the 1997 ATLAS test-beam runs, and again at RAL in February 1998. JPM software for the DAQ, diagnostic, and offline analysis programs were written and merged into the level-1 software, and the full data chain from the CPMs to JPM to CTP was shown to work.

A more detailed study of the JPM functions, however, showed an unacceptably high bit-error rate ($\sim 10^{-6}$) in the JPMD jet algorithm. Analysis of the data to pinpoint the source of the errors is still in progress, but some design issues have emerged which have influenced the planning of the next jet / energy-sum demonstrator:

- The board-level timing and control of the JPMD by an FPGA, which produced all read / write strobes from a single 40 MHz external clock. This gave inadequate control over the relative timing of different components on the board, and some devices (such as the jet processor FPGA) received clock signals with an asymmetric duty cycle.
- The jet processor FPGA (Xilinx 4036EX-4) was selected before the full programming and routing software was available for it. Timing analysis of the jet algorithm showed a maximum speed of 40 MHz, with no margin for error. This, combined with the timing issues above, may account for the error rates observed.

The jet / energy-sum 'minus-2' demonstrator, to be produced in June 1998, will resemble the current jet / energy-sum proposal much more closely than the JPMD, and is being designed with the above issues in mind. The clock distribution will be performed using discrete, high-speed components, and will feature independently-adjustable 40 MHz clocks which will be locally refreshed and multiplied with discrete PLL-based devices. Test configurations have already been written for the Altera devices selected to perform the jet algorithms, and simulations show that the 80 MHz algorithms should be able to run at over 100 MHz, giving a far better timing margin than that of the JPMD.

7.4 Bunch-crossing identification

Although the exact raw data format was not clear at the start of the programme, it was always foreseen that the level-1 trigger would have to perform bunch-crossing identification (BCID). Analogue techniques have been used at previous experiments, but it seemed a good idea to consider digital filtering as an alternative. Therefore the programme included the digitization of the prototype calorimeter signals and an investigation into implementing digital BCID. The digitized signals were useful for offline analysis and optimization of digital filtering techniques, and the BCID hardware was built to demonstrate the reliability of real-time BCID.

All of the demonstrator BCID implementations were based on an FIR filter combined with a simple peak-finder. This algorithm has been implemented in several modules over the course of the test-beam programme. It should be noted that in each case, the full matched filter has not been implemented — only a simplified version with a limited number of components, and with each component having a limited accuracy.

7.4.1 FADC module

In the early test-beam runs, various FADCs developed from previous experiments were used, but from 1995 onwards, new FADC modules custom-built by the Heidelberg group were used [7-10]. They ran at 40 MHz, with four channels per 6U module and generating 8-bit ECL outputs for each channel. They also contained a 256-byte memory for data capture, playback, and test-data injection. These have been used as the digital data source from 1995 to the present.

Much data was gathered with these modules over the course of three years, and we now have considerable confidence in their reliability. There was one genuine problem observed in 1995 and 1996 data where there would be a latency slip of one tick on about 1% of all events, but this problem was fixed in early 1997.

7.4.2 Single-channel BCID module

The first BCID board was built in Birmingham in 1993. It simply implemented the BCID algorithm in discrete hard-wired logic. Variable FIR components were implemented via lookup tables (LUTs) so that different filters could be used. Analysis documented in [7-4] shows that this initial attempt to implement BCID performed at or near 100% efficiency at the required LHC frequency.

7.4.3 Phase-1 BCID demonstrator module

The next stage was to build a full 36-channel system and implement the BCID logic in Xilinx FPGAs incorporated on a BCID demonstrator. Modules to do this were built at RAL, and a detailed description can be found in [7-7]. Each 9U board contained three Xilinx chips, each of which could process four data streams. Thus each board dealt with 12 channels, so that three modules were needed for the full system. The Xilinx logic could not cope with an infinite variety of filters, but several simple filters were implemented and tested. One useful feature of the Xilinx implementation was the capability to switch the Xilinx into a ‘transparent’ mode, where the filter and peak-finding were disabled. This allows data to flow unchanged through the system, making data comparison simpler.

Analysis of the test-beam results showed that the Xilinx implementation was, within limits of low statistics, 100% reliable. This meant that BCID algorithm performance and subsequent cluster-finding performance could be analysed from the test-beam data [7-7].

7.4.4 Phase-2 developments of the BCID module

The main phase-2 BCID module was the Transmitter Module (TXM), which did not develop the BCID implementation further, but concentrated instead on data transmission. This aspect will be described in more detail later. The internal BCID performance of the TXM in the 1996 and 1997 test-beam periods was entirely satisfactory. Within the limits of statistics available, there were no errors in clocking the data into the TXMs, or with the Xilinx and BCID performance. Results from these tests suggest bit-error rates at less than 10^{-8}. Figure 7-4 shows the TXM in BCID mode correctly flagging (with the same latency) a particle which deposited energy in two adjacent trigger towers.

BCID implementation was developed further towards a final ATLAS solution first in a Heidelberg BCID test ASIC [7-11], then in the Heidelberg Front-end Module (FEM) [7-12]. This module encapsulated most of the functionality of the TXM, and added the extra feature of buffered data readout, which is closer to the readout mechanism in the final system. On this module, the BCID algorithm and data buffering were implemented in an ASIC, called the front-end ASIC [7-13], again designed by the Heidelberg group. The FEM was commissioned in 1997 and used in the October test-beam set-up. Results confirmed that the data transfer and BCID processing of the system was very reliable, with bit-error rates measured at less than 10^{-7} (limited by statistics). Some data were also recorded using the buffered readout mechanism. In this case verification of data integrity was difficult, since very few elements of the system had the capability to perform buffered readout, so there was little scope for cross-checking. However, some results were obtained by running with the Central Trigger Processor Demonstrator, as described in Section 7.6. Earlier laboratory tests in 1997 [7-14] showed that the fast readout mode had no bit errors in 10^9 bits when tested at level-1 accept rates up to 128 kHz.

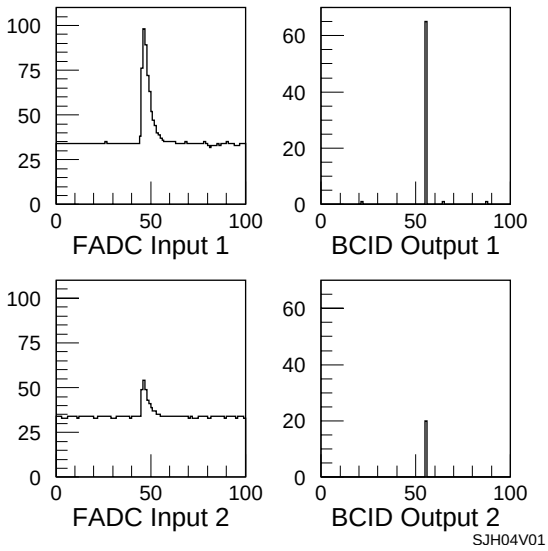

Figure 7-4 Two TXMs performing BCID on the same particle, seen in two adjacent channels of the Tile Calorimeter.

Occasional errors started to occur at higher rates, but 128 kHz is already higher than the ATLAS specifications (75–100 kHz).

7.5 High-speed data transmission

The trigger processor must be split over many crates of electronics, meaning that fast digital data transmission between modules and crates is a potential problem area. For crate-to-crate transmission, several different digital options were considered, including both electrical and optical solutions [7-15]. In order to minimize input pins, serial links running at both 800 MBd and 1600 MBd were investigated at the test beam, the results confirming that the option of G-links running at these rates with copper cables is a safe baseline.

For the hardware that performs the trigger algorithms, the module-to-module interconnectivity within a crate becomes an issue. Both the cluster and jet-finding algorithms require massive fan-out of data between modules. For Cluster Processor crates, it is foreseen that this will be done using a fast backplane. In order to keep pin-counts down, the data will be fanned out serially using single-ended 160 Mbit/s bit-streams between boards. This will therefore require both a reliable multiplexing/demultiplexing process, and a dense backplane that will work with high data integrity at these speeds. One of the major tasks of the phase-2 demonstrator was to demonstrate that this is feasible. In fact, the current processor architecture means that the backplane will now be simpler than envisaged at the time of the demonstrator programme. In the final design, modules will only have to communicate with two nearest neighbours, whereas the demonstrator backplane had to transport signals between modules up to four slots apart. Therefore, proof that the demonstrator backplanes work means that the ATLAS trigger backplane should be easily achievable.

7.5.1 Inter-crate serial links

In the ATLAS system, a fast serial link will be needed from the Preprocessor Module to the Cluster Processor Module. To investigate various options for this link, a Transmitter Module (TXM) was designed at RAL. This implemented four-channel BCID on a single-slot 6U VME board. The data transmission aspect was separated off onto a replaceable single-slot daughter-board, allowing for different options to be tested. The original specifications can be found in Ref. [7-16]. The data were received on the phase-2 Cluster Processor Modules (CPM) described below.

The original daughter-board, used in 1996, used the HP G-link chip-set (HDMP-1014D) for data transmission. This could be run at either 800 MBd or1600 MBd, allowing the four channels of 8-bit data to be transmitted on two or one link(s) respectively. The link could be either a coaxial cable or an optical fibre via a commercial optical converter. The details of the implementation of the various options are shown in Figure 7-5.

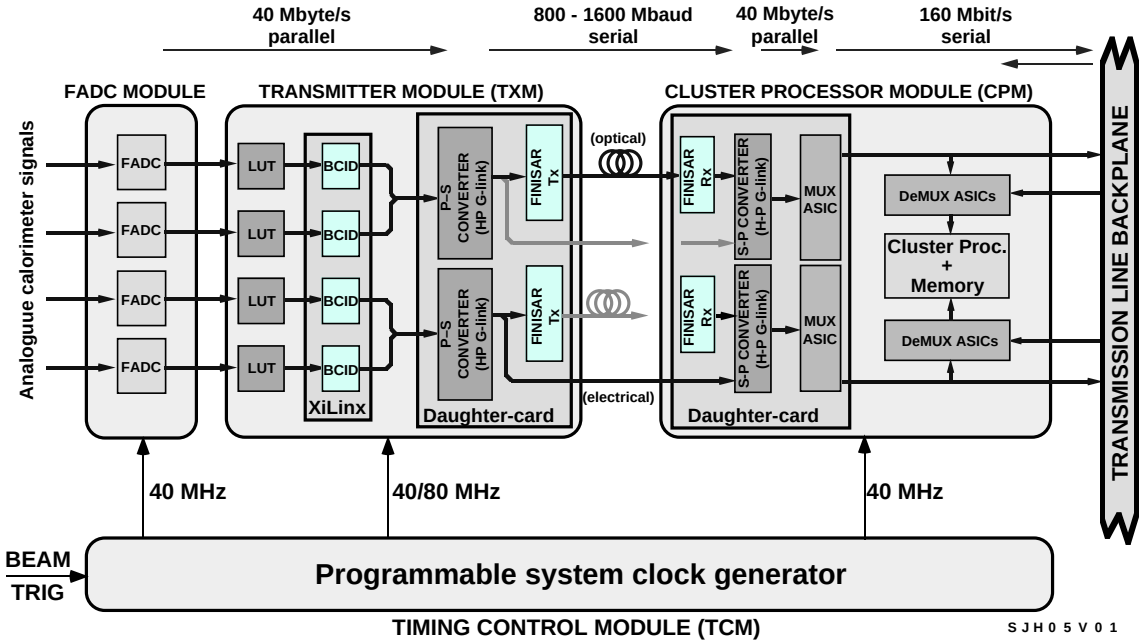

Figure 7-5 Schematic diagram of the 1996 demonstrator.

Early tests in 1996 suggested that only optical transmission at 800 MBd performed well. All other modes had unacceptably large bit-error rates, or failed to work at all. However, as the problems of reliably driving G-links came to be understood, particularly the importance of reducing power and ground noise by good isolation, better results were obtained. Probably the best performance seen in the September 1996 test-beam run was no errors in about 10^{13} bits with optical transmission at 800 MBd. The experimental set-up in shown in Figure 7-6. Electrical transmission was best when no VME access was being made to the modules for the purposes of data readout. With event readout taking place, the best performance gave bit errors at one in 10^{10}. After the test-beam period, the 1600 MBd mode was also made to work, and electrical isolation of the G-links gave confidence that there was no need for optical transmission. For more detail on results see Ref. [7-8].

Using insight gained in the 1996 run, new daughter-boards were designed for the 1997 run. These were built with the possibility of the two speed modes, but no optical link capability. This

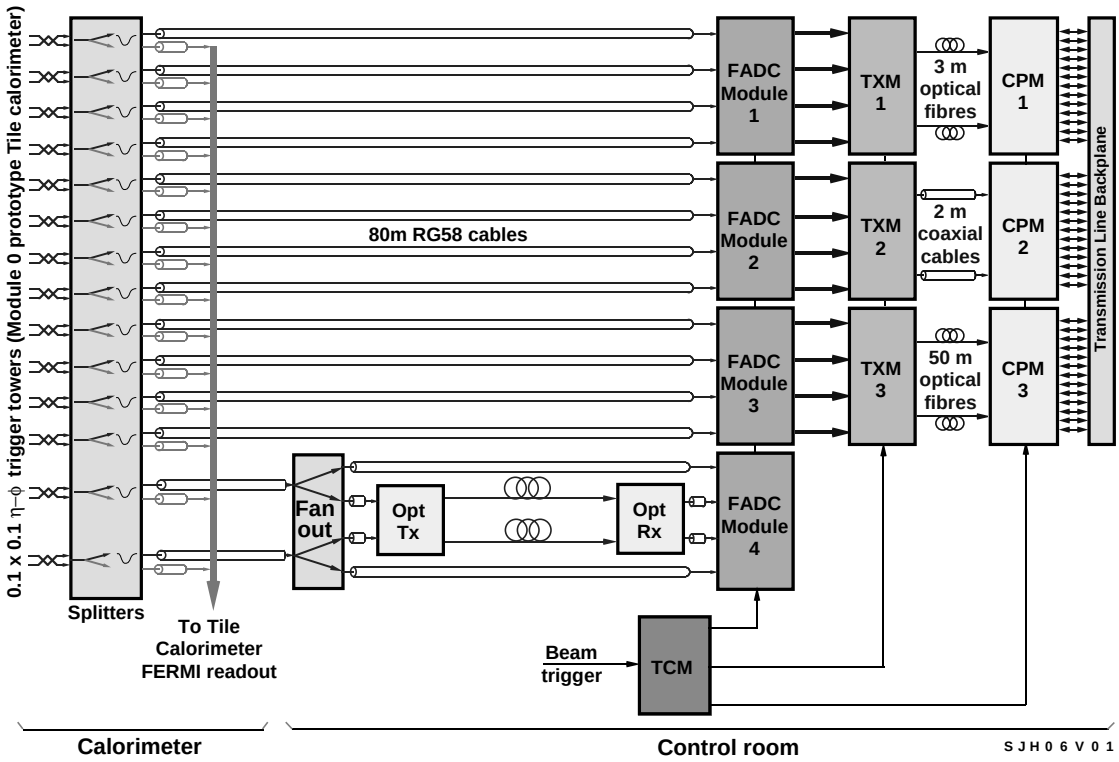

Figure 7-6 Experimental set-up in 1996, showing the optical test and the reduced 12-channel demonstrator system.

system was seen to work very reliably in the 1997 test-beam run, with typical bit-error rates of less than 10^{-15} in the 800 MBd mode, as measured by the G-link internal error-checking. Some low-level problems were observed with the 1600 MBd mode, but again these were associated with events being read out. From the behaviour of the system with different cables, it was inferred that the choice of connectors and cables was also paramount to the reliable performance of G-links, so careful decisions will have to be made with the choice of connectors in the final system.

7.5.2 Cluster Processor fast backplane

The feasibility of the Cluster Processor crate backplane communication was a critical feature that needed to be demonstrated, to which end the backplane density in the phase-2 demonstrator was made to be approximately the same as that then envisaged for the final system. In order to achieve the level of backplane fan-out required without having an enormous number of pins, a single 40 MHz 8-bit signal was multiplexed into two 160 Mbit/s streams. These were transmitted as single-ended signals across the backplane. This is a factor of eight saving over differential 40 Mbit/s signals. However this format clearly requires more care and a custom-designed PCB backplane. Careful timing, signal isolation, and good grounding were needed to prevent high levels of data corruption and crosstalk. The full cluster-finding crate consisted of nine Cluster Processor Modules (CPMs) communicating via the backplane. Each module had to output data from a maximum of 21 trigger towers onto the backplane, and receive up to 21 trigger towers from the backplane. This adds up to 84 signal pins per CPM connecting with the custom backplane, plus several ground and power lines.

Another aspect of the backplane problem was how to perform the required multiplexing and demultiplexing. Again, custom-built components in the final system would be needed for this task. To demonstrate this operation, another ASIC was designed, known as the dual-function ASIC RAL163 [7-17]. As the name suggests, two of these could be used as a pair to perform both the 40 MHz to 160 MHz multiplexing, and also the conversion back down to 40 MHz. It also contained internal memory for test purposes and timing self-calibration to make the bit alignment easy to handle. These ASICs had been tested stand-alone in 1995, and were ready for integration with the CPMs and backplane in 1996.

7.5.3 Phase-2 Cluster Processor Module

The first CPM boards were two-slot 9U VME modules designed by Birmingham and RAL in 1995/1996 [7-18]. They had to receive data from the phase-2 TXMs so, in an equivalent way to the TXMs, they used daughter-boards to receive front-panel signals to allow testing of several protocol options. The signals received via the front panel, called the 'core cells', were first multiplexed up to 160 Mbit/s to be passed onto the motherboard, which then fanned them out onto the backplane and also into the core-cell CF ASICs. The motherboard also received fanned-out signal data from other CPMs via the backplane, which were also sent to the CF ASICs. Only four channels were received directly onto the CPMs, the others having to come from the backplane. In all 25 channels were available on each CPM, meaning that four full versions of the cluster-finding algorithm could be performed. Each CPM had four CF ASICs, and the whole system fully covered the 6×6 tower matrix. The layout of the channels meant that the maximum distance travelled by a signal over the backplane was four double-width VME slots. Trigger-flag output bits and the four core-cell trigger-tower data bytes were available on the front panel.

Three of these CPMs were tested in the test-beam run in September 1996. The set-up has already been shown in Figure 7-6. After stable G-link operation was established, the performance of the multiplexing and the backplane was investigated. Deficiencies in the timing set-up led to less than perfect performance in these first CPMs. The best-optimized performance seen at the 1996 test-beam for an individual trigger tower was no errors in about 10^8 bits [7-8], ignoring occasional latency slips due to the timing set-up. This relatively high data integrity was only observed on the trigger towers that were entirely internal to a CPM, confirming that the dual-function ASICs were performing well. A typical event, showing the flow of data through the whole system from FADC to CPM, is shown in Figure 7-7a.

The backplane, which used BTL signal protocol, was far more of a problem in the 1996 test-beam run. The best error rates observed were at the level of 5%, which was completely unacceptable. Despite this, many of the calorimeter pulses appeared quite clean across the backplane, since only errors in the higher-order bits have an obvious effect on the pulse shape, as is shown in Figure 7-7b. The signals themselves suffered bad reflections and crosstalk, and in general the BTL logic was found to have signal rise and fall times that were too long for 160 Mbit/s bit-streams arriving from multiple sources.

Because of the failure of the BTL version of the backplane logic to reach the required specifications, a second version of the CPM was designed in 1997. ECL logic was used for the backplane protocol, and timing improvements were included on the board. The other major difference was the capability to use an MCM on the CPM daughter-board to receive the serial signals from the TXM. This would form our first demonstration of the use of MCM technology (Section 8.4.4).

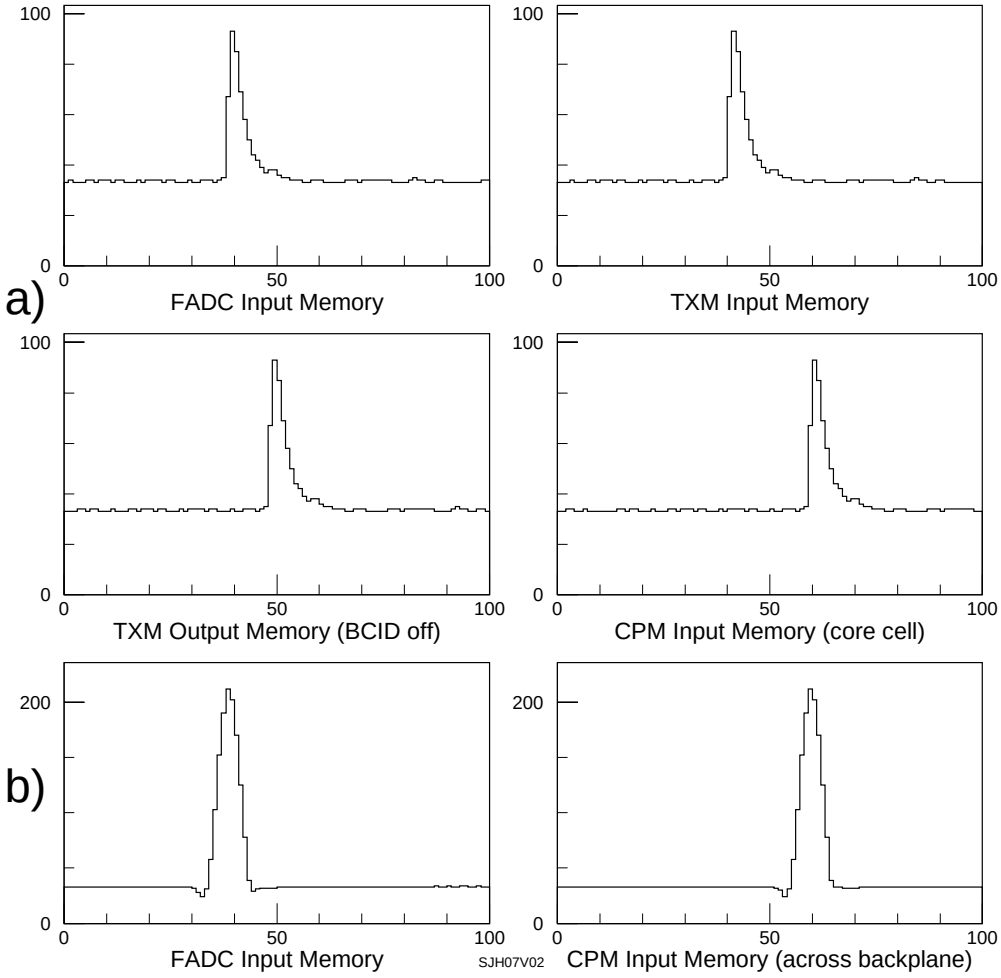

Figure 7-7 Data flow through the complete system. (a) Tile Calorimeter pulse flowing from FADC to TXM input, through the BCID (switched off) to TXM output, and finally to CPM — note the latency at each stage. (b) a test pulse seen across the backplane.

A full system of nine TXMs and nine of the new CPMs was tested at CERN in October 1997. This meant that the full chain of processing could be tested for 36 channels from FADC through to CPM output, and measurements made of errors between each stage. Also the full latency of the demonstrator system could be directly measured. A typical pulse flowing with no errors onto both the on-board CPM memories, and across the backplane, can be seen in Figure 7-8.

All parts of the chain up to the CPM worked perfectly, as far as could be seen from the data recorded. Error rates in the whole of the chain from FADC through BCID in the TXMs and G-link data transmission at 800 MBd worked with bit-error rates of less than 10^{-8}. However, low-level problems were observed with the CPM system. For the ‘core cells’, i.e. those not going across the backplane, bit errors were observed at a rate of about 3×10^{-5} averaged over all channels. Also, occasional latency slips were observed (at a rate of about one in 2,000 events). It was concluded that these errors were due to control problems rather than difficulty with the multiplexing process, and the rate was very board-dependent. There were also problems in timing the backplane data since the ECL parts were slightly slower than expected. This meant that the timing for some data coming across the backplane was difficult or impossible to optimize. This, and other minor problems, caused the bit-error rate on backplane data to be about 5×10^{-5}.

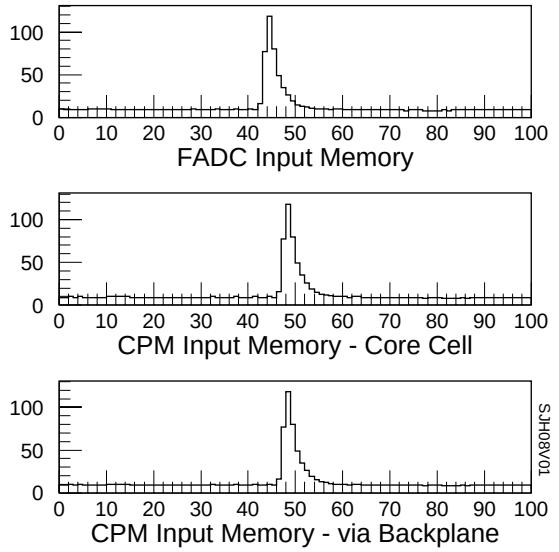

Figure 7-8 Tile Calorimeter data from one FADC flowing perfectly onto two CPM boards, one directly and one via the backplane.

Since there were clearly curable problems in the test beam, it was decided to reconstruct the whole system in the laboratory for final tests. This has the advantage that pseudo-random or particularly difficult data patterns could be used to exercise the system far more stringently than test-beam data. The system was put together in February 1998 after several minor faults had been discovered and fixed in the CPMs. Running the system in the usual way suggested a very high level of data integrity.

Analysing events taken by the DAQ system requires huge numbers of events to prove that the system is working to the level demanded for the final ATLAS system, i.e. bit-error rates less than about 10^{-10}. In order to bypass that problem, an alternative mechanism was devised using a hardware module to perform the necessary comparisons in real time. This custom-built module compared data from the input and output of the system and produced an error count that could simply be scaled. The system is illustrated in Figure 7-9. The results obtained via the DAQ system could be duplicated using this module in a matter of seconds, with no extra analysis required. Data passing through the whole system were measured to have bit error rates less than 10^{-13}, showing that each stage of the processing was performing to this level. Data going across the backplane were measured to have equivalent error rates showing that negligible data corruption and crosstalk was taking place. A full series of tests will be performed at RAL to give confidence in the choices of technology made for the trigger system.

7.6 Full-slice test

An additional important aspect of the demonstrator programme was to show that once all the individual components were working, they could actually all be put together and work in a stable fashion as a genuine trigger. That is to say that using a common clock (equivalent to the final LHC clock), the entire system would work properly to produce a reliable and correct trigger with a constant latency. The test-beam environment, and working with the CTPD, was a vital aspect to showing that the final system acted as a working trigger.

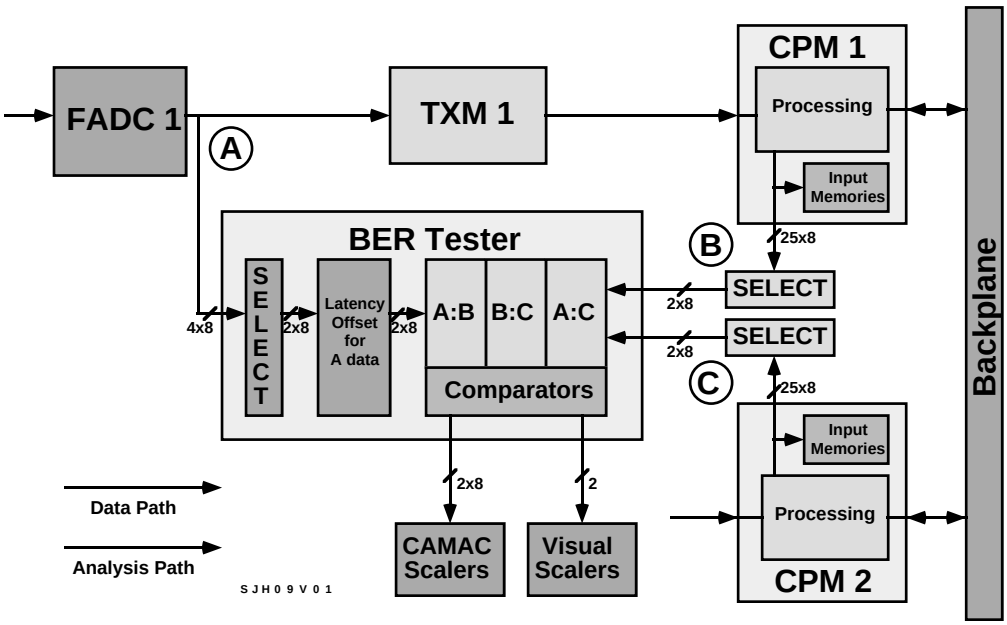

Figure 7-9 Block diagram of the bit-error rate tester and experimental setup.

7.6.1 Timing calibration and stability

The timing control in ATLAS will be provided by the TTC system. This will provide a very stable clock locked to the phase of the LHC bunch-crossings. However, this was not available for use at the test beam, so a module was designed and built at RAL to provide all the necessary timing and synchronization signals. This was the Timing Control Module (TCM) [7-19] used in the phase-2 demonstrator.

The TCM provided fan-out for all the 40 MHz clocks needed by each module in the system, and also generated the stop signals and VME interrupts for event readout purposes. It could be programmed to accept different triggers and also has the flexibility to either use an external clock or generate its own clock. The phase of clocks and stop signals to each of the different types of modules could be programmed, which was necessary to ensure that data were being strobed at the optimum time throughout the system. The success of the TCM is seen by the integrity and latency stability of the data in all parts of the full system.

Along with the programmable delays on the TCM, several other points in the data stream required timing delays. For example, to allow for the possibility of different length cables between crates, the data transmitted from the TXMs could be delayed by more than 1 μs in steps of 1 ns. Delays of this sort are generally quantized in units of 25 ns, corresponding to the time between LHC bunch-crossings. Clearly, the 1 ns steps below 25 ns are needed to align the data to be strobed correctly into the next stage, whereas the ‘tick’ delays (in steps of 25 ns) are needed to equalize the latency of the data.

With the TCM delays and all the individual module delays, the full phase-2 demonstrator had about 100 parameters to calibrate. In principle, many of these delays could be set automatically via software, but in fact most were done by hand as experience was gained with how to control the timing. In designing the system, the policy with delays was to err on the side of caution and provide a delay mechanism in most conceivable places. Much was learnt about which delays were necessary, and how much delays were likely to vary between data streams. Generally, it

was encouraging that once a stable timing regime was found for one data chain, it was usually easy to time the other modules in with the same, or very similar, parameters. This led to a full system with data transfer stable over the full 36 channels for several thousand events. As for the latency stability of individual channels, the bit-error rate measurements given above would have been dominated by latency errors had they been present, so the conclusion is that we could, by use of a few timing parameters, tune the system to be completely stable.

7.6.2 Full-slice tests with FEM and CTPD

The availability of the CTPD module [7-20] during the beam tests allowed the assembly and operation for the first time of a complete slice through the calorimeter trigger. The organization of modules for this is illustrated in Figure 7-10. Data digitized in the 36 FADC modules were passed via TXMs and FEMs to the cluster-processing system, and 16 core-cell energy values were passed to the JPMD. The nine hit outputs from the cluster-processing system and the JPMD were sent to the CTPD, where a trigger decision was taken and a level-1 accept (L1A) signal was formed. Since this was the first SPS run following the SPS fire, beam availability was below average, so the complete slice tests were made using simulated calorimeter pulses derived from a pulse generator. The system was operated in three modes:

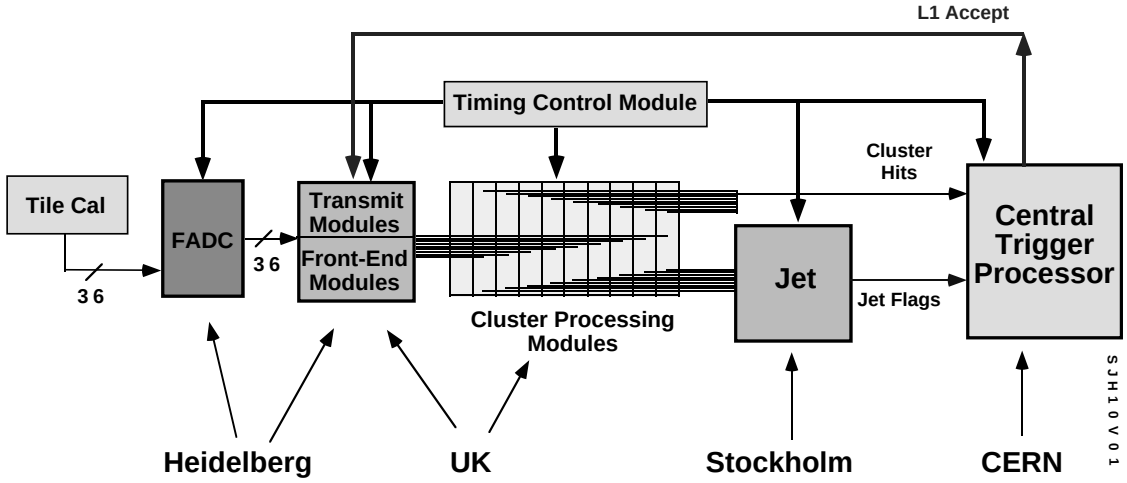

Figure 7-10 Schematic diagram of the complete slice tests using JPMD and CTPD.

- An output from the pulse generator was used as an ‘event’ input to the TCM, stopping the scrolling memories and initiating computer readout of all modules after each analogue pulse was applied to the FADC input. This allowed checks that the data passed correctly through the system, that the correct data were sampled in the FEM pipeline, and that an L1A was generated correctly for each analogue input pulse. The timing relationship between the different parts of the system is shown in Figure 7-11a.
- The system was modified so that the L1A signal provided the ‘event’ signal. In this mode, the system operated as a true trigger for the first time, so that, for example, all readout activity could be suspended by lowering the input analogue pulse height. At this stage, the frequency of L1A was restricted to the maximum rate at which the online computer could record data from all modules. The altered timing relationships are shown in Figure 7-11b.

- The L1A signal was used to initiate the fast pipeline readout in the FEM. Running in this mode, data from the eight time-slices surrounding the L1A were captured in an FEM double buffer without stopping the trigger pipeline. A signal indicating FEM buffer full was used to initiate readout of the FEM only, again without stopping the trigger pipeline. Running in this mode, the frequency of L1A could be raised to many kHz, finally being limited by the ability of the computer to empty the FEM buffer. The contents of the buffer memory are illustrated in Figure 7-12. The analogue input pulse produced a digitized FADC output after pedestal subtraction of around 60 counts. However, since the analogue signal was not phase-locked to the 40 MHz system clock, successive pulses might be digitized at the peak or on rising or falling edges. After BCID, this led to digitizations ranging from 55 to 70 counts. It may be seen that a regular sequence of one in eight time slices is occupied with a value in the expected range, other time-slices having been set to zero by the BCID [7-14].

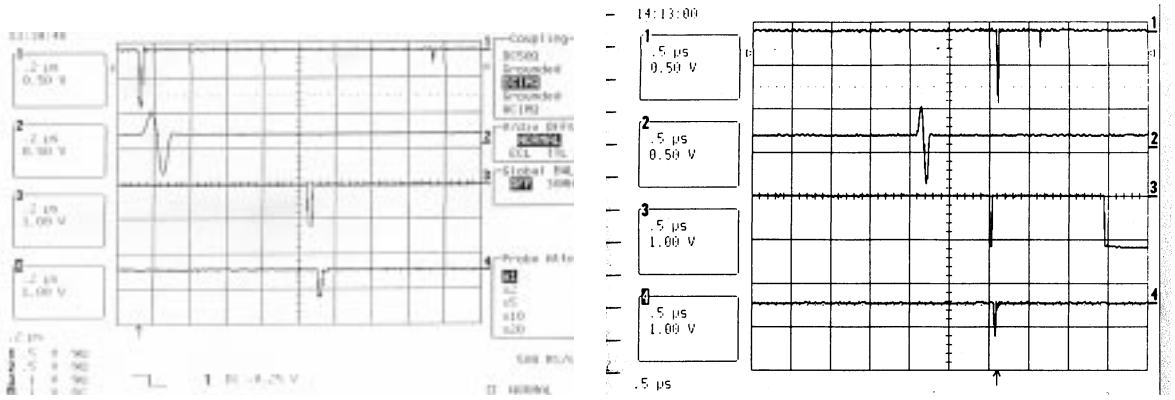

Figure 7-11 Captured oscilloscope traces showing the progress of a trigger pulse through the system. (a) using external trigger, (b) using self-generated trigger. From top to bottom, the traces are: system trigger, generated pulse, CPM hit output, and level-1 accept from CTPD.

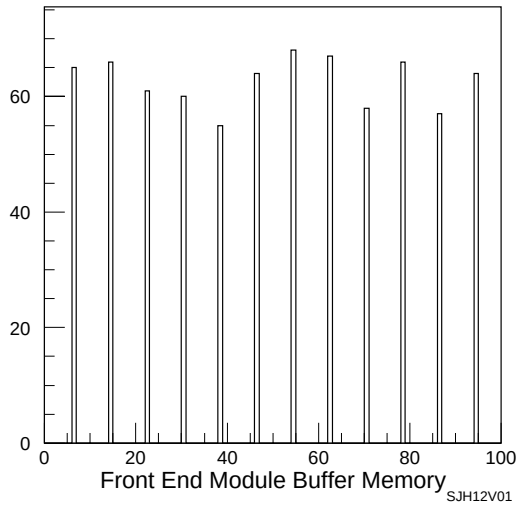

Figure 7-12 BCID processed data captured in FEM memory using level-1 accept to trigger data buffering.

7.7 Conclusions from the test-beam programme

The group has performed a sequence of successful test-beam runs using signals from prototype ATLAS calorimeters. The initially-simple system has evolved into one of considerable complexity which contains most of the essential elements of the proposed final trigger system. During this time many options have been studied and decisions regarding preferred solutions have been made. The final options, as built into the full phase-2 demonstrator, were tested to a high level, and found to perform in a satisfactory manner. There have been some difficulties encountered, as was only to be expected when building such a large integrated system, but these have been overcome in the final system. The experience gained by this programme has led to the current proposal for the level-1 calorimeter trigger, and there can be considerable confidence in the feasibility of each part of this design.

7.8 Computing

The main focus of our demonstrator programme has been to investigate the technologies, such as ASICs, MCMs, fast backplanes and serial links, which are essential components of the trigger. This has also involved developing software, some of which can be considered as a prototyping and testing exercise, particularly in regard to object oriented (OO) techniques. In this section we briefly review the software from our demonstrator programme and summarize what relevant experience has been gained for the design and implementation of software for the final trigger system.

7.8.1 Data acquisition and online monitoring

The DAQ software used in the laboratory and beam tests of the demonstrator system evolved from the CERN Spider system. It has been ported from OS9 to LynxOS and considerably upgraded. In addition to the usual readout and data logging, the DAQ software provides simple histogramming and limited online analysis of the data. While this was essential for our development programme, any further DAQ work should be done in the context of the ATLAS DAQ framework. Discussion will be needed to determine how the ATLAS DAQ should be used for laboratory tests outside CERN on our installed processor base.

Our existing DAQ system provides fast VME readout but has limited computing speed and memory. It allows us to monitor histograms of pulse shapes and the comparisons of pulses from different stages in the processor chain. It also tracks a variety of errors in the readout and the transmission of data across the system. We will still require this kind of low-level monitoring in the final system, but far more sophisticated checking of the performance of the trigger will also be required. Some of the analysis we currently only perform offline should be moved online. Better graphical display of status and errors will also be desirable. The higher level online monitoring will have to be written afresh as a suitable component of the ATLAS DAQ system.

7.8.2 Diagnostics and test package

For the most recent phase of our demonstrator programme, we invested considerable effort in providing a user-friendly diagnostics package to allow physicists and engineers an easy way to manipulate and test the various modules we built. The package presents a graphical view of the

whole configuration, and of the registers and memories on each board and ASIC. The behaviour of modules and their individual registers can be verified in detail. The software has proven to be very useful, and we will certainly require something similar for our future prototyping and testing phase.

The diagnostic software was written mostly in C++, with the front-end user interface written in the Tcl/Tk scripting language which provides a Motif-like 'look and feel'. The software has been described in an ATLAS note [7-21] and more detailed documentation is also available.

In the light of our experience, some redesign of the diagnostics software will probably be desirable. In particular the use of Tcl/Tk to provide the user interface part should be reconsidered as it has become rather cumbersome to maintain the scripts (which are not OO in style). We have also experienced problems (due to limited memory size) in running a large X-based application on 'front-end' processors such as the Motorola MVME167 under LynxOS. Alternative architectures, where the user interface and hardware access are split, will also need to be considered for the future. Thirdly, the existing software provides a detailed model of the behaviour of the hardware. This has been very useful for the debugging of individual cards, but adds a large overhead in memory when a complete configuration consisting of many modules is being used. Some way of dynamically enabling and disabling the detailed model would be useful. Finally, the diagnostic software was written assuming a unified design philosophy for the modules. However, one prototype module contained an internal 'firmware' layer which does not fit naturally into the existing software design.

The diagnostics software was well developed for interactive diagnosis and configuration of individual modules. The architecture also provided for testing complete chains of modules, and there were also facilities for more automated testing of the complete functionality of boards or the links between them. However these aspects have not yet been fully exploited. In the final system, much more emphasis on automated tests will be required.

7.8.3 Calibration

During the demonstrator programme we have developed a number of small stand-alone calibration programs. These have addressed issues such as the tuning of positive ECL (PECL) voltage thresholds and backplane timing between CPMs. However, in practice most of these calibrations have been performed by hand. This proved to be easier in the laboratory during debugging and also in the test beam environment. This will not be feasible in the final system, so much more work on calibration software will be required.

7.8.4 Test-beam data analysis

We have written a number of programs to analyse the data we have recorded in successive beam tests. The latest of these is fully object-oriented and is written in C++. This program unpacks the data read out from each module and calculates the data transformation from input data to expected output performed by each module in a processor chain. It can then compare the actual and expected data at each stage to verify the correct operation of the boards and their various cable and backplane connections. The simulation of each module is simpler than the detailed model implemented in the diagnostic package, but allows for a quick analysis of error sources. This functionality is a large part of what will be required for the online monitoring of the DAQ data from the trigger in ATLAS.

7.9 References

7-1 *A demonstrator programme for the ATLAS level-1 calorimeter trigger*, RD27 note 35, ATLAS note DAQ–NO–031, 1995.

7-2 R. E. Carney et al., *The level-1 bit-parallel electromagnetic trigger processor — CERN beam tests, November 1992*, RD27 note 12, 1993.

7-3 U. Dickow, *An analysis of the RD27 level-1 prototype calorimeter trigger processor test beam data — April/May 1993*, RD27 note 23, 1994.

7-4 I. P. Brawn et al., *Beam tests of a single-channel bunch-crossing identification module for the level-1 trigger — November 1993*, RD27 note 31, 1994.

7-5 I. P. Brawn et al., *Beam tests of the bit-parallel trigger processor with the tile hadronic and liquid argon electro-magnetic calorimeters — September 1994*, Internal note, July 1995.

7-6 I. P. Brawn et al., *Beam tests of a prototype level-1 calorimeter trigger processor for LHC experiments*, Nucl. Instrum. Methods A349 (1994) 356–366.

7-7 I. P. Brawn, *Bunch-crossing identification for the ATLAS level-1 calorimeter trigger*, Ph.D. Thesis, University of Birmingham, 1996.

7-8 S. J. Hillier, *Summary of September 1996 test-beam*, Internal note, October 1996.

7-9 V. Perera, *Data on the LHC electromagnetic cluster-finding ASIC RAL114*, RAL hardware write-up, February 1992.

7-10 P. Hanke et al., *A flash analog-to-digital converter for trigger tests with the ATLAS calorimetry*, RD27 note 41, 1995.

7-11 A. Mass et al., *Design and test of an ASIC for bunch-crossing identification in the ATLAS level-1 Trigger.* CERN / LHCC / 95–36, First Workshop on Electronics for LHC Experiments, Lisbon, 1995.

7-12 C. Schumacher, Diplomarbeit Universität Heidelberg: *Der ATLAS Level-1 Trigger, Auslese des Front-End*, IHEP 97–05, 1997.
http://www.ihep.uni-heidelberg.de/Veroeffentlichungen/ps/schumacher.diplom.ps.gz

7-13 A. Mass et al., *A front end digitization and readout system for the ATLAS level-1 calorimeter trigger.* CERN / LHCC / 96–39, Second Workshop on Electronics for LHC Experiments, Balatonfüred, Hungary, September 1997.
http://wwwasic.ihep.uni-heidelberg.de/atlas/PAPERS/lerb-96.ps.gz

7-14 U. Pfeiffer et al., *Performance of the front-end demonstrator system of the ATLAS level-1 calorimeter trigger*, CERN / LHCC / 97–60, Third Workshop on Electronics for LHC Experiments, London, UK, September, 1997.
http://wwwasic.ihep.uni-heidelberg.de/atlas/PAPERS/LEB97.ps.gz

7-15 U. Pfeiffer, Diplomarbeit Universität Heidelberg: *Analoge optische Signalübertragung vom Kalorimeter zum ATLAS Level-1-Trigger.* IHEP 96–16, 1996.
http://www.ihep.uni-heidelberg.de/Veroeffentlichungen/ps/pfeiffer.diplom.ps.gz

7-16 G. C. Lewis, *ATLAS transmitter module demonstrator phase-2*, RAL hardware write-up, January 1996.

7-17 J. Edwards and V. Perera, *User guide for the first level calorimeter trigger processor demonstrator dual function ASIC RAL163*, Hardware write-up, 1994.

7-18 R. Staley, *Cluster processing module — user guide*, Hardware write-up, University of Birmingham, 1996.

7-19 B. Hatley and J. Leake, *Programming model for the trigger timing control module used in the phase-2 ATLAS bit-parallel calorimeter trigger demonstrator*, RAL hardware write-up, 1996.

7-20 G. Schuler, *Central Trigger Processor*, RD27 working document, August 1996.

7-21 J. Leake et al., *ATLAS level-1 calorimeter trigger diagnostic software*. ATLAS note DAQ–NO–72, August 1997.

8 Calorimeter trigger implementation

8.1 Introduction

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

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

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

8.2 Interfaces

8.2.1 Central Trigger Processor

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

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

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

8.2.2 Level-2 trigger

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

8.2.3 Data acquisition system

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

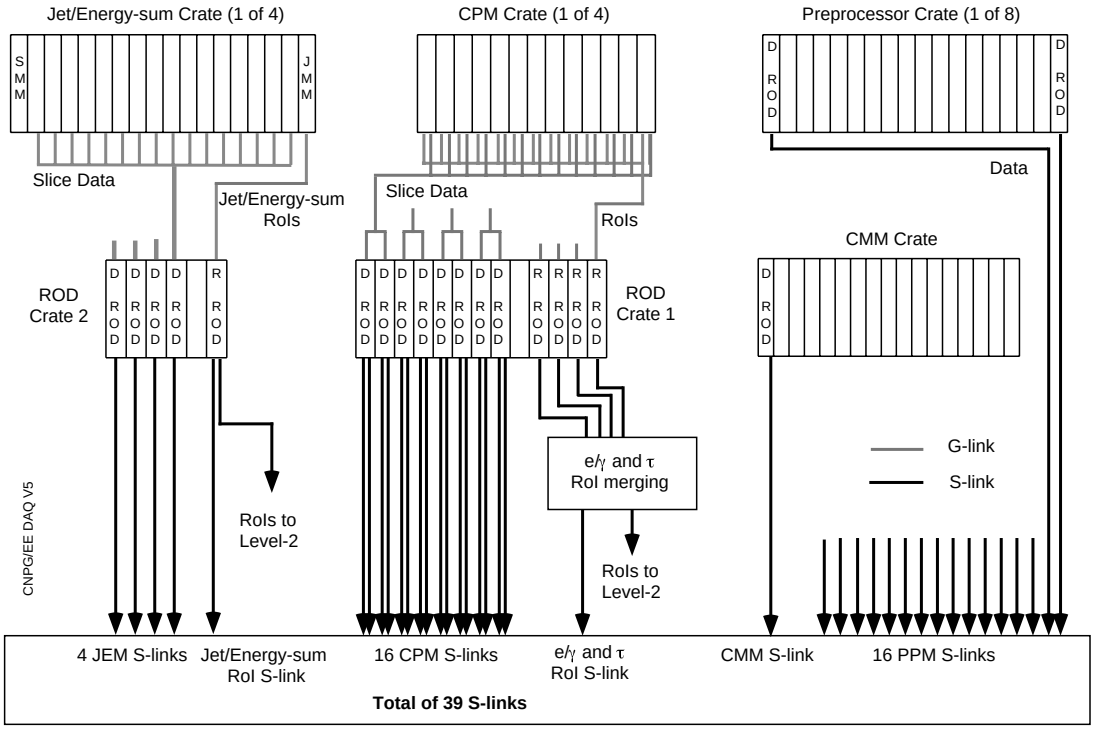

Figure 8-1 Diagram of RODs for data acquisition and RoIs.

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

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

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

8.2.4 Detector Control System

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

8.3 Latency estimate

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

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

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

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

Table 8-3 Latency in analogue signal-handling.

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

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

Table 8-4 Latency in the Preprocessor.

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

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

Table 8-5 Latency in the Cluster Processor.

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

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

Table 8-6 Latency in the jet trigger.

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

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

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

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

8.4 Technologies

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

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

8.4.1 Serial links

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

Gigabit link operation requires consideration of many technical issues:

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

8.4.1.1 Hewlett–Packard G-links

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

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

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

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

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

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

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

8.4.1.2 Alternative serial links

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

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

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

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

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

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

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

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

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

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

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

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

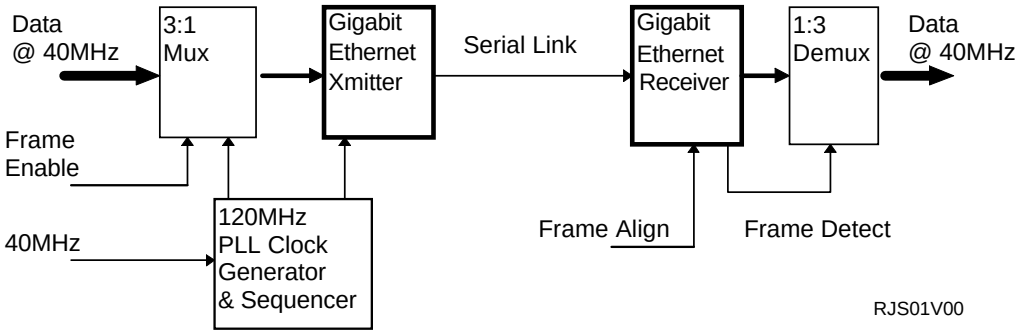

Figure 8-2 Scheme for using Gigabit Ethernet parts to transmit 40 MHz parallel data serially.

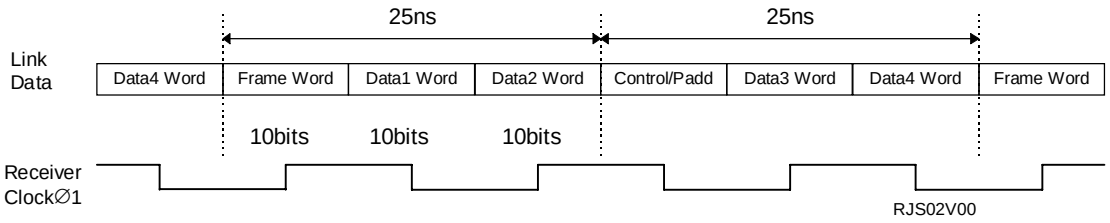

Figure 8-3 Proposed serial link format using Gigabit Ethernet parts.

Custom design

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

8.4.2 ASICs

8.4.2.1 Use of ASICs in the trigger system

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

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

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

The Cluster Processor will use two ASIC designs:

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

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

8.4.2.2 ASIC design experience

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

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

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

FeASIC
This is a prototype of the PPrASIC. It includes the calibration lookup table and the BCID function in the real-time data path, and components of the readout function. The implementation required approximately 80,000 transistors (~20,000 gates) on a 16.4 mm^2 die.

The two preceding ASICs are described further in Section 7.4.4

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

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

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

RAL114 (0.8 micron CMOS technology from Fujitsu)
This was a 1993 gate-array design, using 20,000 gates, to demonstrate the implementation of the cluster-finding algorithm using pipeline processing elements of adders and comparators.

RAL163 (0.7 micron CMOS technology from ES2)
This ASIC is a semi-custom cell-based design serving two functions:

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

RAL215 (0.7 micron CMOS technology from ES2)
The Cluster Processor will contain 6400 trigger towers which will be serialized at 160 Mbit/s on the CPMs. As these serial bit-streams have no clock recovery mechanism, an automated calibration scheme is required to align the 160 MHz clock to the incoming data and synchronize them to the 40 MHz system clock. To evaluate one possible method, the RAL215 ASIC employs phase-locked delay techniques to generate controlled delay-elements and to use them to capture correctly the serial data and synchronize them to the 40 MHz system clock. In addition, the design includes elements for testing fast mulitpliers for BCID FIR-filter coefficients. See Section 6.3.7.3.

8.4.3 FPGAs

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

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

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

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

8.4.4 Multi-chip modules

8.4.4.1 Use of MCMs in the trigger system

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

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

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

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

8.4.4.2 MCM design experience

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

Cluster Processor demonstrator MCM (RAL2403)
The demonstrator MCM incorporates two HP G-link dies (HDMP-1014D) to receive the trigger towers at 800 MBd and convert them to 16-bit parallel words, and two multiplexing ASICs (RAL163) which convert each 16-bit parallel word into four 160 Mbit / s bit-streams for the cluster-finding ASICs.

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

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

Preprocessor demonstrator MCM

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

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

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

8.4.5 Backplanes

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

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

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

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

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

The design features of this prototype were:

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

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

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

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

Table 8-9 Comparison of backplane features.

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

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

8.4.6 Connectors and cables

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

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

8.4.6.1 Cable connectors

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

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

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

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

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

8.4.6.2 Backplane connectors

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

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

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

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

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

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

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

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

8.4.7 Programmable delay elements

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

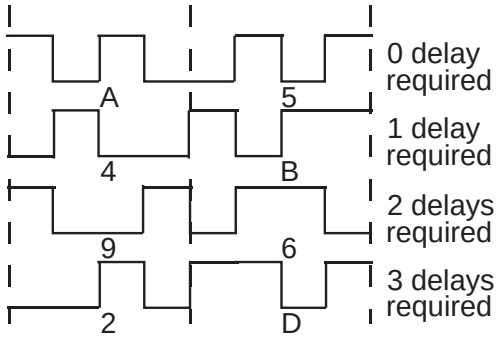

Figure 8-4 Calibration sequence pattern.

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

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

8.5 Rack and crate layout of the calorimeter trigger

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

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

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

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

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

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

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

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

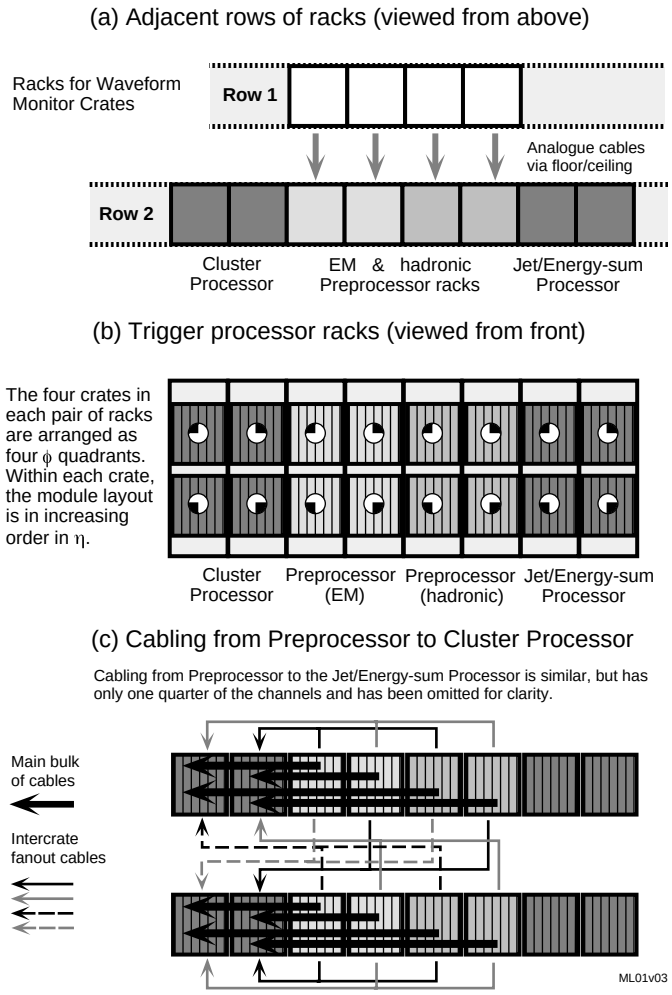

Figure 8-5 Possible layout of calorimeter trigger racks.

8.6 Software

8.6.1 Future development of test and diagnostic software

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

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

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

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

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

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

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

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

8.6.2 Future development of calibration software

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

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

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

8.6.3 Online control and monitoring of the calorimeter trigger

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

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

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

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

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

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

8.7 Strategy for setting up the timing

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

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

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

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

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

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

8.7.1 Synchronization of input signals

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

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

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

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

8.7.2 Use of real particles

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

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

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

8.7.3 Timing of the readout

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

8.8 Design, construction and assembly procedure

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

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

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

8.8.1 Design procedure

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

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

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

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

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

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

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

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

8.8.2 Assembly and pre-testing

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

8.8.3 Installation and integration

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

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

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

8.8.4 Quality assurance and review procedures

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

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

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

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

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

8.9 Timetable

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

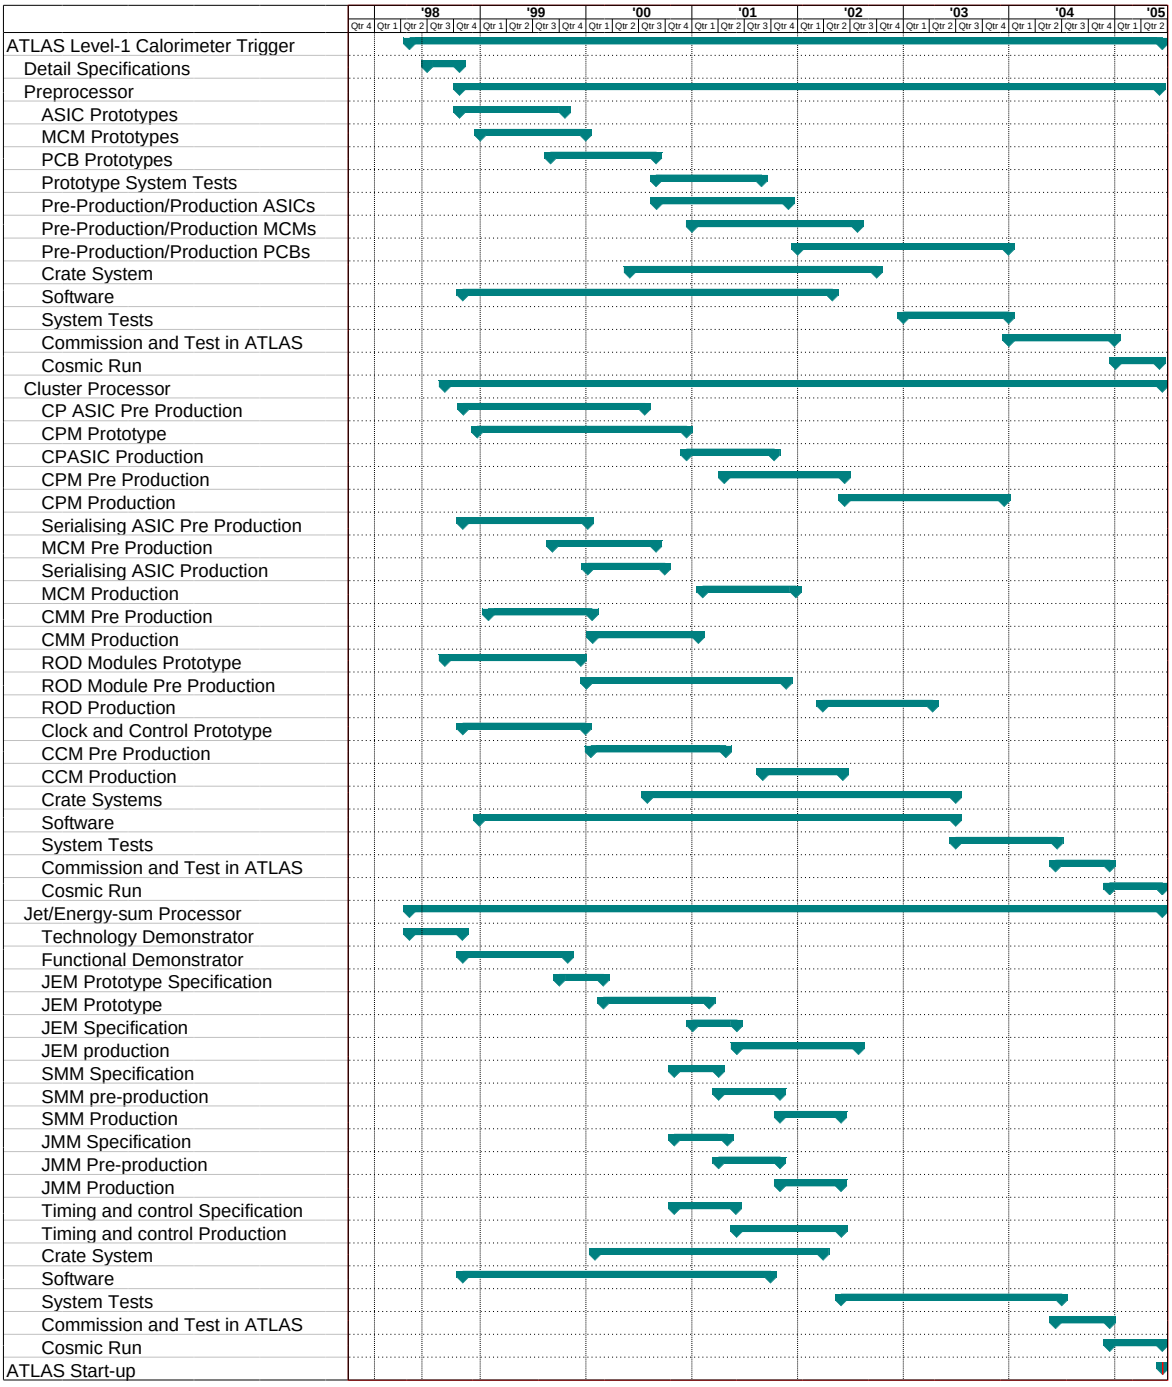

Figure 8-6 Schedule for design and construction of Level-1 Calorimeter Trigger.

8.10 References

8-1 A.J. Maddox, *Simple-links.* RAL hardware write-up, January 1998.
http://hepwww.ph.qmw.ac.uk/~efe/Simple-Links.html

8-2 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, 1994

8-3 A.J. Maddox, *Evaluation summary of the Harting 'harpak' mini-coaxial module with IS23 approved cable.* RAL hardware write-up, February 1998.
http://hepwww.ph.qmw.ac.uk/~efe/coax&harpak.html

8-4 *Back-End DAQ User Requirements Document*. Draft Revision 3, 1996.
http://atddoc.cern.ch/Atlas/DaqSoft/document/draft.ps

8-5 J. Edwards and V. Perera, *User guide for the first level calorimeter trigger processor demonstrator dual function ASIC – RAL163*, RAL hardware write-up, 1994.

8-6 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN / LHCC / 96–41, 1996.

8-7 *ATLAS Tile Calorimeter Technical Design Report*, CERN / LHCC / 96–42, 1996.

9 Muon trigger overview and requirements

9.1 Muon trigger overview

The level-1 muon trigger is based on dedicated, fast and finely segmented muon detectors. The layout of these so-called trigger chambers, Resistive Plate Chamber (RPC) detectors in the barrel and Thin Gap Chamber (TGC) detectors in the end-caps, has been documented elsewhere [9-1]. Very briefly, the RPC-based system covers the pseudorapidity range $|\eta| < 1.05$, while the TGC-based system covers $1.05 < |\eta| < 2.4$. More details on the layout, and in particular on changes compared to the layout in [9-1], are given in Chapter 10.

The RPCs are wireless strip detectors with time resolution $\sigma_t = 1.5$ ns. The TGCs are multi-wire detectors with both wire and induced strip read-out having a finer segmentation albeit a larger timing-resolution. As discussed in following chapters, the timing resolution of both kinds of detectors is sufficient to provide unambiguous identification of the bunch crossing containing a high-p_T muon candidate.

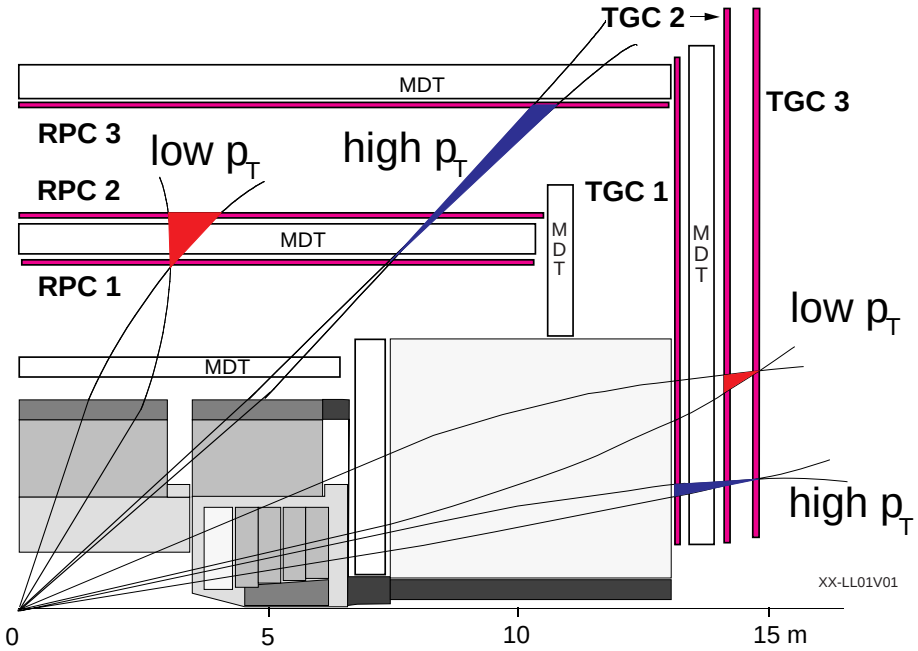

Figure 9-1 Longitudinal view of the end-cap and barrel muon trigger systems.

As illustrated in Figure 9-1, the LVL1 trigger is based on three trigger stations. Two stations are used for low-p_T muon triggers (threshold range approximately 6–10 GeV), while the third station is used in addition for high-p_T triggers (threshold range approximately 8–35 GeV). Each station is composed of two detector planes (with the exception of the innermost TGC station that has three planes). Each detector plane is read out in two orthogonal projections, η and ϕ, that will be referred to as the bending and non-bending projections, even though in the 'non-bending projection' there is some bending in the barrel and substantial bending in the end-cap. (Note that in the innermost TGC triplet, only two of the three planes are read out in the non-bending projection.)

The sharpness of the p_T cut applied by the trigger is mainly given by the information read out from the detectors in the bending projection. However, the information in the non-bending view helps to reduce the background trigger rate from noise hits in the chambers produced by low-energy photons, neutrons and charged particles, as well as localizing the track candidates in space as required for the LVL2 trigger. In addition, the trigger chamber information in the non-bending view provides the second coordinate measurement for offline reconstruction of muons (the precision chambers give information only in the bending projection).

As indicated in Figure 9-1, the basic principle of the algorithm is to require a coincidence of hits in the different chamber layers within a road. The width of the road is related to the p_T threshold to be applied. Space coincidences are required in both views, with a time gate close to the bunch-crossing period (25ns). The coincidence requirement allows for missing layers due to detector inefficiencies, dead regions, etc. For the low-p_T trigger, hits are required within the road in at least three of the four layers, in each of the two projections. For the high-p_T trigger, an additional requirement is made, demanding hits in at least one of the two layers in each of the two projections of the third station. (In the case of the third station for the TGCs, there are three active detector layers in the bending plane, and a two-out-of-three requirement is made.)

As discussed below, a system of programmable coincidence logic allows concurrent operation with a total of six thresholds, three associated with the low-p_T trigger and three associated with the high-p_T trigger. Each of the six thresholds is independently programmable.

9.2 Requirements analysis

A detailed study has been made of the requirements for the LVL1 muon trigger (L1MT). In this section we give a summary of the main requirements; details can be found in [9-2]. Here we discuss what the L1MT has to do, without discussing how it should do it. Subsequent chapters address algorithm and implementation issues.

A context diagram for the L1MT is shown in Figure 9-2, showing its connections to external systems.

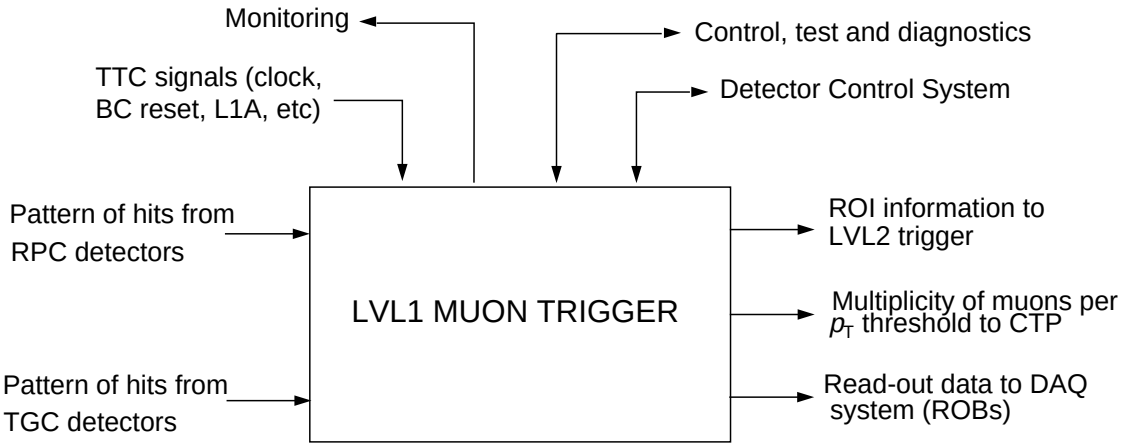

Figure 9-2 Context diagram for the level-1 muon trigger.

The L1MT has to process hit information from the so-called Trigger chambers, RPCs in the barrel region and TGCs in the end-caps [9-1], identifying candidate muon tracks. The hit signals are provided by Amplifier-Shaper-Discriminator (ASD) circuits that are part of the RPC and TGC front-end electronic systems. The muon track candidates must be assigned to a unique bunch crossing. The trigger must provide six independently-programmable p_T thresholds; a p_T range has to be indicated for each track candidate.

The L1MT has to provide data to a number of systems: the LVL1 central trigger processor (CTP), the LVL2 trigger and the DAQ. The data to be sent to the CTP are the muon-candidate multiplicity values for each of the six p_T thresholds. Data to be sent to the LVL2 trigger and the DAQ include the multiplicity values and information about each of the candidate tracks (p_T range; position in η, ϕ). The pattern of hit strips and wire-groups also has to be sent to the DAQ. The data sent to the CTP, which are used to make the overall LVL1 trigger decision, must be provided with the smallest possible latency. A specific requirement is that the information from the L1MT must arrive at the CTP within 51 bunch-crossings following the collision that produces the muon(s); this figure includes cable propagation delays.

9.2.1 Requirements on trigger performance

There are a number of requirements on the performance of the L1MT which are summarized in this subsection. As discussed later in Chapter 14, simulations have been performed to demonstrate that the proposed trigger design meets the requirements.

The trigger has to be able to operate with p_T thresholds in the range 6–35 GeV. The threshold value is defined such that muons with transverse momentum greater than the quoted value shall be triggered with > 90% efficiency if they fall within the acceptance of the trigger.

An average acceptance of at least 90% is required for muons in the pseudorapidity range $|\eta| < 2.4$. Note that the p_T-dependent trigger efficiency and the acceptance have to be combined when evaluating the overall probability that a muon a of given p_T will be identified by the trigger.

The trigger has to make a reasonably sharp cut on the p_T of muon candidates. It is required that, for a variety of threshold settings, there should be at least a specified fraction of muons with true p_T above the threshold p_T value. For example, for a trigger threshold setting of 20 GeV, at least 25% of the muons selected by the trigger should have true p_T greater than 20 GeV. These requirements take into account real muons from a number of sources: decays of bottom and charm quarks, decays of W and Z particles and decays in flight of charged pions and kaons.

The step size by which the trigger thresholds can be adjusted has to be reasonably small. Quantitatively, it is required that the increase in rate for the inclusive muon trigger, due to decreasing the threshold by one step, shall not be more than a factor of 1.5 within the p_T range 6–20 GeV (factor 2 within the p_T range 20–35 GeV).

The rate of triggers due to background sources has to be kept as small as possible. It is required that the rate of fake triggers, due to correlated and / or uncorrelated noise hits in the trigger chambers, shall be much less than the rate from real muons. It is also required that the trigger rates due to cosmic rays, beam-halo particles and due to hadronic shower leakage into the trigger chambers shall be small.

The L1MT has to provide to the CTP the multiplicity of muon candidates for each of the six p_T thresholds. This is important so that, for example, a low-p_T dimuon trigger can be maintained at high luminosity. It is foreseen that the threshold on the dimuon trigger will be kept at about 6GeV per muon for $L = 10^{34}\,cm^{-2}\,s^{-1}$, while the threshold for the single-muon trigger will have to be about 20GeV for an acceptable trigger rate. Given the steeply falling muon p_T distribution, it is important that muons should only very rarely be double counted (e.g. in areas of overlapping chambers), giving fake dimuon triggers from single-muon events. It is required that at most 10% of the dimuon triggers shall be due to doubly-counted single muons.

The trigger system may not be able to resolve pairs of muons that are not well separated in the detector. It is required that the fraction of events that are lost will be small for a number of representative physics processes.

It should be possible to use the L1MT to provide a cosmic-ray and beam-halo trigger for the experiment, for use in calibration and alignment studies. In the cosmic-ray mode, the system would be driven by a free-running 40MHz clock.

9.2.2 Other requirements

The detailed requirements for the data to be sent, for LVL1-accepted events, to the LVL2 trigger and to the DAQ can be found in Ref. [9-2]. Very briefly, the required data (pattern of hit strips and wire groups, information on each muon candidate and the multiplicities for the various thresholds) have to be retained in pipeline memories during the latency of the LVL1 trigger. When an event is accepted by the LVL1 trigger, the corresponding data have to be sent to the DAQ system using the ATLAS standard read-out link / read-out buffer chain. A subset of the data, including the information on each of the muon candidates, has to be sent via a separate path to the region-of-interest builder of the LVL2 trigger. The data that are sent to the DAQ will be used for monitoring the performance of the trigger, and also for calculation of trigger efficiencies and acceptances. The data that are sent to the LVL2 trigger will be used to guide the LVL2 processing by defining roads in which to search for muon-track candidates.

The L1MT is a data source that, from the point of view of the DAQ, resembles a front-end system. It must therefore conform to the ATLAS requirements for the interface between front-end systems and the DAQ [9-3]. Similarly, it must take into account the needs of the LVL2 trigger [9-4].

There are numerous other requirements concerning the ability to control, test, and monitor the L1MT system, and also on error handling and fault tolerance, and maintenance. These are documented in [9-2].

An important issue for all parts of the LVL1 trigger (and also for detector front-end systems) is the need for a strategy for setting up the timing for various modes of operation: beam–beam collisions, cosmic-ray and beam-halo triggers, test and calibration running. The system contains numerous programmable delay and phase-adjustment circuits, and a means has to be worked out for determining the appropriate values of all the parameters of these components. The procedure for setting up the timing should not rely on physical access to the electronic systems, some of which are in any case on the detector and therefore inaccessible when there is beam in the machine.

9.3 System implementation overview

The electronics implementation of the level-1 muon trigger is shown in Figure 9-3 below. The detector and front-end electronics are described in the Muon Spectrometer TDR [9-1]. Summary and update information can be found in Chapter 10. Despite the problems of radiation, magnetic fields, power distribution and cooling, a large fraction of the system is mounted on the detectors in order to enable the accurate timing of detector signals into the coincidence logic and to reduce cabling – the on-detector coincidence logic reduces about 800,000 channels to about 1000 optical links to the off-detector trigger electronics located outside the ATLAS cavern in underground area USA15. In both the barrel and end-cap systems track candidates found independently in η an ϕ are subsequently combined. The η-ϕ area is divided into sectors, 64 for the barrel and 144 for the end-cap. Each sector provides track candidates to the Muon to Central Trigger Interface (MUCTPI). In addition to providing trigger signals, the system reads out all hits in the muon trigger chambers for all ATLAS level-1 triggers. The MUCTPI counts muon candidates for six thresholds, removing doubly counted tracks in overlap regions, and generates the Region-of-Interest (RoI) list for the level-2 trigger system. A RoI is generated for each candidate.

Differences between the barrel and end-cap region requirements result in different capabilities and different implementations. These are summarized in Section 2.4.2.

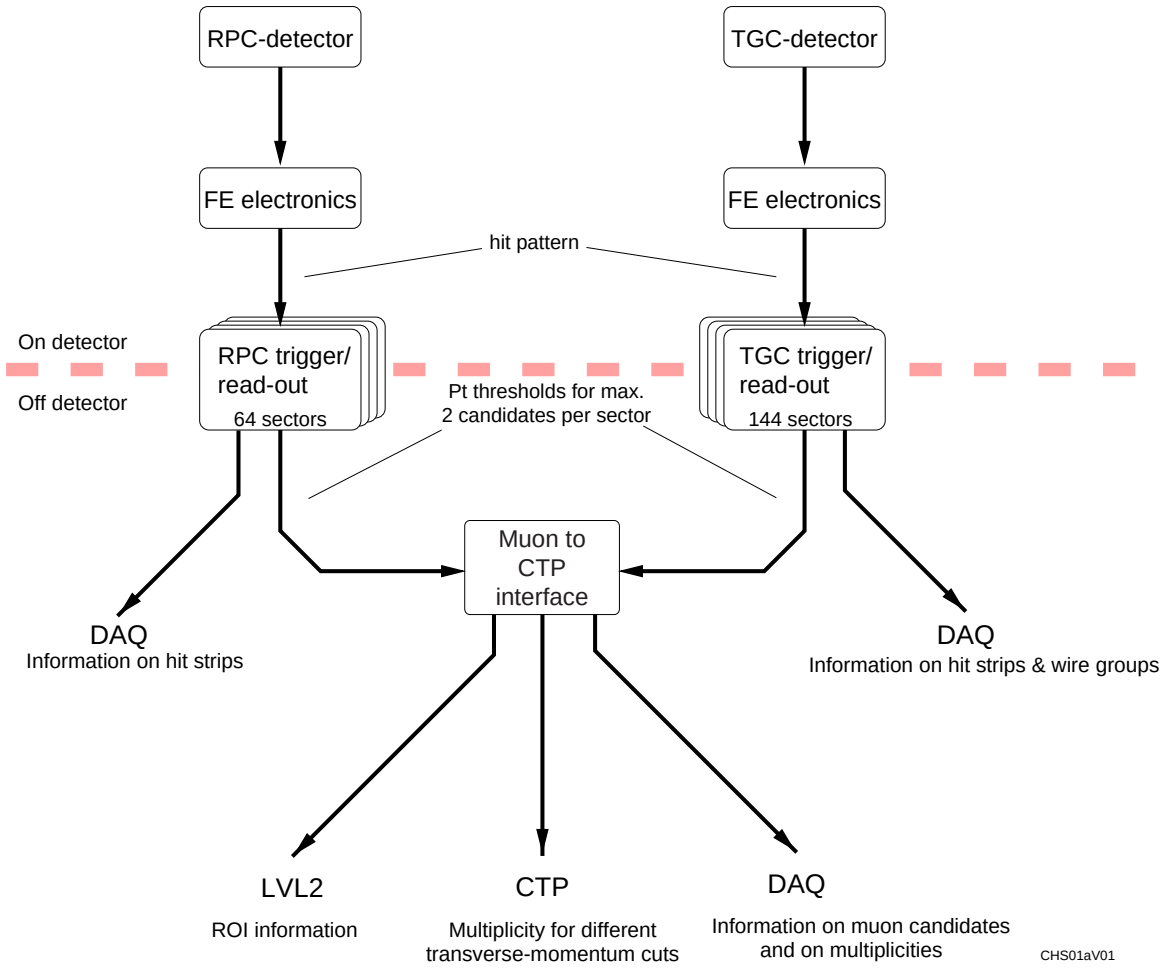

Figure 9-3 Block diagram of the level-1 muon trigger system.

9.4 References

9-1 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

9-2 *LVL1 Muon Trigger User Requirements Document (Draft Version 1.4)*, ATLAS working document, ATL–DA–ES–0002, March 1998.
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/muons/L1MT980309.ps

9-3 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

9-4 *ATLAS Level-2 Trigger User Requirements Document*, ATLAS note DAQ–NO–079, November 1997.

10 Muon trigger chambers

10.1 Introduction

The detectors that provide signals to the level-1 muon trigger electronics must provide fast information on muon tracks traversing the detector in an environment of very high background. The trigger detectors must provide information that allows:

- discrimination on muon transverse momentum, p_T
- bunch-crossing identification
- fast, coarse tracking for input to higher-level triggers
- second coordinate measurement in the non-bending projection with a resolution ~10 mm

These demands define the requirements on both the time resolution (less than the bunch-crossing interval) and the spatial resolution of the detectors. Stand-alone trigger detectors are necessary since the precision muon chambers have maximum drift times much longer than the LHC bunch-crossing period of 25 ns and are therefore unsuitable for bunch-crossing identification. The trigger detectors must provide acceptance in pseudo-rapidity out to $|\eta| \approx 2.4$, and over the full ϕ range. The soft background in the cavern will induce random hit rates in the trigger detectors in the range $20-60\,\mathrm{Hz/cm^2}$ at high luminosity. This requires using a coincidence of two or more planes for at least one of the coordinates of a triggering track.

The technologies used for the trigger detectors must satisfy the above criteria and yet be suitable for mass-production in order to cover the large areas required. The solutions chosen optimize the cost of the total system, matching the rate capability and the spatial resolution to the above requirements. Since the operating conditions in barrel ($|\eta| < 1.05$) and end-cap ($1.05 < |\eta| < 2.4$) regions are different, different technologies are chosen for these two regions.

In the barrel, the Resistive Plate Chambers (RPCs) provide the needed time and space resolution. RPCs do not have wires and are therefore easy to construct. This allows covering large areas simply and inexpensively. They have an adequate rate capability and a very good time resolution, $\sigma = 1.5$ ns, that allows bunch-crossing identification.

In the end-cap region, where the trigger chambers are located outside the toroidal field with a smaller separation between trigger stations, a higher granularity is needed. This need is enhanced by the higher particle momenta for the same p_T and the lower $\int B\mathrm{d}l$ in certain η–ϕ regions. Thin Gap Chambers (TGCs) provide this capability. Higher granularity is easily attained by the use of wires for read-out. Despite the need for high granularity, the number of electronic channels can be kept reasonably low by shifting the read-out channels by a half or a third of a channel width between layers. This ‘staggering’ of wire-groups is possible because TGCs have small group-to-group cross-talk. A smaller effective granularity is therefore achieved with a reasonable number of channels. Thin Gap Chambers also have the good time resolution required for bunch-crossing identification and a high rate capability to cope with the expected higher backgrounds in the forward region. TGCs are extremely robust and have been proven in a previous long-running experiment [10-1].

The trigger scheme used is similar in both barrel and end-cap regions, as described in Chapter 9 and in more detail in Chapters 11 and 12. For more details on the Muon Spectrometer including the RPC and TGC chambers and front-end electronics, see the Muon Spectrometer TDR[10-2].

10.2 Layout

10.2.1 Barrel

The RPCs are used in the barrel region up to $|\eta|=1.05$. The trigger detector is organized in three stations, two of which are located above and below the MDT chambers in the middle station, and the third in the outer station as shown in Figure 10-1.

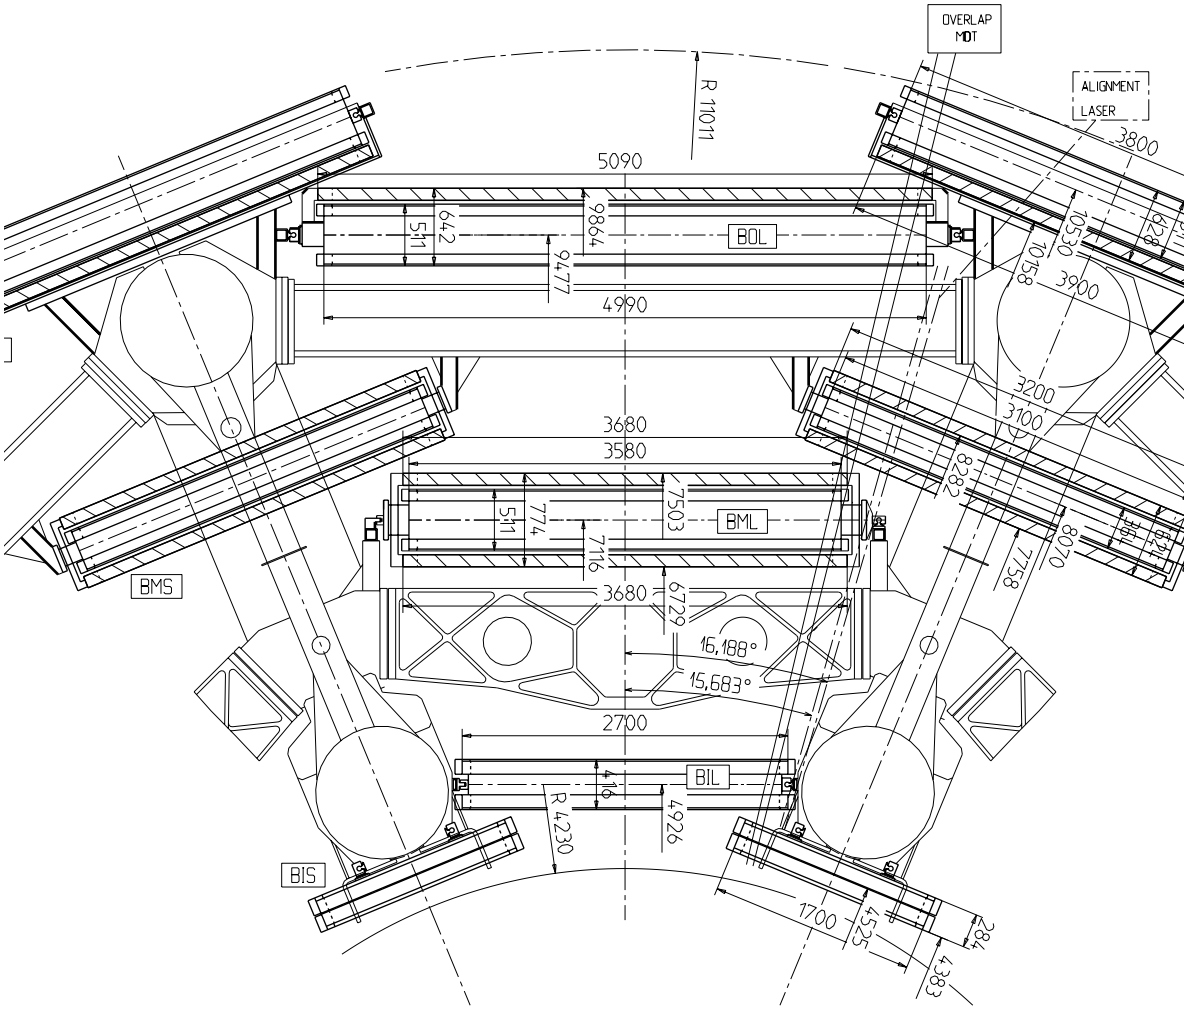

Figure 10-1 View of a standard barrel sector. The RPCs are represented by shaded boxes.

Each RPC station consists of two gas gaps with two planes of read-out strips each: one in the transverse and one in the longitudinal direction. Both planes are used in the trigger: the so called ‘η strips’ are parallel to the MDT wires and provide the bending view of the trigger detector; the ‘ϕ strips’ are orthogonal to the MDT wires and provide the second-coordinate measurement. The ϕ-strips are also needed for the pattern recognition. The RPCs are organized in several modules and their dimensions having been chosen to match those of the corresponding MDTs. In most of the stations, RPCs are composed of two units along both the azimuthal and the beam direction. To avoid dead areas between adjacent units, the active zones of neighbouring RPCs are partially overlapped in η as shown schematically in Figure 10-2. There is a total of 1052 chambers with close to 430,000 read-out channels. Table 10-1 summarizes the RPC strip dimensions and Figure 10-3 details the placement of the front-end boards.

Table 10-1 RPC dimensions

number of chambers	1052
number of channels	430,000
η strip pitch	26.2–30.0 mm
ϕ strip pitch	23.1–26.8 mm
η strip length	480–1200 mm
ϕ strip length	970–2425 mm

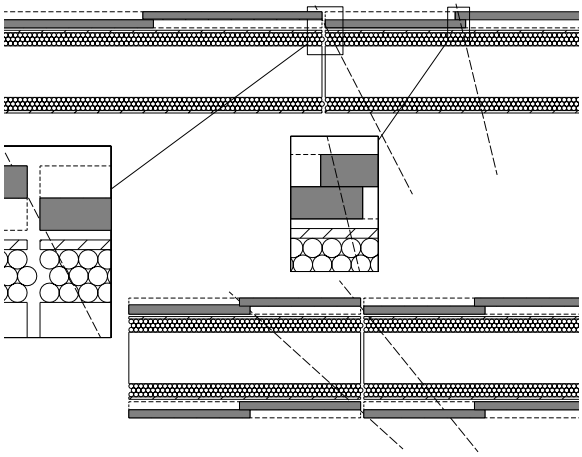

Figure 10-2 Schematic view of the border in the η direction between two neighbouring chambers for the middle and outer stations.

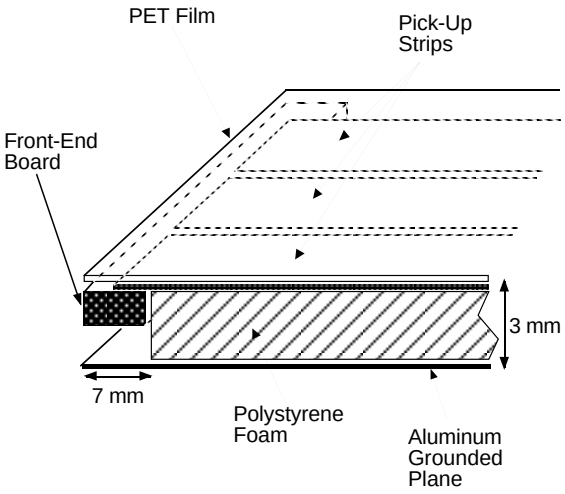

Figure 10-3 Detail of RPC strip and front-end board placement.

10.2.2 End-cap

As shown in Figure 10-4, TGC trigger chambers are arranged in the ATLAS detector in seven layers in each end-cap at $|z| \sim 14$ m. They are grouped in three planes in z: one plane of triplet units and two of doublet units. The doublet forming the plane farthest from the interaction point in each end-cap is referred to as the ‘pivot plane’, and the chamber layout and electronics are arranged such that there are no overlaps or holes in this plane. The second doublet plane and the triplet plane are separated from the pivot plane by 50 cm and 160 cm in z respectively. Neighbouring units within a plane are staggered in z to produce overlaps that prevent inactive regions creating acceptance gaps. For triggering, the TGCs cover a pseudorapidity range $1.05 < |\eta| < 2.4$. Each trigger plane of TGCs consists of a ‘wheel’ of eight octants of chambers symmetric in ϕ. As shown Figure 10-5, each octant contains three sets of units (doublets or

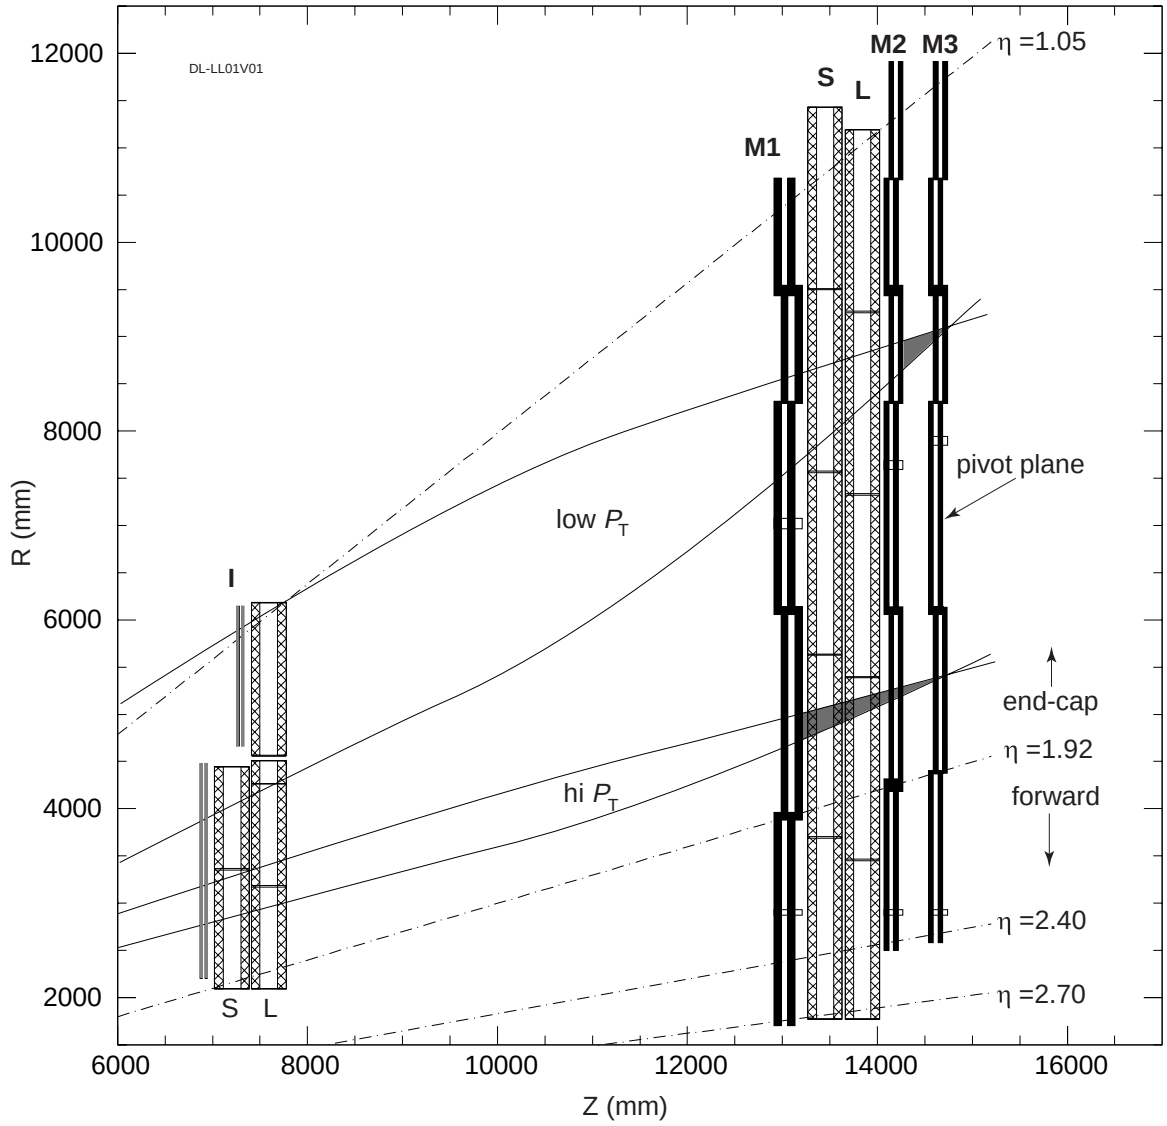

Figure 10-4 Longitudinal view of the TGC system, showing trajectories of high and low p_T muons and their trigger windows. M1 is the TGC triplet, M2, M3 are the TGC doublets, S and L are the small and large MDTs, I is the inner-station TGC (not used in the trigger). To allow overlapping of the physical chambers, doublets and triplets at adjacent ϕ are at slightly different z; chambers at two adjacent ϕ's are shown.

triplets), where a set consists of one unit in the Forward region at high η, and eight or 10 units in the End-cap region. Coincidences between hits in the pivot plane and the adjacent doublet plane are used to make a low-p_T threshold trigger, and those between hits in all three planes and to make a high-p_T threshold trigger. The layout in ϕ is such that the two end-caps are mirror images of each other so that the detector has CP symmetry.

The TGCs provide a measurement of the ϕ coordinate in the inner layer, I, of the muon spectrometer, where soft backgrounds induce a hit rate exceeding 200 Hz / cm^2.

10.2.2.1 Layout changes since the Muon TDR

Some improvements have been made to the layout of the end-cap muon system since the Muon Technical Design Report [10-2]. Since each TGC station is now mounted on a single, large circular wheel, there is no longer any shift in *z* between the Forward and End-cap TGC sections. The consequent reduction in the number of chambers and their increased size has reduced the total area of their overlap regions. A further benefit is the reduction in the number of channels without loss of resolution. The ϕ segmentation is now rationalized such that a span in ϕ of two End-cap chambers is now matched by one Forward chamber, as shown in Figure 10-5.

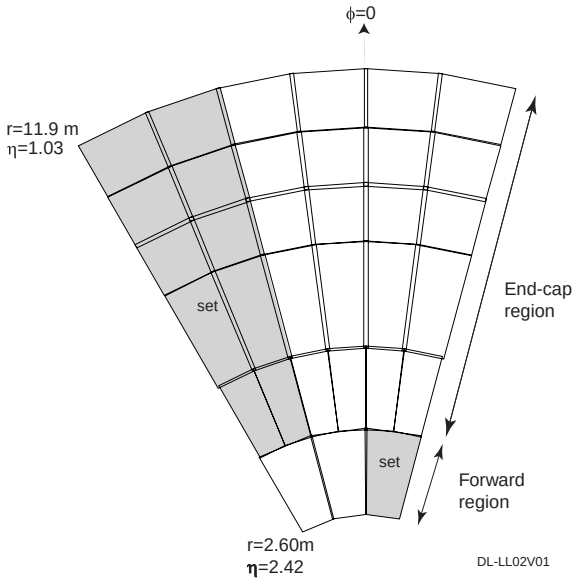

Figure 10-5 A TGC doublet octant showing its component doublet units. The shaded areas denote a Forward set and an End-cap set. The 'double lines' are the chamber overlaps.

10.2.3 Combined system

The geometry of the two independent systems that constitute the muon trigger detectors has been designed to produce a global layout where acceptance is optimized without compromising the system performance. The demands for acceptance require that the two systems, barrel and end-cap, overlap in both η and ϕ in the transition region. Care has been taken to ensure that the two sub-systems can be combined in a way that ensures good trigger performance in the transition regions, while avoiding double counting of muons. Figure 10-6 shows how the overlap is treated in ϕ, where the segmentation of the TGC counters is shown for the sub-sectors of the end-cap detector. The boundaries between the RPC sectors of the barrel have been designed to map onto the boundaries between sub-sectors of the TGC. For more details, see Ref. [10-3]. The overlap in η can be seen in Figure 9-1 in Chapter 9.

Triggers from both sub-systems are treated similarly in the Muon to Central Trigger Processor Interface (MUCTPI) described in Chapter 13 from where they are passed to the Central Trigger Processor. In both η and ϕ, trigger logic exists in the MUCTPI to handle triggers occurring in the overlap region to avoid double triggers from single muons.

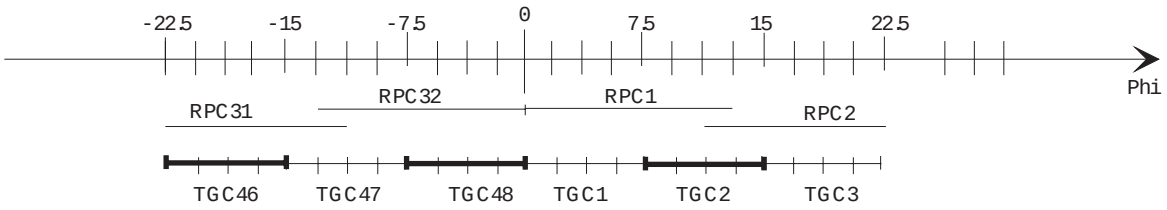

Figure 10-6 The overlapping of the four RPC sectors in the barrel with the TGC sub-sectors of the end-cap for an octant. The TGC 'tick' marks indicate the TGC sub-sector boundaries. The boundaries between the barrel sectors have been designed to map onto the boundaries between the end-cap sub-sectors.

10.3 Detector characteristics

The use of two different chamber technologies in the barrel and end-cap is suited to the different granularity, rate capabilities, production capabilities and cost optimization. It also requires different matching electronics which are described in Section 10.4.

10.3.1 Barrel

The RPC is a gaseous detector providing a typical space–time resolution of $1\,\text{cm} \times 1\,\text{ns}$ with digital read-out. The basic RPC unit is a narrow gas gap formed by two parallel resistive Bakelite plates, separated by insulating spacers. The primary ionization electrons are multiplied into avalanches by a high, uniform electric field of typically 4.5 kV / mm. Amplification in avalanche mode produces pulses of typically 0.5 pC. The candidate gas mixture is based on tetra-fluoro-ethane ($C_2H_2F_4$), a nonflammable and environmentally safe gas. The signal is read out via capacitive coupling by metal strips on both sides of the detector. A trigger chamber is made from two rectangular detector layers, each one read out by two orthogonal series of pick-up strips as described in Section 10.4.1. RPCs have a simple mechanical structure, use no wires, and are therefore simple to manufacture. The 2 mm thick Bakelite plates are separated by polycarbonate spacers of 2 mm thickness which define the size of the gas gap. The spacers are glued on both plates at 10 cm intervals. A 9 mm wide frame of the same material and thickness as the spacers is used to seal the gas gap at all four edges. The RPC performance specifications are summarized in Table 10-2.

Table 10-2 RPC performance specifications

Intrinsic detector efficiency (single layer)	98.5%
Efficiency including spacers and frames	>97%
Intrinsic time jitter (within 3σ)	<5 ns
Time jitter including strip propagation	<10 ns
Rate capability	~1 kHz / cm^2
1 MeV photon sensitivity	1%
Gas gap	2 mm
Operating voltage	~9.0–10 KV

The outside surfaces of the resistive plates are coated with thin layers of graphite paint which are connected to the high-voltage supply. These graphite electrodes are separated from the pick-up strips by 200 μm thick insulating films which are glued on both graphite layers. The read-out strips are arranged with a pitch varying from 23.0 mm to 30.0 mm.

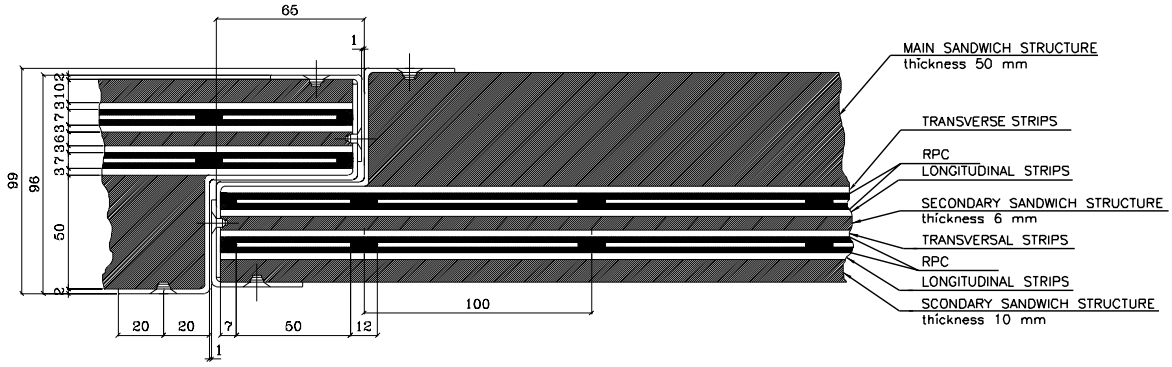

Figure 10-7 Schematic representation of a trigger station in the barrel system.

Each trigger station is made from two detector layers and four read-out strip panels. These elements are held together rigidly by two support panels which provide the required mechanical stiffness of the chambers. The panels are made of polystyrene sandwiched between two aluminium sheets. The mechanical structure of an RPC chamber is shown in Figure 10-7; the total thickness is 100 mm.

The RPCs will be operated with a gas mixture of tetra-fluoro-ethane, $C_2H_2F_4$, and a small component of isobutane, C_4H_{10}, (97:3), with a total gas volume of 18 m^3. Recent laboratory tests have shown that the addition of a small amount of sulphur hexafluoride (SF_6) makes the mixture streamer-less for a large voltage range [10-4]. In Figure 10-8 the RPC efficiency and streamer probability are plotted as a function of the applied voltage for the $C_2H_2F_4/C_4H_{10}$ mixture, with and without SF_6. The gas mixture with e.g. 1% of SF_6 gives about 1kV high-voltage plateau in pure avalanche operation.

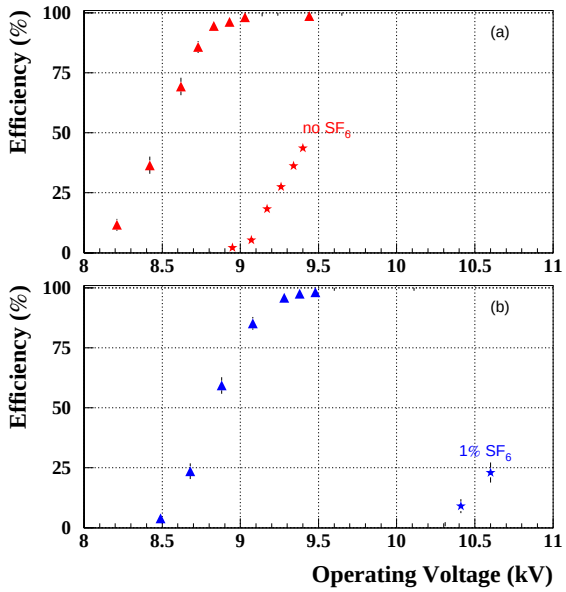

Figure 10-8 RPC efficiency and streamer probability as a function of operating voltage, without (a), and with (b), the addition of 1% SF_6 to the standard gas mixture.

10.3.2 End-cap

Thin Gap Chambers are similar to multi-wire proportional chambers but with an anode wire pitch greater than the anode–cathode distance. The gas used in the chambers is a highly quenching mixture of CO_2 (55%) and *n*-pentane (45%) permitting operation in a saturated mode. This combination of gas and geometry offers small sensitivity to deformation, small dependence of pulse height on incidence angle and a Gaussian pulse-height distribution with no streamer formation. Signals from anode wires and from read-out strips orthogonal to these wires provide bunch-crossing identification with an efficiency greater than 99% for the 25ns gate of ATLAS [10-5]. The use of wires enables high granularity with very small cross-talk between wire-groups (<3%). The use of a strongly quenching gas combined with the fast response of the TGCs allow operating at rates exceeding 20kHz/cm^2.

Table 10-3 Principal TGC parameters.

Gas gap	2.8mm
Anode wire pitch	1.8mm
Wire diameter	50µm
Wire potential	3100V
Gas mixture	CO_2 / *n*-C_5H_{12} (55% / 45%)
Gas amplification	10^6
Anode r/o pitch	7.2 – 39.0mm
Time resolution	> 99% efficiency for 25ns gate
Rate capability	tested at 30kHz/cm^2
Read-out strip width	14.6–49.1mm
Total number of anode r/o channels	280 000
Total number of strip r/o channels	95 000

The chambers are constructed from a plane of anode wires sandwiched between cathode planes of graphite coated glass-epoxy laminate (FR-4). TGCs are built as units of doublets or triplets, with the two or three gas gaps within a unit separated by 20mm thick paper honeycomb. Several anode wires are fed to a common amplifier-shaper-discriminator for input to the trigger electronics. The number of wires per group varies according to the desired granularity as a function of pseudorapidity. The structure of the chambers is shown in Figure 10-9 and the parameters of the chambers are summarized in Table 10-3.

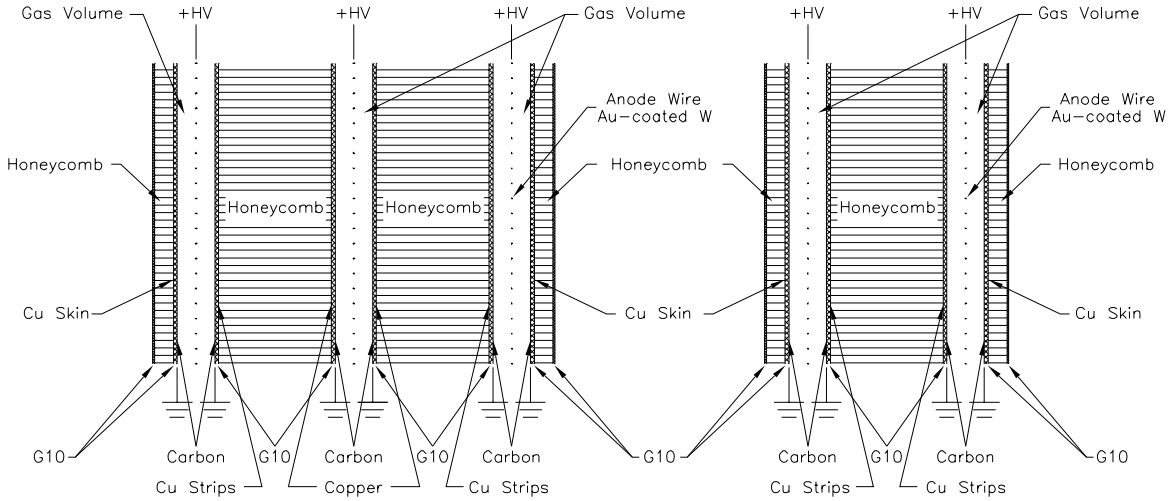

Figure 10-9 Cross-section of a triplet (left) and of a doublet (right) of TGCs. Note that the triplet has strips for the azimuthal coordinate read-out on the two outer layers only. (The gas gap is not drawn on the same scale as the other elements.)

10.4 Front-end electronics

10.4.1 Barrel

A RPC operating in avalanche mode produces typically a single signal of 5ns FWHM and 1.5ns time jitter. A typical RPC timing distribution for the measured time-difference between two detectors for normally-incident particles in a beam-test is shown in Figure 10-10.

Table 10-4 RPC front-end electronics and signal specifications.

Input voltage threshold	150 mV
Shaping time	5 ns
Time skew	<100 ps
Dissipated power	<50 mW/ch

To preserve the excellent intrinsic time resolution of the RPCs, the read-out strips are optimized for good transmission properties and are terminated at both ends to avoid signal reflections. The strips have a pitch between 2.3cm and 3.0cm. The front-end electronics is based on an eight channel chip which is mounted on boards fixed at the edge of the read-out panel between the strip and the ground planes. The boards are 3mm thick to fit the read-out panel thickness and 1cm wide. The length of the boards is tuned with the width of the strips so that each board input is essentially in contact with the corresponding strip and no connection wires are required. The input-to-output delay is the same for all eight channels. The front-end circuit is a three-stage voltage amplifier connected to a variable-threshold comparator. Figure 10-11 shows the block scheme of the front-end amplifier. It is implemented in GaAs technology. The amplifier frequency response was optimized for the typical time structure of the RPC avalanche signal according to the following conditions:

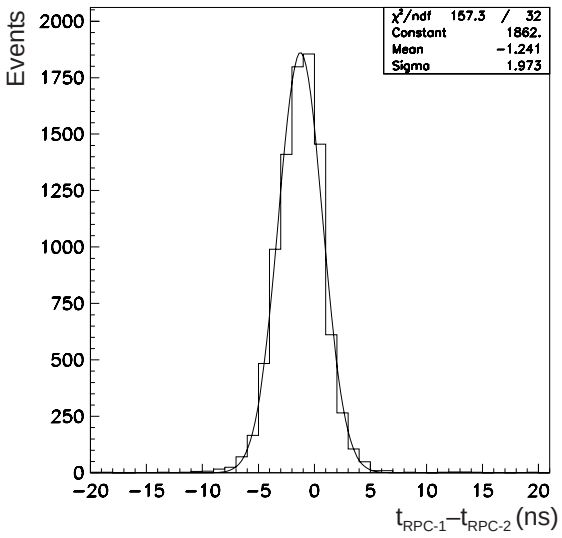

Figure 10-10 Time-of-flight between two RPCs. The resolution of one chamber is the σ of the Gaussian fit divided by $\sqrt{2}$.

- same rise-time for the amplified and the input signals
- minimum return-to-zero time for the output signal

The resulting frequency response has a maximum at 100MHz and a 3dB bandwidth of 160MHz. The amplifier output is bipolar giving zero integrated charge thus avoiding a possible dependence of the steady output voltage on the counting rate. A two-channel, full-custom prototype chip of the front-end circuit has been manufactured in 0.7 mm GaAs technology. This chip, which is a preliminary stage before the final eight channel version, has been successfully tested on board and also mounted on the detector for cosmic-ray and beam tests [10-6]. Electronic and signal specifications are summarized in Table 10-4.

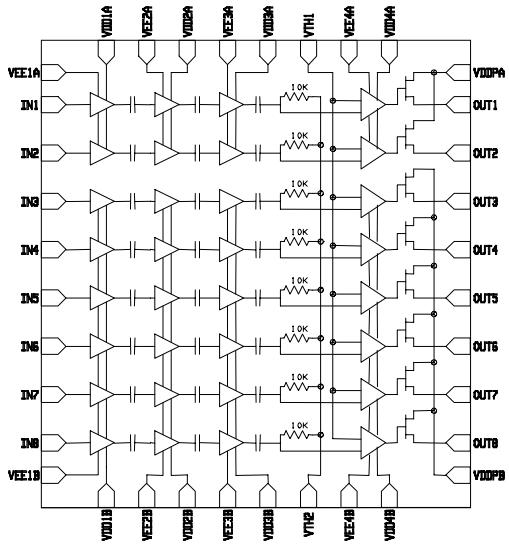

Figure 10-11 Block diagram of the 8-channel chip front-end amplifier for the RPC system.

10.4.2 End-cap

Signals for the level-1 muon trigger system in the end-cap come from TGC chamber anode wires and from strips. Several wires (the number depending on the granularity required as a function of η) are ganged together to provide an anode signal. Wires are spaced 1.8 mm apart and the cathode plane is 1.4 mm from the anode wires. Detector and signal parameters of TGCs are shown in Table 10-5.

Table 10-5 Detector and signal parameters of TGCs.

Anode wire length	39–167 cm
Number of wires per group	6–20
Strip length	104–216 cm
Wire-group capacitance	22–390 pF
Strip capacitance	50–60 pF
Wire propagation velocity	27 cm / ns
Total wire signal (electrons in 5 ionization clusters)	1×10^7
Total strip signal (electrons in 5 ionization clusters)	$2-5 \times 10^6$

The timing resolution of the system is such that there is a 99% efficiency for signals arriving within the 25 ns gate (for perpendicular tracks). The measured arrival time distribution is shown in Figure 10-12. The tails of the timing distribution are mostly due to particles traversing the chamber in the low electric field regions midway between wires, a phenomenon that will be minimized in ATLAS due to the angle of incidence of muons into the TGC counters (>10°). The ASD circuit has a negligible effect on the timing resolution.

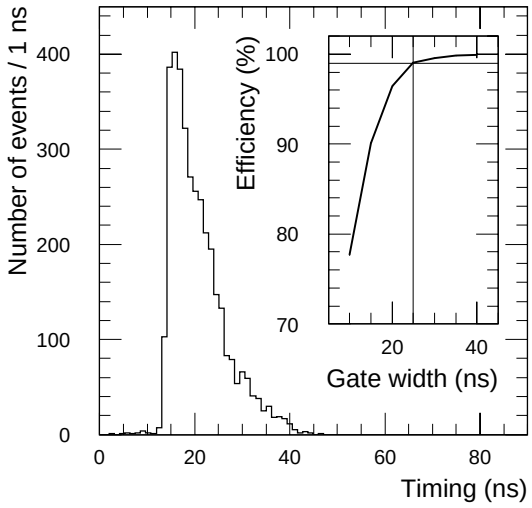

Figure 10-12 Measured arrival time distribution for a minimum-ionizing particle, and the fraction of the arrival times that are within the given time width. There is 99% efficiency for a gate width of 25 ns, at 0° incident angle.

The signals emerging from the chambers are fed into a two-stage amplifier in an amplifier-shaper-discriminator (ASD) circuit [10-7]. Four such circuits are built into a single ASD chip and four ASD chips incorporated into an ASD board. Each such 16 channel ASD board is attached to the edge of a TGC. The final prototype of the ASD chips has been produced and incorporated into the prototype of the ASD board. A number of such ASD boards were tested together with a full-size prototype TGC doublet and triplet and operated successfully at a threshold of 3σ above the thermal noise. There will be a total of ~24500 ASD boards corresponding to ~375,000 signals input to the trigger logic from the TGC wires and strips.

10.5 References

10-1 G. Mikenberg et al., *Thin-gap gas chambers for hadronic calorimetry*, Nucl. Instrum. Methods A265 (1988) 223.

10-2 *ATLAS Muon Technical Design Report*, CERN / LHCC / 97–22, May 1997.

10-3 P. Farthouat, *Current Understanding of the Muon Level-1 System*, Revision 2, ATLAS note DAQ–NO–089, May 1998.

10-4 P. Camarri et al., *Streamer suppression with SF_6 in RPCs operated in avalanche mode*, (to be published in Nucl. Instrum. Methods).

10-5 Y. Arai et al., *Timing optimization of thin gap chambers for the use in the ATLAS muon end-cap trigger*, Nucl. Instrum. Methods A367 (1995) 398–401.

10-6 R. Cardarelli et al., *RPC front-end electronics for the ATLAS Level-1 Trigger Detector*, Proc. of the 7th Pisa Meeting on Advanced Detectors, Isola d'Elba, May 1997, to be published in Nucl. Instrum. Methods.

10-7 O. Sasaki et al., *Front-end electroncis for the Thin Gap Chambers (TGC) in ATLAS forward muon trigger system*, CERN / LHCC / 96–39, Second Workshop on Electronics for LHC Experiments, Balatonfüred, Hungary, September 1997.

11 Muon trigger for the barrel

11.1 Hardware description of the system

As already described in Chapters 9 and 10, the level-1 muon trigger makes use of dedicated RPC detectors in the barrel region [11-1]. The trigger processing includes a low-p_T trigger and a high-p_T trigger [11-2]–[11-7]. To reduce the rate of accidental triggers, due to low-energy background particles in the ATLAS cavern, the algorithm is performed in both the η and ϕ projections for both low-p_T and high-p_T triggers. A valid trigger is generated only if the trigger conditions are satisfied for both projections. In addition, the ϕ projection gives to the experiment the second muon coordinate with a resolution of 1 cm.

As shown in Figure 11-1, the low-p_T algorithm makes use of information generated from the

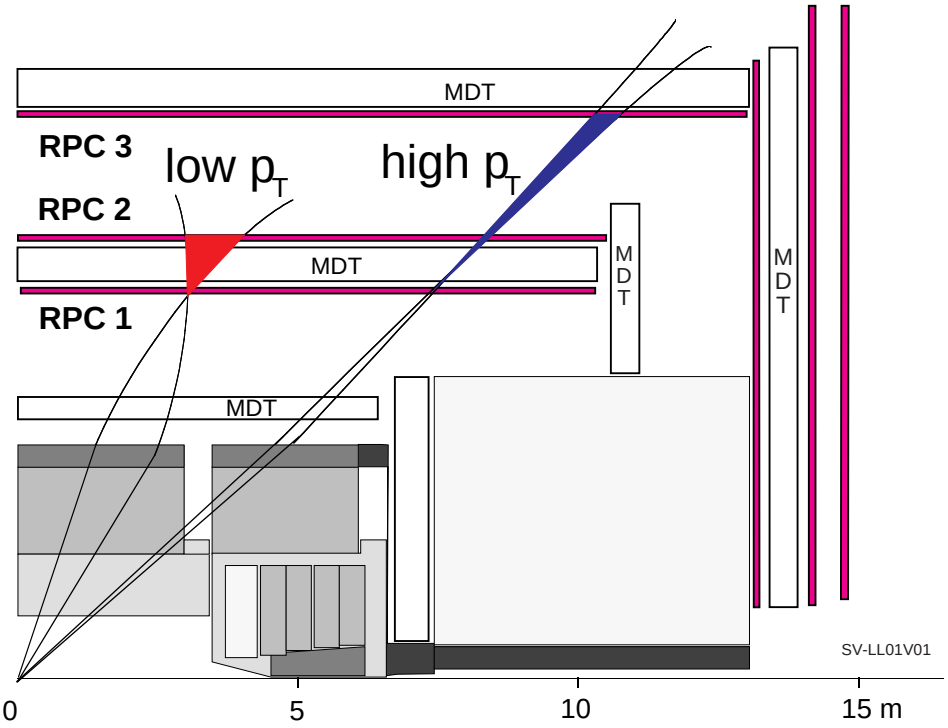

Figure 11-1 Schematic view of the trigger principle.

two Barrel Middle (BM) RPC1 and RPC2 stations. Each station is made of two RPC doublets, one in the η (bending) projection and one in the ϕ (non-bending) projection, and the first stage of the trigger algorithm is performed separately and independently for the two projections [11-8].

The algorithm operates in the following way: if a track hit is generated in the first RPC1 doublet (so-called pivot plane), a search for the same track is made in the second RPC2 doublet, within a road whose centre is defined by the line of conjunction of the hit in the first doublet with the interaction point. The width of the cone is a function of the desired cut on p_T: the smaller the

cone, the higher the cut on p_T. The system is designed so that three p_T thresholds in each projection can be applied simultaneously. In addition, to cope with background from low-energy particles in the cavern, a 3/4 majority coincidence of the four hits of the two doublets is required [11-9].

The η and ϕ trigger information is combined to generate the so-called 'Regions-of-Interest' (RoI), identifying areas in the apparatus in which valid triggers were produced with a spatial resolution of ~ 0.1 × 0.1 in η-ϕ.

The high-p_T algorithm makes use of the result of the low-p_T trigger and of the information generated in the RPC3 Barrel Outer (BO) station. This station is made of two RPC doublets, one in the η projection and one in the ϕ projection. The algorithm operates in a similar way to the low-p_T one. The centre of the road is determined in the same way as for the low-p_T trigger and a 1/2 majority coincidence of the RPC3 doublet and of the low-p_T trigger pattern result is required. As for the low-p_T trigger, three p_T thresholds operate simultaneously and a 0.1 × 0.1 η-ϕ RoI is generated.

11.1.1 Overview of the electronics system

As discussed in Chapter 10, the signals from the RPC detector are amplified, discriminated and digitally shaped on the detector. The Amplifier–Shaper–Discriminator (ASD) boards, each containing eight ADS channels, are attached to the chamber at the end of the RPC strips.

Figure 11-2 shows the context diagram of the overall trigger and read-out system, indicating which of the two main functionalities each board performs.

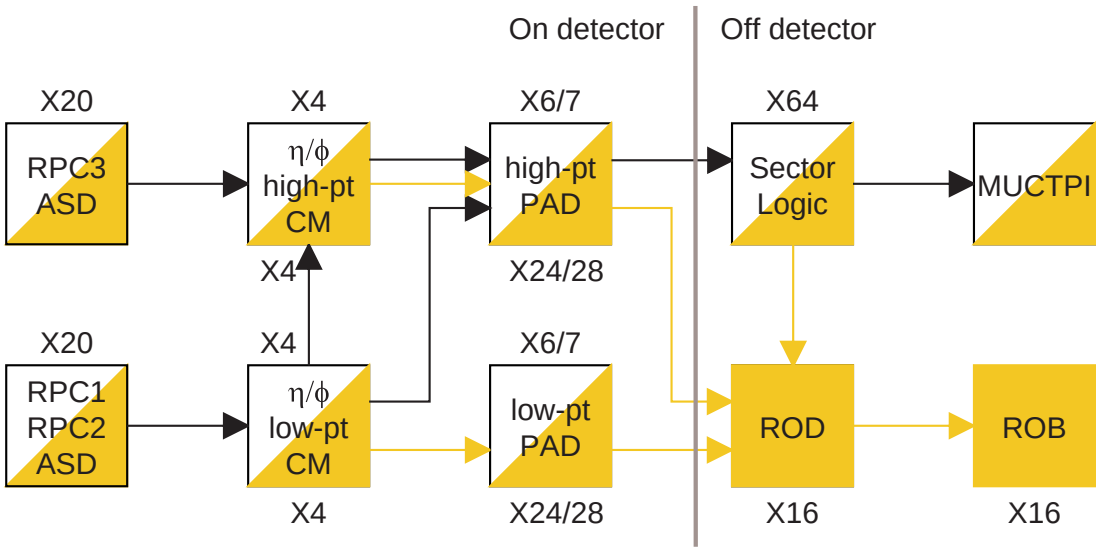

Figure 11-2 Context diagram of the barrel RPC trigger system. Each system component functionality is also indicated (white for trigger, grey for read-out and white/grey for trigger and read-out functionality.

In the low-p_T trigger, for each of the η and the ϕ projections, the RPC signals of the two detector doublets, RPC1 and RPC2, are sent to a Coincidence-matrix (CM) board, that contains a CM chip. This chip performs almost all of the functions needed for the trigger algorithm and also for the read-out of the strips. It aligns the timing of the input signals, performs the coincidence and

majority operations, and makes the p_T cut on three different thresholds. It also contains the level-1 latency pipeline memory and derandomizing buffer. A detailed description of the chip is given in Section 11.1.3.1.

The CM board produces an output pattern containing the low-p_T trigger results for each pair of RPC doublets in the η or ϕ projection. The information of two adjacent CM boards in the η projection, and the corresponding information of the two CM boards in the ϕ projection, are combined together in the low-p_T Pad Logic (Pad) board. The four low-p_T CM boards and the corresponding Pad board are mounted on top of the RPC2 detector, as shown schematically in Figure 11-3.

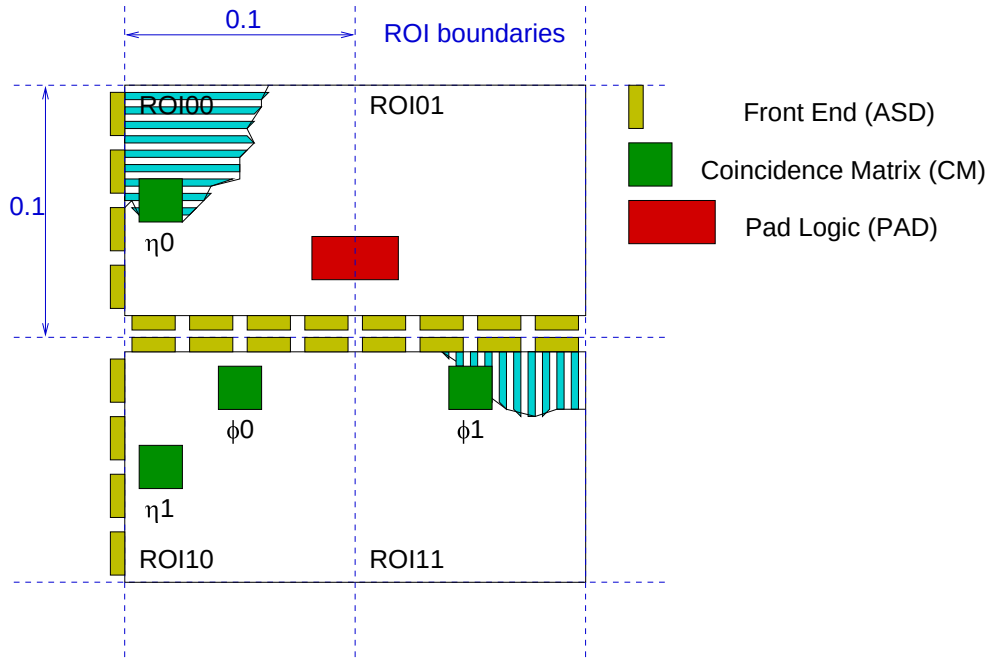

Figure 11-3 On-detector electronics for low- and high-p_T triggers.

The low-p_T Pad board generates the low-p_T trigger result and the associated RoI information. This information is transferred, synchronously at 40MHz, to the corresponding high-p_T Pad board, that collects the overall result for low- and high-p_T.

In the high-p_T trigger, for each of the η and ϕ projections, the RPC signals from the RPC3 doublet, and the corresponding pattern result of the low-p_T trigger, are sent to a CM board, very similar to the one used in the low-p_T trigger. This board contains the same coincidence-matrix chip as in the low-p_T board, programmed for the high-p_T algorithm. The high-p_T CM board produces an output pattern containing the high-p_T trigger results for a given RPC doublet in the η or ϕ projection. The information of two adjacent CM boards in the η projection and the corresponding information of the two CM boards in the ϕ projection are combined in the high-p_T Pad Logic board. The four high-p_T CM boards and the corresponding Pad board are mounted on top of the RPC3 detector.

The high-p_T Pad board combines the low- and high-p_T trigger results. The combined information is sent, synchronously at 40MHz, via optical links, to a Sector Logic (SL) board, located in the USA15 counting room. Each SL board receives inputs from seven (six) low-p_T (high-p_T) Pad boards, combining and encoding the trigger results of one of the 64 sectors into

which the barrel trigger system is subdivided. The trigger data elaborated by the Sector Logic is sent, again synchronously at 40MHz, to the Muon Interface to the Central Trigger Processor (MUCTPI), located in the same counting room.

Data are read out from both the low- and high-p_T Pad boards. These data include the RPC strip pattern and some additional information used in the LVL2 trigger. The read-out data for events accepted by the LVL1 trigger are sent asynchronously to Read-Out Drivers (RODs) located in the USA15 underground counting room and from here to the Read-Out Buffers (ROBs). The data links for the read-out data are independent of the ones used to transfer partial trigger results to the SL boards.

Pad, SL and MUCTPI modules generate themselves read-out data on partial trigger results, in order to monitor the system.

11.1.2 System segmentation

From the trigger point of view, the barrel system is segmented into 64 logically (but not physically) identical sectors.

The barrel is divided in two parts, $\eta<0$ and $\eta>0$. Within each half barrel, 32 sectors are defined. The correspondence between these logical sectors and physical chambers is indicated in Figure 11-4. The Barrel Large (BL) chambers and the Barrel Small (BS) chambers of both middle and outer RPC stations are logically divided in two to produce two large sectors and two small sectors per half-barrel octant.

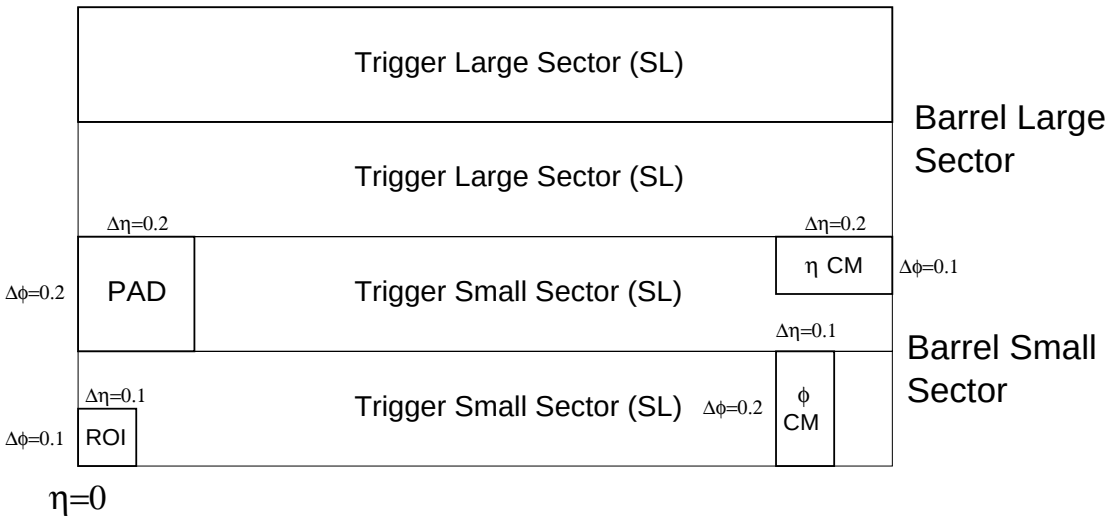

Figure 11-4 Barrel trigger segmentation. Also indicated in the figure are the areas covered by η and φ CM boards, by an RoI, by a Pad Logic (PL) board and by a Sector Logic (SL) board.

Inside a sector, the trigger is segmented in Pads and RoIs. The Pad segmentation is different for large and small sectors; this difference follows the muon spectrometer layout. Two relevant factors that affect the performance and implementation of the trigger system were taken into account in the segmentation: the RPC strip length and the front-end modularity.

The strip should not be too long (the length never exceeds 2.5 m in the bending projection) to avoid a large spread in the signal propagation delay that would make the bunch-crossing identification (BCID) difficult or impossible.

Concerning the front-end modularity, the basic segmentation is 32 channels (two sets of 32 inputs received per RPC doublet) generated by four eight-channel ASD boards. This makes it possible to have the same dimension of the coincidence-matrix inputs in all the apparatus and an easier cabling of the system. Because of layout constraints, in some part of the apparatus it was impossible to follow this segmentation. To handle this case, the CM chip can be programmed to simulate inputs that do not exist.

A large sector contains six Pad regions, while a small sector contains seven Pad regions. Table 11-1 gives the total number of each type of RPC chamber, used in the barrel trigger. The number of strips per chamber is matched to the CM dimensions. The strip length is chosen to give a maximum propagation delay spread of ±7 ns.

Table 11-1 Number of RPC chambers of each type used in the barrel trigger (dimensions in mm).

Name	Nº. chamber	η pitch	η strip length	Nº. chan / chamber	Nº. chan.	ϕ pitch	ϕ strip length	Nº. chan. / chamber	Nº. chan.
BMSa	144	30.0	1480	24	13824	23.1	720	128	36864
BMSb	96	26.2	1480	32	12288	23.1	840	128	24576
BMSc	24	30.0	1480	32	3072	23.1	960	128	6144
BMFa	16	30.0	1480	24	1536	23.1	720	128	4096
BMFb	16	30.0	1480	32	2048	23.1	960	128	4096
BMFc	16	33.8	1480	32	2048	23.1	1080	128	4096
BOSa	8	26.2	1830	32	1024	25.4	840	144	2304
BOSb	16	30.0	1830	32	2048	25.4	960	144	4608
BOSc	120	27.0	1830	40	19200	25.4	1080	144	34560
BOFa	16	30.0	1830	24	1536	25.4	720	144	4608
BOFb	8	30.0	1830	32	1024	25.4	960	144	2304
BOFc	8	27.0	1830	40	1280	25.4	1080	144	2304
BOHa	18	30.0	970	40	1440	24.2	1200	40	1440
BOGa	18	30.0	1105	40	2880	25.1	1200	88	3168
BMLa	64	30.0	1720	24	6144	26.8	720	128	16384
BMLb	192	26.0	1720	32	24576	26.8	840	128	49152
BMLc	64	37.5	1720	32	8192	26.8	1200	128	16384
BML1	16	30.0	1720	16	1024	26.8	480	128	4096
BOLa	16	26.2	2425	32	8192	25.3	840	192	24576
BOLb	128	27.0	2425	40	20480	25.3	1080	192	49152

The region covered by a Pad is ~ 0.2×0.2 in $\Delta\eta \times \Delta\phi$. Inside the Pad the trigger is segmented into RoIs. An RoI is a region given by the overlap of an η coincidence-matrix and a ϕ coincidence-matrix. The dimension of the RoI is ~ 0.1×0.1 in $\Delta\eta \times \Delta\phi$.

The total number of Pads is $7 \times 2 \times 32$ for the small sectors and $6 \times 2 \times 32$ for the large ones, giving 416 Pads altogether. Since one Pad covers four RoIs, the total number of RoIs is 1664.

Concerning system segmentation, a relevant issue is the overlap handling [11-10]. Overlap between different parts of the apparatus cannot be avoided without losing efficiency in the trigger system due to uncovered regions. On the other end, overlap can cause double counting of muon candidates. (The same muon can be counted twice, by adjacent parts of the apparatus.) In the barrel trigger system, overlap is treated and solved at three different levels (Figure 11-5).

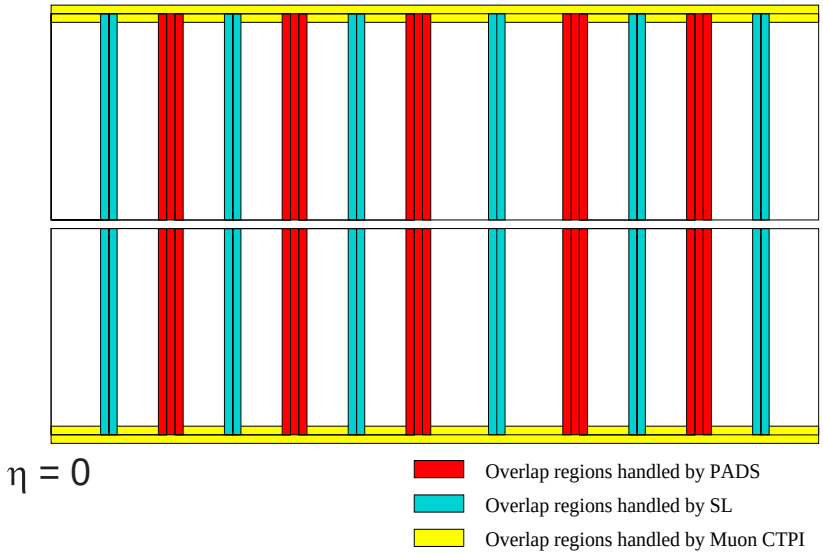

Figure 11-5 Overlap handling.

11.1.2.1 Overlap handling

Overlap within a Pad region

The problem of overlap inside a Pad region is solved by the Pad Logic that removes double counting of tracks between the four RoIs of the region. In addition, if it is found that a trigger was generated in a zone of overlap with another Pad region, this trigger is flagged as ‘border’ trigger and any overlap will be solved later on by the Sector Logic or by the MUCTPI.

Overlap within a sector

As stated above, the triggers generated in a zone of overlap between different Pads are flagged by the Pad logic and sent to the Sector Logic; the Sector Logic prevents double counting of triggers inside a sector. In addition, if it is found that a trigger was generated in a zone of overlap with another sector of the barrel or the end-cap, this trigger is flagged as a ‘border’ trigger and the overlap will be solved later on by the MUCTPI.

Overlap between sectors

Triggers generated in zones of overlap between different Sectors are flagged by the Sector Logic and sent to the MUCTPI which prevents double counting between sectors.

11.1.3 Coincidence-matrix board

The CM board performs almost all of the functions needed for the trigger algorithm and for the read-out of the system. It receives the signals from the RPC doublets and performs the trigger algorithm and read-out functions using a dedicated coincidence-matrix chip. Two separate boards make the low- and high-p_T triggers, using the same CM chip programmed in a different way. In terms of connections to the rest of the system (see Section 11.1.8), both boards have a separate connection to the read-out system, via the RODs (Figure 11-15), while the trigger information of the low-p_T and high-p_T is concentrated on the high-p_T board and then sent to the Sector Logic (Figure 11-17).

11.1.3.1 The coincidence-matrix chip

This chip performs most of the functions needed for the low-p_T and high-p_T triggers and for the read-out [11-11]. These functions are:

- timing and digital shaping of the signals coming from the RPC doublets.
- execution of the trigger algorithm.
- data storage during LVL1 latency.
- trigger and read-out data generation.
- storage of read-out data in derandomizing memory.

Figure 11-6 shows a block diagram of the coincidence-matrix chip. The chip receives the 40 MHz machine clock and through a Delay Locked Loop (DLL) system, it generates a 320 MHz internal clock that synchronizes all the pipeline operations inside the chip.

The estimated dimension of the matrix is 32×48 inputs. In the low-p_T trigger, the signals coming from the RPC1 doublet are connected to the 32 I0 and 32 I1 inputs, and the signals coming from the RPC2 doublet are connected to the 48 J1 and 48 J2 inputs. In the case of the high-p_T trigger the I0 inputs are used to connect the low-p_T trigger result (the I1 inputs are not used), while the J0 and J1 inputs are both used to connect the high-p_T RPC doublet.

The first block at the input of the chip is an edge detector and dead-time circuit. This block detects the rising edge of the arriving signals and, for each detected signal, sets a programmable dead time, of the order of a few hundred nanoseconds, to avoid the arrival of extra pulses in the same RPC channel.

The second block at the inputs allows for the possibility of masking RPC noisy channels.

The third input block is a programmable-depth pipeline that is very important for setting up the timing of the system. Since the cables that bring the signals to the chip could have different lengths, corresponding to different paths, this block allows for timing adjustment of the different groups of signals in such a way that their coincidence is in time. The pipeline can be programmed in groups of eight RPC signals, with a minimum step corresponding to the period

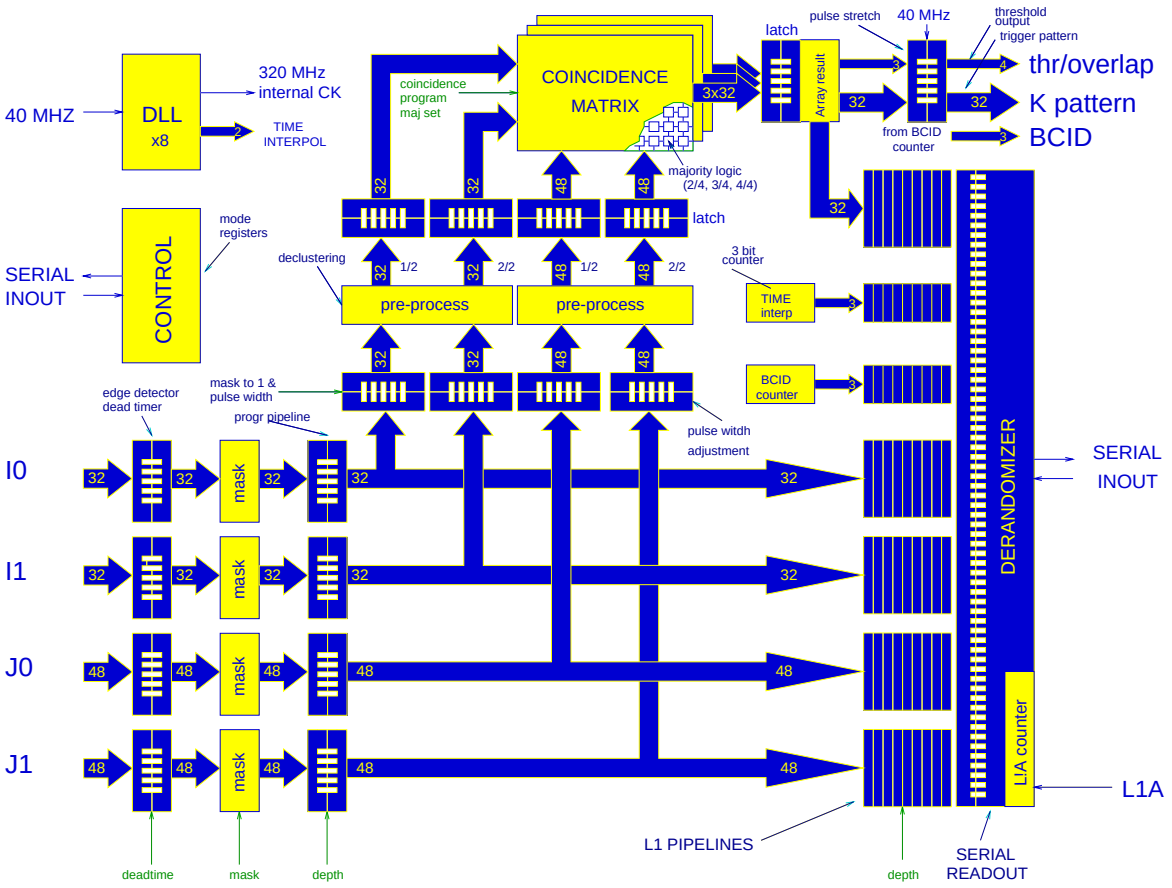

Figure 11-6 Coincidence-matrix chip block diagram.

of the 320 MHz internal clock (~ 3.1 ns). The RPC front-end takes care of aligning in time a group of eight signals, with a maximum spread of 0.5 ns.

At the output of the programmable-depth pipeline, the signals follow two different paths: one to the read-out part and one to the trigger part.

In the trigger part, the first block is the ‘mask to 1 and pulse width’. This block allows for masking to ‘1’ unused inputs to the matrix and for the digital shaping of the signal before making the trigger coincidence. Digital shaping is necessary since the signals must be made as short as possible before the coincidence, to minimize the rate of fake triggers due to low-energy background particles in the ATLAS cavern (uncorrelated noise). This block allows for a programmable digital shaping of the signals, adjusting the duration in steps of ~3.1 ns, corresponding to the internal clock period.

The next block is the preprocessing block. Here the declustering of the signals and the 1/2 and 2/2 majority logic are performed. The RPC detector has an intrinsic cluster size of about 1.5 strips and a declustering algorithm is needed to define the cluster centre.

A latch follows the pre-processing block. This circuit sets the timing at the input of the coincidence-matrix.

The coincidence-matrix has three times 32×48 cells, since the trigger must be performed on three different p_T thresholds simultaneously. The cells have the possibility to be programmed at

will, through a dedicated serial line, to perform 2/4, 3/4, or 4/4 majority logic, and to be set according to the required cut on the p_T threshold.

The output of the coincidence-matrix is sent to the read-out part, and to the trigger output part of the chip. Before producing the trigger output, the trigger pattern is synchronized back to the 40 MHz machine clock.

The trigger result (total of 39 bits) is produced synchronously at 40 MHz. The 32-bit 'K pattern' contains the RPC1 doublet pattern that has generated the trigger. Two bits are used for coding the highest p_T threshold that was validated by the trigger (in case of no trigger both bits are zero), and two bits are used as a flag to indicate whether the trigger occurred in a possible overlap zone (left or right) within the Pad region. The remaining three bits contain the low-order bits of the BCID number.

The main element in the read-out part of the chip is the level-1 pipeline memory. This memory stores the information for a time corresponding to the ATLAS level-1 trigger latency, which is 2.5μs, including some reserve. Since the memory is clocked at 320 MHz, to reach 2.5 μs of latency the depth of the memory must be ~ 800 steps.

The level-1 latency memory is subdivided in seven blocks: two 48-bit blocks and two 32-bit blocks contain the information from the RPC doublets; one 32-bit block contains the trigger array output result; one 3-bit block contains the BCID counter. The last block contains a 3-bit time interpolator that gives the subdivision by eight of the period and tells the read-out at which moment, within the period, the trigger occurred. This information is needed for the timing calibration of the system (see Section 11.2.4.1).

For events selected by the LVL1 trigger, the data from the level-1 latency memory are transferred to the derandomizing memory from here they are read out serially.

The coincidence-matrix chip will be designed in 0.25μm CMOS technology that has some desirable features for the design of the chip. The technology is radiation tolerant for the levels of neutron and gamma radiation (~10 krad over 10 years) expected in the muon spectrometer. In addition, the deep sub-micron technology allows the 320 MHz working frequency, necessary to maintain the design performances of the trigger. Another relevant issue of this technology, is the low-voltage supply that allows for reduced power consumption. This is very important in the design of the barrel system, since the goal is to put the electronics on the detector without any cooling. Table 11-2 gives the relevant chip specifications.

Table 11-2 Coincidence-matrix chip specifications

Number of logic gates	~ 120000
Memory dimension	~ 120 kbit
External clock	40 MHz
Internal working frequency	320 MHz
Power dissipation	< 1 W max
Number of I/Os	~ 210
Package type	BGA

11.1.4 Low-p_T coincidence-matrix board

The low-p_T CM board performs the low-p_T trigger algorithm in the η or in the φ projection. Figure 11-7 shows a block diagram of the board. IL0 and IL1 are the front-end signals from the RPC1 doublet, and JL0 and JL1 are the front-end signals from the RPC2 doublet. The signals are transmitted from the front-end electronics to the board in Low Voltage Differential Signalling

(LVDS) standard. The signals enter the board via LVDS receivers, where they are translated into CMOS signals, suitable for the CM chip inputs.

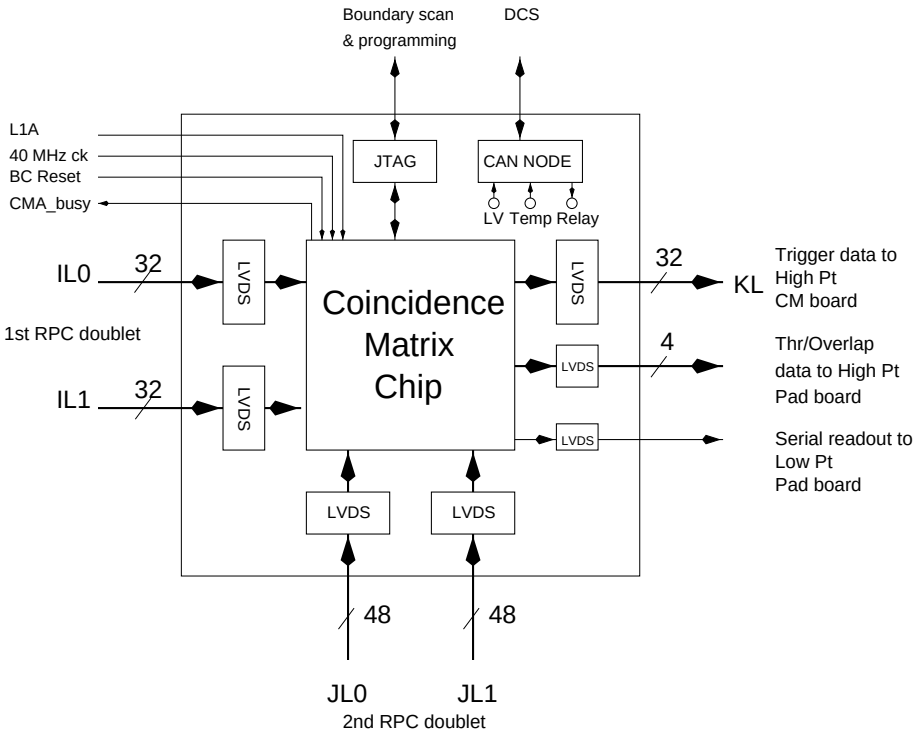

Figure 11-7 Block diagram of the low-p_T coincidence-matrix board.

As already explained above, the trigger algorithm and read-out functions are performed in the coincidence-matrix chip. The chip produces the 32-bit KL trigger pattern to be transmitted to the high-p_T coincidence-matrix board and the 4-bit trigger data pattern to be transmitted to the high-p_T Pad board. The read-out information is transmitted serially to the low-p_T Pad board. All output signals are in LVDS standard.

The CM boards of the ϕ projections must also take into account the bending in the η projection. Since no p_T cut is imposed on the ϕ projection, and since the muon tracks of low momentum can spread over more than one chamber, the OR-ing of the 48 JL0 and JL1 inputs with the corresponding inputs of the adjacent matrices is required, as shown in Figure 11-8.

The control and programming functions of the board are performed by the I2C interface and by the Detector Control System (DCS CAN node). The I2C interface allows programming of the functions inside the coincidence-matrix chip and checking of the functionalities of the circuits through the boundary-scan facility.

The board also receives the Level-1 Accept (L1A), the machine clock and the Bunch-Crossing Reset (BCR) signals of the ATLAS Timing, Trigger and Control System (TTC). It produces a CM Busy (CM_BUSY) signal that is sent back to the Pad Logic to indicated if the derandomizer buffer of the chip is almost full.

The DCS system monitors the temperature of the board and the low-voltage supply. If necessary, the system can switch off the board through a relay. The DCS is based on the use of the CAN field-bus standard, following ATLAS policy.

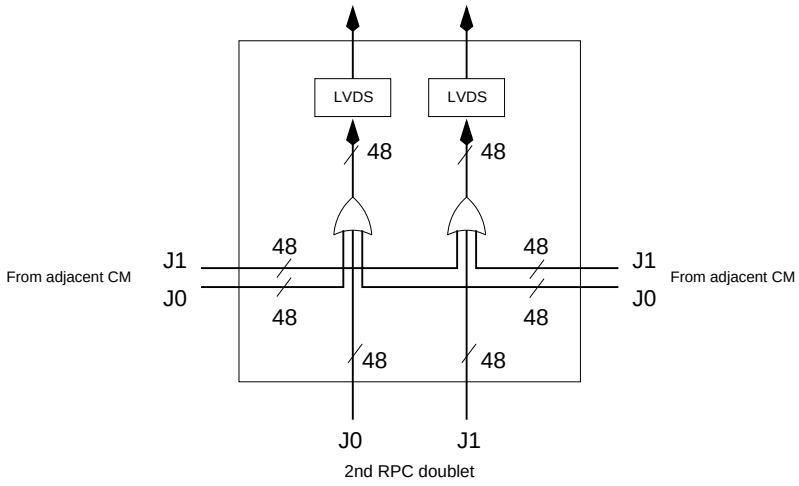

Figure 11-8 Input circuit for the ϕ low- and high-p_T CM boards.

11.1.5 High-p_T coincidence-matrix board

The high-p_T CM board performs the high-p_T trigger algorithm in the η or in the ϕ projection. Figure 11-9 shows a block diagram of the board. The IH input is the pattern result (KL) of the low-p_T trigger, and the JH1 and JH2 inputs are the front-end signals from the RPC3 doublet. The signals are transmitted in LVDS standard and then translated into CMOS signals, suitable for the coincidence-matrix chip inputs.

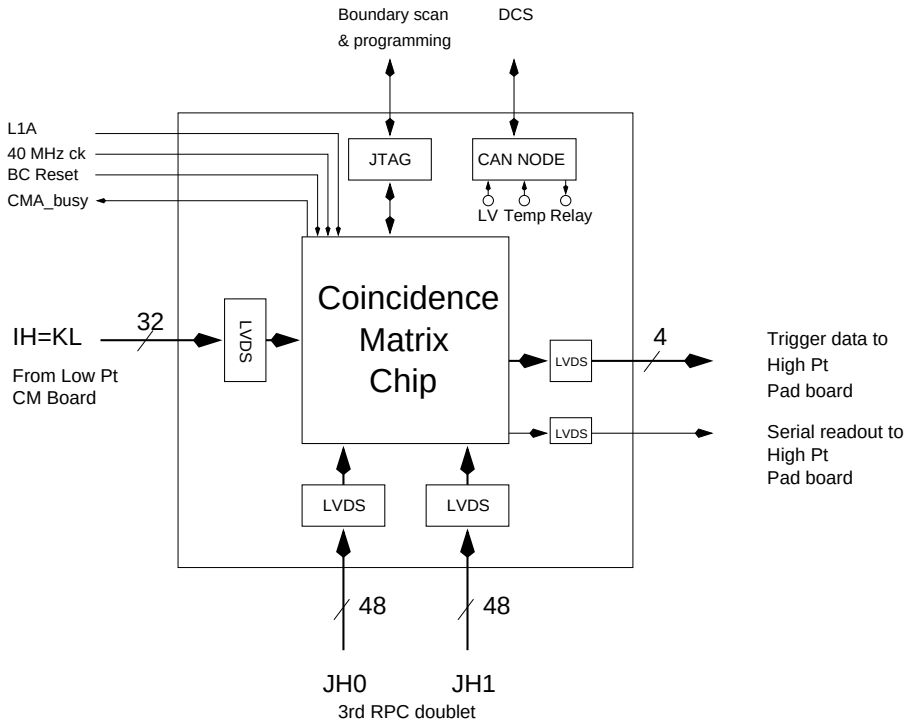

Figure 11-9 Block diagram of the high-p_T coincidence-matrix board.

As for the low-p_T board, the trigger algorithm and read-out functions are performed in the coincidence-matrix chip. The chip produces the 4-bit trigger data and the serial read-out information to be transmitted to the high-p_T Pad board.

All the functions of the board are performed in the same way as for the low-p_T board.

11.1.6 Pad logic board

The information from two adjacent CM boards in the η projection and the information from the two corresponding CM boards in the ϕ projection are combined in the Pad Logic (PL) board. Two different PL boards will be used: one for the low-p_T trigger and one for the high-p_T one. The low-p_T PL board is much simpler than the high-p_T PL board. It only combines the low-p_T read-out information of the two η and two ϕ low-p_T CM boards. The high-p_T PL board combines the high-p_T read-out information of the two η and two ϕ high-p_T CM boards and, in addition, it combines all of the trigger information generated by the four low-p_T CM boards and the four high-p_T CM boards.

The design of the PL board is based on a dedicated gate-array ASIC, developed in radiation-tolerant technology. Although the functions required for the low-p_T PL board are much simpler than those required for the high-p_T PL board, the design of both is based on the same chip, to avoid the overheads and costs of developing two circuits.

11.1.6.1 Pad-logic chip

The most important functions performed by the chip are the combination of the trigger information and of the read-out information from the four CM boards connected to the Pad. The logic associated with these two functions is described below.

PL chip read-out logic

The combination of the read-out information is shown in Figure 11-10. The four serial inputs of the two η and two ϕ matrices are transformed into parallel patterns and then combined together and formatted. The combined pattern is retransformed into a bit-stream suitable for serial transmission to the ROD.

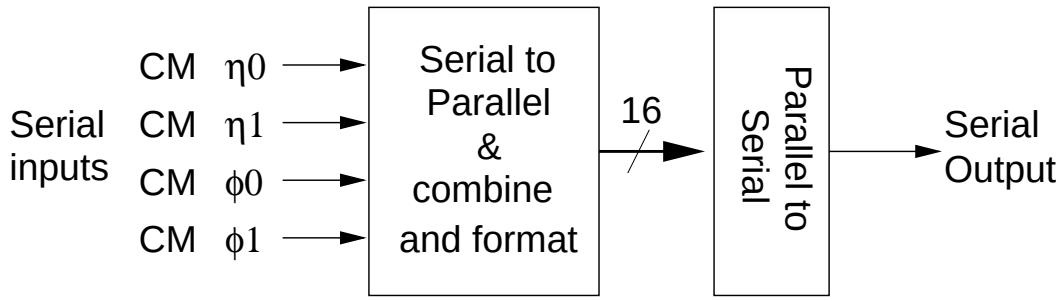

Figure 11-10 Block diagram of the PL chip read-out logic.

PL chip trigger logic

The Pad chip trigger logic combines the low- and high-p_T trigger information in the η and φ projections, assigning track candidates to RoIs. It also deals with the overlap inside the Pad region and tags possible ambiguities in case of more than one track in the region.

Table 11-3 pad-logic trigger output

Bit	Description
0-1	RoI number inside the Pad region
2-5	p_T coding (three low-p_T thresholds + three high-p_T thresholds)
6	Overlap φ flag to be solved by the MUCTPI
7	Overlap η flag to be solved by the Sector Logi
8-9	Ambiguity bits to be sent only to the read-out

The Pad trigger logic is shown in Figure 11-11. The logic works in pipeline mode at the machine clock frequency. The operations are performed in three pipeline steps. In the first step the input trigger data coming from the CMs is aligned in time. In the second step the highest threshold is found and the pad overlaps are solved independently for each projection. Overlap and ambiguity flags are set at this stage. During the third step the two views are combined and the 10-bit output word is produced.

The trigger logic receives as input eight times four bits coming from the four low-p_T coincidence-matrices (two per projection) and the four high-p_T coincidence-matrices.

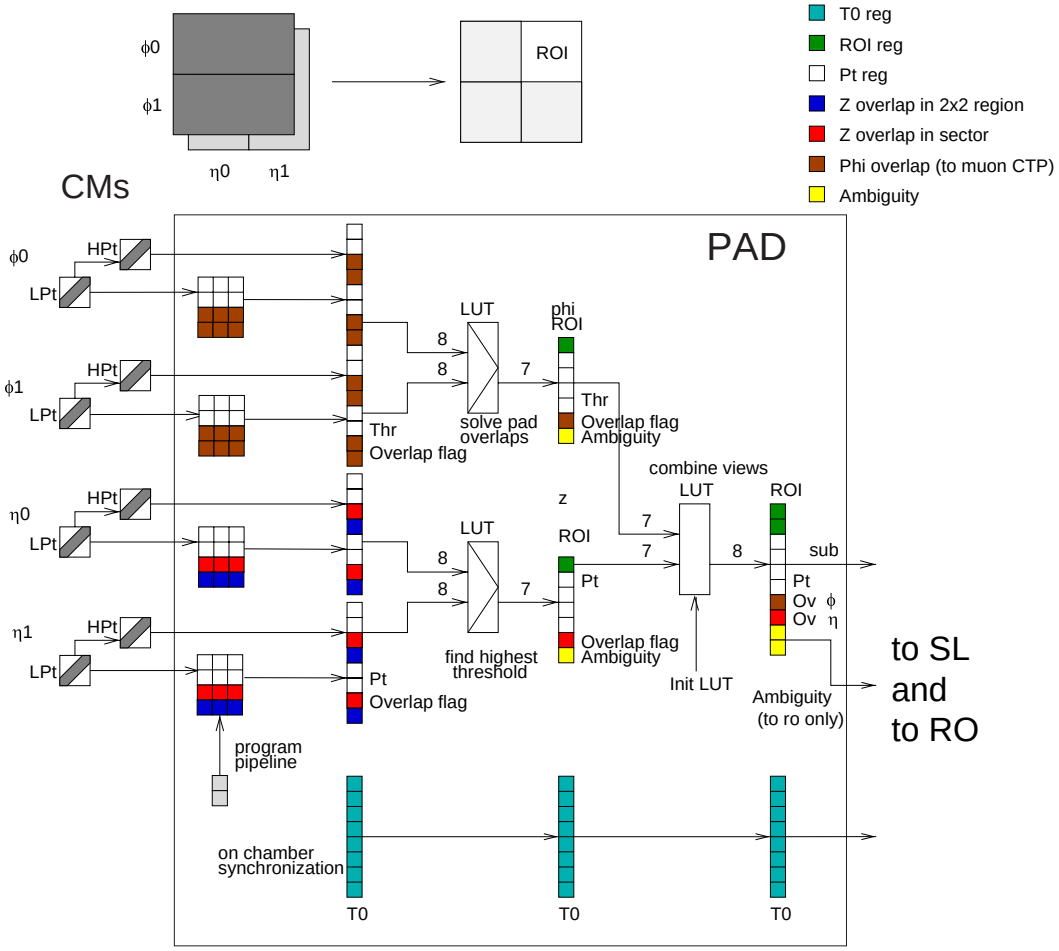

Figure 11-11 Block diagram of the Pad chip trigger logic.

The trigger logic produces a ten-bit output pattern. Eight bits are sent to the trigger Sector Logic, while the full ten bits are sent to the read-out. The content of the output pattern is illustrated in Table 11-3. In case of more than one muon candidate in the pad region, the logic transfers only the highest-p_T candidate and informs the level-2 trigger of the possibility of a second candidate using the ambiguity bits.

11.1.6.2 The Low-p_T Pad-Logic board

The low-p_T pad board combines the read-out information of two η and two ϕ low-p_T CM boards using the read-out logic part of the pad logic chip described above. A block diagram of the low-p_T Pad board is shown in Figure 11-12. The read-out input data are received serially from the

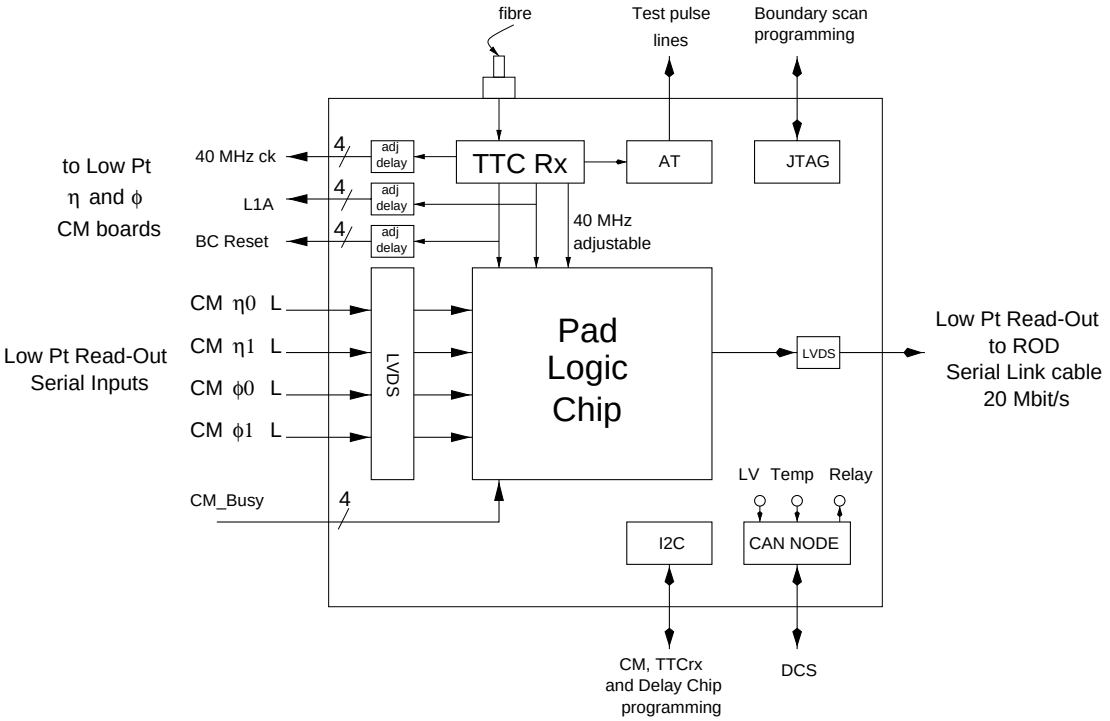

Figure 11-12 Block diagram of the low-p_T Pad board.

CM boards. The output read-out data are sent to the ROD via serial copper cable, asynchronously at 20 Mbit/s maximum data rate.

The high-p_T Pad board also takes care of receiving the trigger and timing signals and of distributing them to the four CM boards to which it is connected. Each Pad board receives the TTC signals over the standard TTC system described in Chapter 16. Details of how the trigger and timing signals are controlled and distributed are given in Section 11.1.10.

All the programming, monitoring and control functions of the board are performed through either the I2C, JTAG or the DCS and the TTC system.

11.1.6.3 High-p_T Pad board

The high-p_T Pad board combines the read-out information of two η and two ϕ high-p_T CM boards and all of the trigger information generated by the four low-p_T CM boards and the four high-p_T CM boards. A block diagram of the high-p_T Pad board is shown in Figure 11-13.

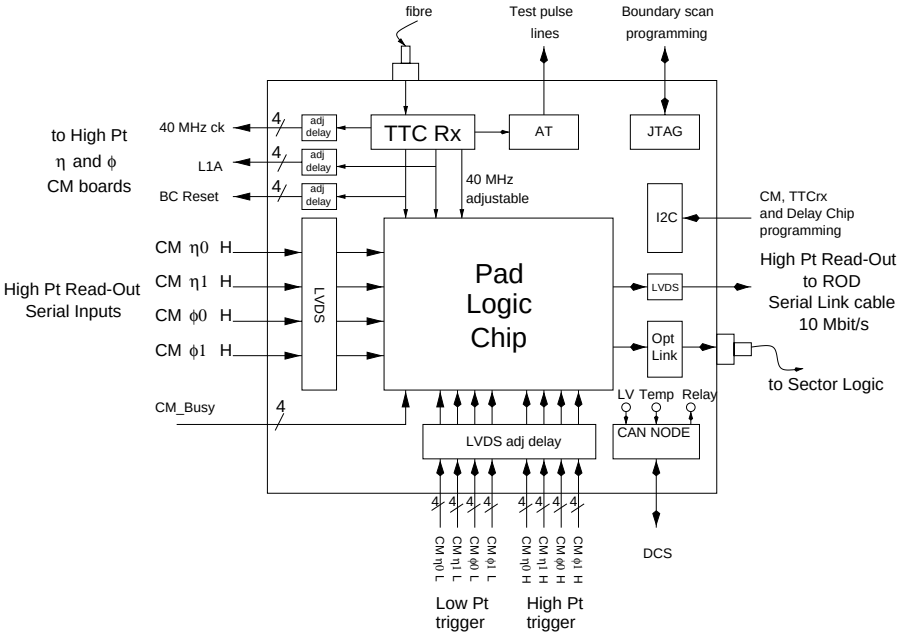

Figure 11-13 Block diagram of the High-p_T Pad board.

The read-out input data are received serially from the CM boards. The output read-out data are sent to the ROD via serial copper cable, asynchronously, at 20 Mbit/s maximum data rate. The trigger input data come in parallel from the low- and high-p_T CM boards. The 8-bit trigger pattern output is sent to the Sector Logic via an optical link, synchronously, at 320 Mbit/s. All the programming, monitoring and control functions, and timing settings of the board, are performed, as for the low-p_T pad board, through the TTC, I2C, JTAG and DCS systems.

11.1.7 Sector logic

The Sector Logic (SL) combines the trigger results for one of the 64 sectors in which the barrel trigger system is subdivided. It is located in the USA15 underground counting room at a maximum distance of 80 m (cable length) from the apparatus. Each SL board receives information from six optical links in the large sectors, and seven optical links in the small sectors. A block diagram of the SL is shown in Figure 11-14.

The logic operates in pipeline mode at 40 MHz. It receives as input seven times (six for the large sectors) an 8-bit pattern, generated by the high-p_T PL boards. In the first clock period, the logic deals with the η overlap within the sector (overlaps between different sectors are solved by the MUCTPI). In the second and third clock periods, the logic sorts the two highest-p_T muon candidates. To reduce the amount of information sent from the SL to the MUCTPI, only two muon candidates per sector are retained. In case of more than two candidates, the SL retains the two highest-p_T tracks and flags that more than two candidates were found in the sector.

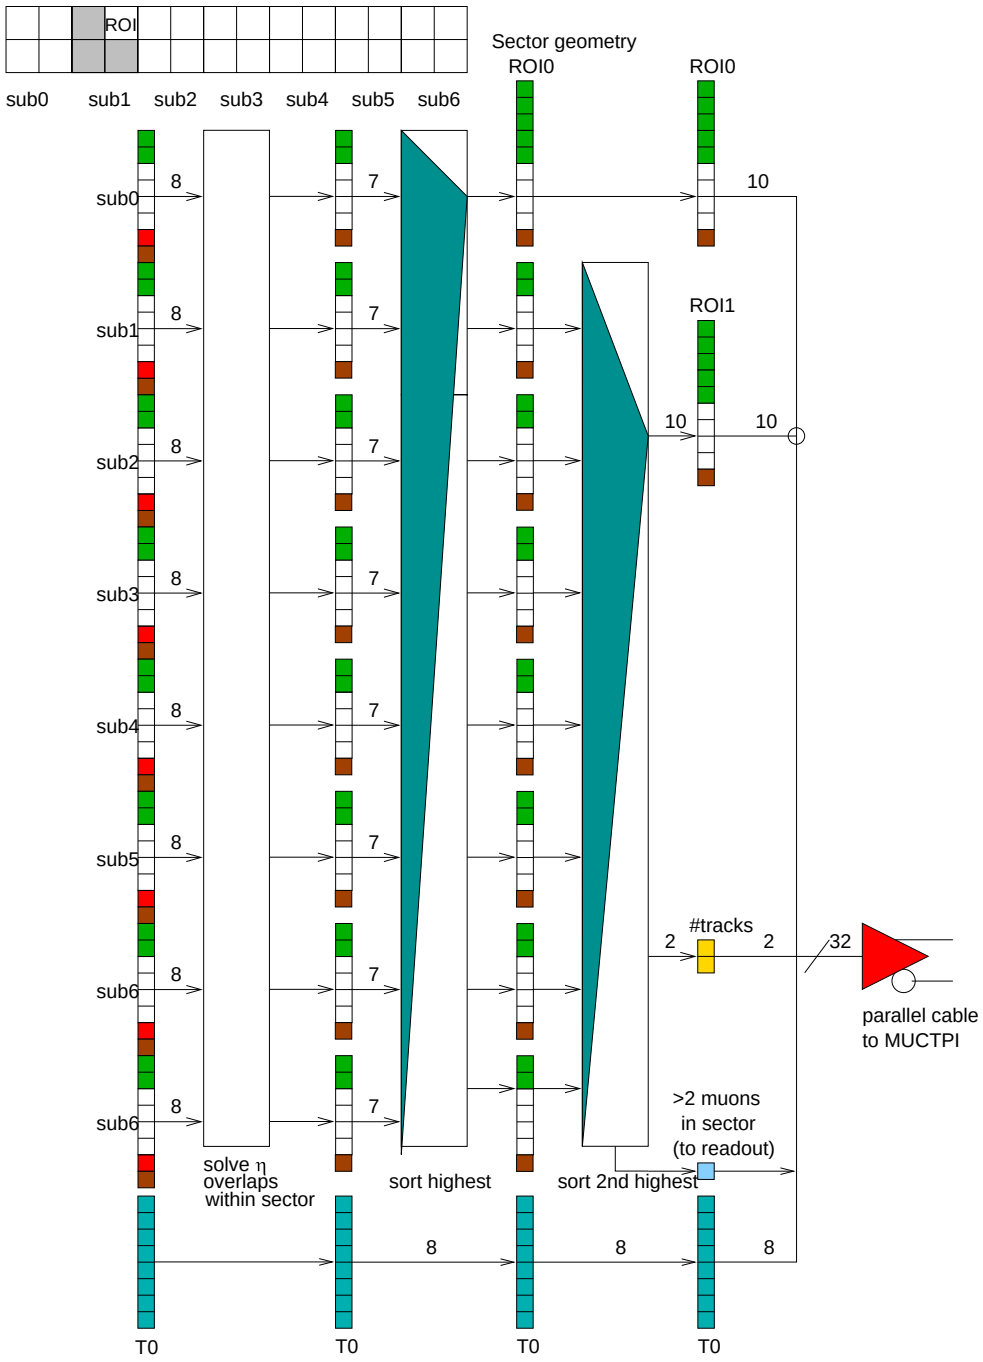

Figure 11-14 Block diagram of the Sector Logic.

As output, the SL produces a 32-bit pattern that is sent synchronously, at 40 MHz to the MUCTPI. The content of the 32-bit pattern is illustrated in Table 11-4.

The Sector Logic will be probably constructed in 9U VME standard and it will be located in USA15, very close to the MUCTPI.

Table 11-4 Sector Logic output format

Bit #	Meaning
1–5	1st candidate RoI number inside the sector
6–9	1st candidate validated p_T threshold
10	1st candidate overlap flag
11–15	2nd candidate RoI number inside the sector
16–19	2nd candidate validated p_T threshold
20	2nd candidate overlap flag
21–22	Number of triggered tracks
23–30	Bunch-crossing number
31–32	Reserved

11.1.8 Read-out scheme

The read-out of the RPC detector that is used for the level-1 muon barrel trigger is described in the ATLAS Muon Spectrometer TDR, but since a large fraction of the read-out is strongly interconnected to the trigger system, we give here, for completeness, a brief description of the read-out scheme.

The whole RPC barrel trigger system is read-out by 16 RODs. The read-out scheme is illustrated in Figure 11-15. The ROD collects the information, of both middle and outer chambers, of two half-barrel large octants. (The same scheme is also valid for the small octants, but here the number of Pads connected to the ROD is smaller.)

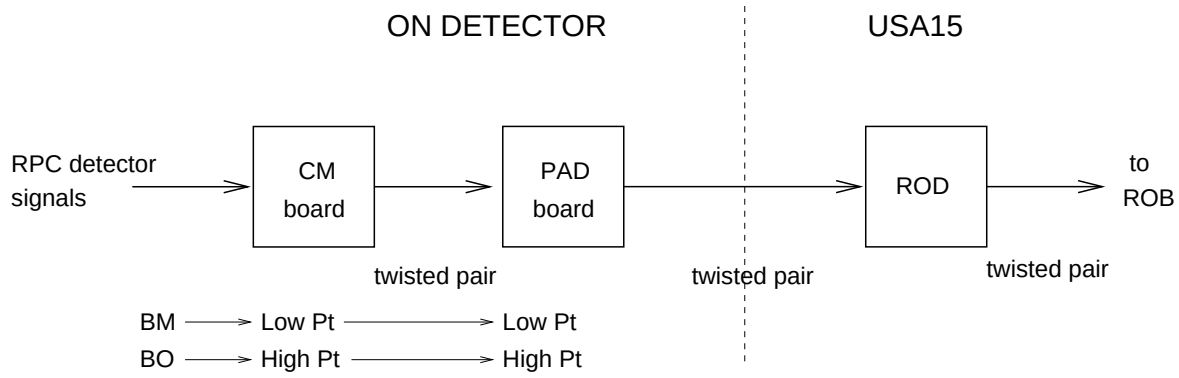

Figure 11-15 Read-out general scheme.

In the large octants, the number of Pads connected to a ROD is $6 \times 8 = 48$. In the small octants, the number of Pads is $7 \times 8 = 56$.

The connection from the Pad to the ROD is implemented through a 20 Mbit/s copper serial link, over a distance of ~ 80 m. The ROD receives the data from 56 (48 for the large octants) links, does the formatting, and sends the data to the ROB via the ATLAS standard read-out link. Figure 11-16 shows a block diagram of the ROD.

The ROD will be located in the USA15 underground counting room. A very preliminary design has been made in which the ROD is composed of four 6U VME slave modules, and one

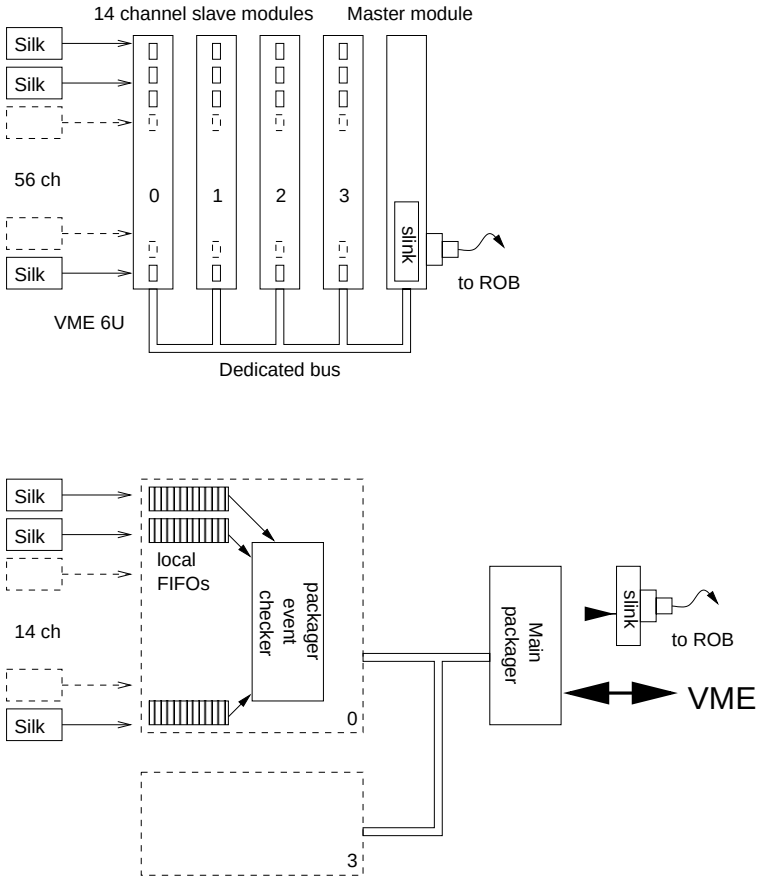

Figure 11-16 Read-Out Driver block diagram.

6U VME master module that controls the activity of the slave modules through a dedicated bus. A slave module houses 14 serial link receivers. The data of each receiver are transferred into a local FIFO from where they are moved to the master module, by a packager circuit. The master module has interfaces to the VME bus and the read-out link, that sends the data to the ROB at a maximum data rate of 1.2 Gbit / s.

11.1.9 System interconnections

The general interconnection scheme is shown in Figure 11-15 for the read-out and in Figure 11-17 for the trigger. The front-end boards are attached to the detector at the end of the RPC strip. The modularity of the front-end boards is eight channels per board. From the front-end, the RPC signals are brought to the CM boards which pass them on to the low-p_T and / or high-p_T Pad Logic boards.

From the Pad Logic boards, the trigger data are sent to the Sector Logic and then to the MUCTPI, while the read-out data are sent to the ROD.

Connection from RPC front-end to CM boards

The connection from the front-end to CM boards is implemented on twisted-pair cables; the numbers and types of cables are shown in Table 11-5.

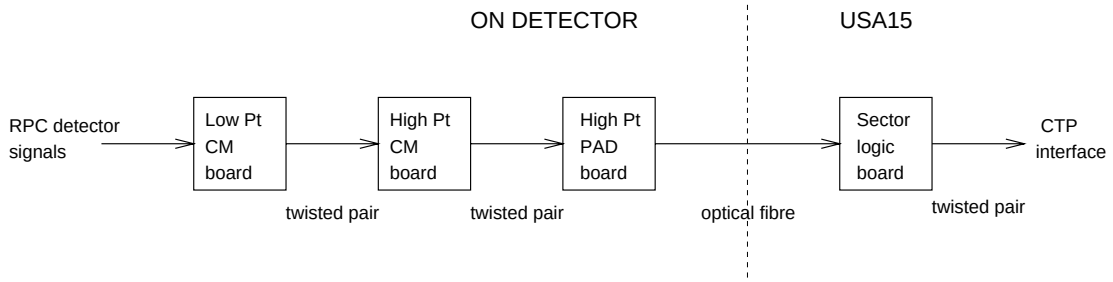

Figure 11-17 Trigger interconnection general scheme.

Table 11-5 Connection from RPC front-end to CM boards.

	Type	Length	Nº. of pairs	Nº. of cables	Signals
Low-p_T	Twisted	2 m	32	3328	IL0, IL1
Low-p_T	Twisted	2 m	48	3328	JL0, JL1
Low-p_T	Twisted	4 m	48	1664	JL0, JL1 from adjacent ϕ CMs
High-p_T	Twisted	2 m	48	3328	JH0, JH1
High-p_T	Twisted	4 m	48	1664	JH0, JH1 from adjacent ϕ CMs

Connection from CM boards to Pad boards and from low-p_T CM boards to high-p_TCM boards

The connection from CM boards to Pad boards is implemented on twisted-pair cables. The numbers and types of cables are shown in Table 11-6.

Table 11-6 Connection from CM to Pad and from low-p_T CM to high-p_T CM boards.

	Type	Length	Nº. of pairs	Nº. of cables	Signals
Low-p_T	Twisted	5 m	4	1664	Low-p_T trigger result
Low-p_T	Twisted	< 1 m	1	1664	Low-p_T serial out
Low-p_T	Twisted	5 m	32	1664	Low-p_T trigger pattern (KL)
High-p_T	Twisted	< 1 m	4	1664	High-p_T trigger result
High-p_T	Twisted	< 1 m	1	1664	High-p_T serial out

Connection from Pad board to ROD

The connection from the Pad to the ROD is implemented using a copper serial link, over a distance of ~ 80 m. The main features of the prototype link (SILK) are shown in Table 11-7.

Four differential pair lines, driving LVDS signals, are used for data, control and clock transmission. A block diagram is shown in Figure 11-18.

Table 11-7 Copper-link prototype main features

Data word size	16 bit
Transmission type	synchronous NRZ
Transmission speed	10 Mbit / s
Board clock frequency	40 MHz
Remote test capability	CCITT 0.151 0.150

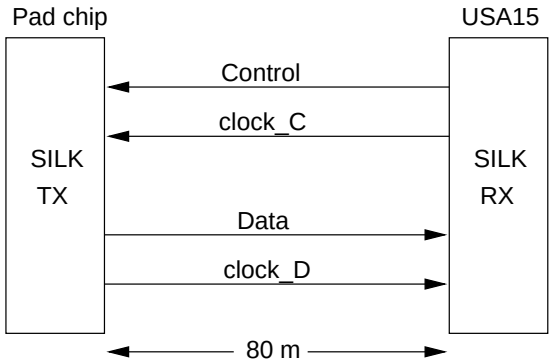

Figure 11-18 Link block diagram.

Connection from high-p_T Pad board to Sector Logic

The connection from the high-p_T Pad board to the Sector Logic is implemented using an optical link. The baseline design is based on a radiation tolerant GaAs transceiver chip developed in Rome by INFN (MATCH chip). This device works on 32-bit words, in asynchronous mode, at a maximum user bit rate of 800 Mbit/s. Table 11-8 shows the chip specifications.

Table 11-8 MATCH chip specifications

32-bit parallel word asynchronous transmission
Auto-reset at each transmitted word
Full duplex
ECL serial Manchester coded data
0.6 to 1 Gbit/s tunable frequency range
Transmission error flag
132-pin PGA package
3.5 Watt power dissipation

Figure 11-19 shows the implementation of the link. Using the MATCH chip, one 32-bit word can be sent every two bunch crossings. Data from the high-p_T pad board from two bunch crossings are packed into the 32-bit word, allowing up to 16 bits per bunch crossing. This leaves ample room to add extra information, such as error detection codes, to the 8-bit output of the PL. Note that the extra latency introduced using this scheme has no effect on the overall LVL1 latency for ATLAS since the RPC trigger result still arrives considerably earlier than those from some other parts of the trigger. Nevertheless, alternatives to the MATCH chip will be considered if they meet the requirements.

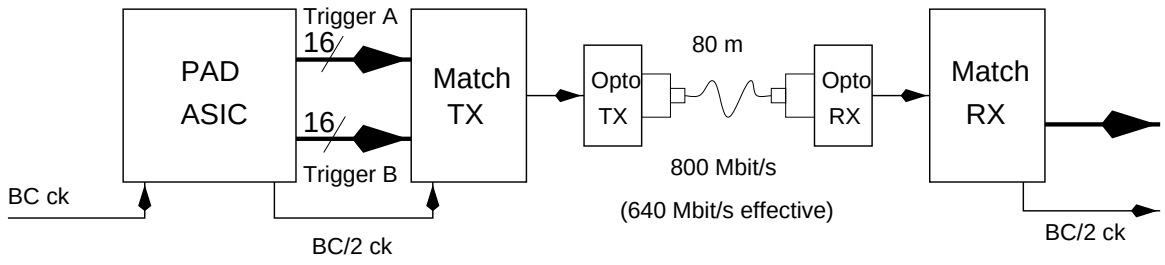

Figure 11-19 Block diagram of the optical link from high-p_T Pad to Sector Logic.

Connection from Sector Logic to the MUCTPI

The Sector Logic output is a 32-bit pattern, synchronous at 40 MHz with the machine clock. Since the SL is located in USA15 very close to the MUCTPI, this connection is implemented on a 32-pair twisted cable carrying differential ECL level signals.

11.1.10 Timing and trigger distribution

The distribution of the timing and trigger is made through the ATLAS TTC system [11-12], based on optical-fibre signal distribution as discussed in Chapter 16. The TTC system for the barrel trigger is segmented in three partitions (Table 11-9). Each partition allows independent running of parts of the trigger system for test purposes.

Table 11-9 TTC system partitioning.

Partition #1	On-detector $\eta > 0$	416 destinations	one per Pad
Partition #2	On-detector $\eta < 0$	416 destinations	one per Pad
Partition #3	USA15	80 destinations	16 RODs + 64 SL
Total		912 destinations	

The on-detector timing and trigger distribution is shown in Figure 11-20. From the TTCrx chip

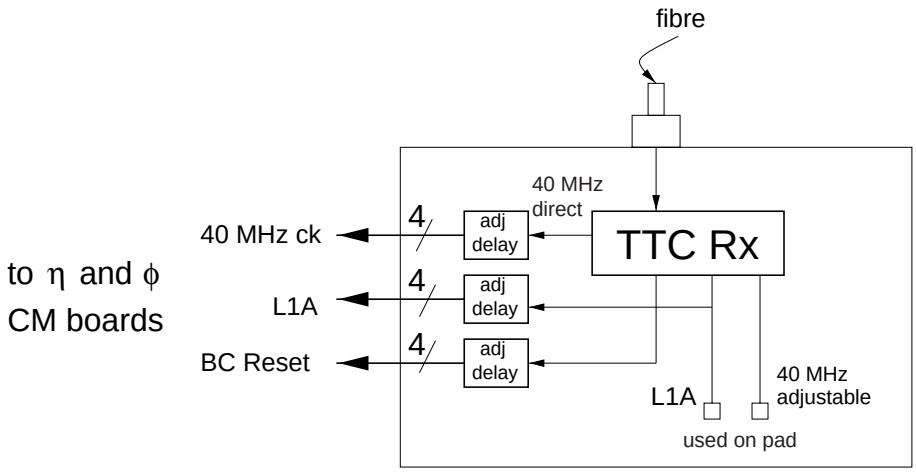

Figure 11-20 On-detector trigger and timing distribution.

mounted on each board, the L1A, BC Reset and three machine-clock signals (one direct and two with adjustable delays) are available.

The L1A signal is used inside the Pad board and it is also sent via a four-channel adjustable delay chip to the four CM boards connected to the Pad board. One of the two adjustable clock signals is used inside the Pad board, while the direct clock signal is sent to a four-channel adjustable delay chip and then to the four CM boards. Also the BC Reset signal is sent through a four-channel delay chip to the CM boards.

A test-pulse command (PPS) is also decoded from the TTCrx chip and then delayed and fanned out to eight RPC test-pulse inputs.

The use of the TTC signals and commands will be explained in Section 11.2.

11.1.11 The trigger latency

A summary of the contributions to the latency for the barrel trigger is given in Table 11-10. As discussed in Chapters 13 and 18, the trigger results from the barrel trigger arrive earlier at the MUCTPI than those of the end-cap system.

Table 11-10 Level-1 muon barrel latency (in bunch-crossings).

Muon barrel latency	BCs	Total BCs
Time-of-flight	3	3
Chamber response and shaping	1	4
Propagation inside RoI	2	6
Local processing for high- and low-p_T	6	12
Pad Logic	3	15
Connection from Pad to Sector Logic	16	31
Sector Logic processing time	5	36
Time at input to MUCTPI		36 (950 ns)

11.1.12 Environmental issues

In the design of the board trigger system two major environmental problems in the ATLAS experimental cavern have to be faced: damage of the electronics due to radiation and the presence of a high magnetic field. Although the radiation level in the muon spectrometer is less than in other parts of the ATLAS apparatus, the problem of radiation damage has to be addressed. Also, the magnetic field problem cannot be neglected since it could impose some constraints on the location of the electronics.

11.1.12.1 Radiation tolerance

The radiation dose accumulated on the muon spectrometer over ten years of running is indicated in Table 11-11. To limit the problem of radiation damage of the electronics, the first fundamental rule is to reduce, as much as possible, the amount of electronics located in the cavern. A large effort has therefore been made to locate the Sector Logic electronics in USA15. This has increased the number and length of links, but has partially solved the radiation-damage problem.

Table 11-11 Radiation doses accumulated in 10 years of running in the muon spectrometer (worst location).

Neutron radiation	$\sim 6.8 \times 10^{11}$ cm^{-2}
Gamma radiation	~ 5.6 Gy

For a safe operation of the rest of the electronics, located in the cavern, commercial components must be avoided and intensive use must be made of dedicated radiation-tolerant ASICs. As many functions as possible are implemented in such ASICs. Any commercial components to be used in the cavern will be qualified for radiation tolerance.

Concerning the use of radiation-tolerant processes for the design of ASICs, a choice must be made among some existing technologies [11-13]:

DMILL DMILL is a process developed in the framework of the RD–29 project. Its radiation tolerance is well proven for large radiation doses. The process is less dense than other standard processes and the software tool for designing chips still needs to be improved.

GaAs The radiation tolerance of the GaAs process seems to be proven, although some confirmation must still come from measurements. This semiconductor is more expensive than silicon, and the software tools for the design are relatively poor.

Deep submicron CMOS (< 0.5 μm) From preliminary and very encouraging results of the RD–49 project, deep submicron processes seem to be sufficiently radiation tolerant for the radiation level of the muon spectrometer. These processes are commonly available commercially and the software tools are well developed. This could be an acceptable solution for a large fraction of our on-detector electronics (although the development price is rather high). In Table 11-12 some relevant parameters for deep submicron CMOS processes are presented.

Table 11-12 Deep submicron CMOS main radiation-tolerance parameters

Gamma dose	~ 200 krad
Gamma dose with enclosed geometry	~ 10 Mrad
V_T shift	< 10 mV for any dose
Neutron degradation	8% g_m loss @ 10^{14} n cm^{-2}

In the barrel trigger design two kinds of on-detector boards are used: the CM board and the PL board. The main element of the CM board is the CM chip, that will perform almost all the functions of the board. This chip will be developed in 0.25μm CMOS technology. The main element of the PL board is the Pad Logic chip which will be developed in 0.25μm CMOS.

Concerning the rest of the components located on the boards (TTCrx, LVDS drivers and receivers, programmable delays, JTAG interface, CAN node and voltage regulators, and electro-optical converters), we plan to take advantage of the common LHC developments in these fields (some of them already existing).

For the two types of link transmitters located on the apparatus, we plan to use the GaAs MATCH chip for optical transmission of the trigger data, and to incorporate in the pad logic chip the transmitter of the copper link for the read-out data.

11.1.12.2 Magnetic field

The magnetic field can impose some constraints on the location of the electronics, although it is not a problem for on-board electronics, since the use of components sensitive to magnetic field can be avoided.

The magnetic field is a problem for the power supplies since we plan to install them in the cavern. It is thought that the existing power supplies can work well up to ~ 300 Gauss although this needs to be confirmed. However, in ATLAS, 300 Gauss magnetic field can be reached only at ~ 13 m from the interaction point, very close to the wall of the cavern.

CERN has established a working group on power supplies to solve this problem. We will follow the work of the group. At present the solution seems to be in the use of the right materials and in shielding the supplies in a proper way. The first tests are encouraging and the final results will come at the end of 1998.

11.1.13 Detector Control System

The standard ATLAS Detector Control System (DCS) is used in the barrel trigger. The system will take care of controlling and monitoring some parameters on the on-detector electronics. Each board sitting in the cavern has a voltage and temperature control and an on–off relay to switch off the board in case of need. The block diagram of the control system for the on-detector electronics is shown in Figure 11-21.

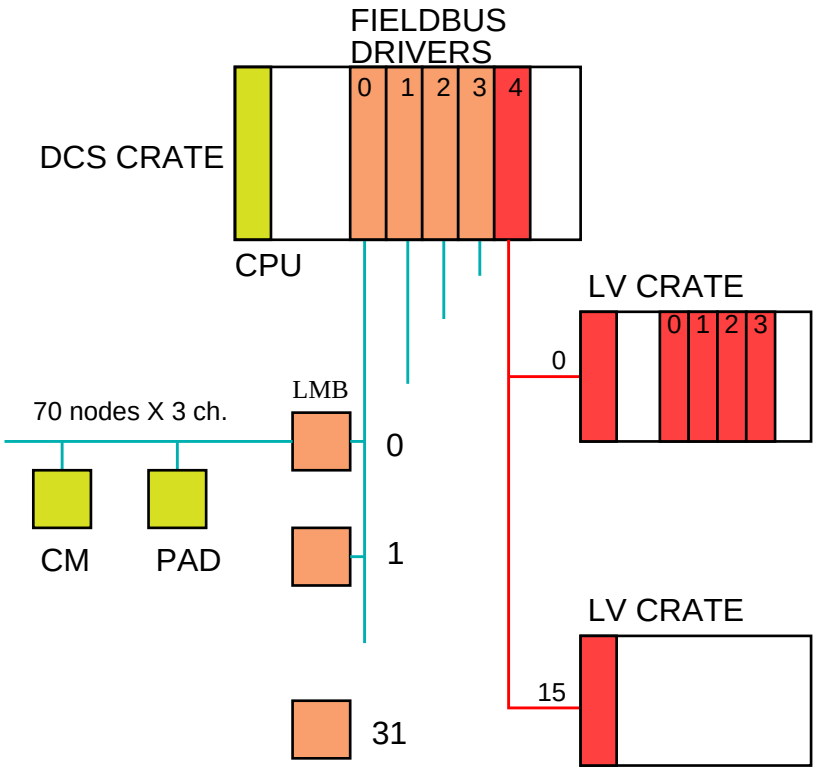

Figure 11-21 DCS block diagram.

One Local Monitor Board (LMB) is installed on the detector for each of the 64 trigger sectors. The LMB will control all the CM boards and PL boards of the sector. The total number of boards for a large (small) sector is 56+14=70 (48+14=62). For each board three DCS channels are needed: one for voltage control, one for temperature control and one for the on–off relay. The communication and control is performed via the CAN standard fieldbus. In total, 12672 DCS channels are required for the whole barrel trigger.

The LMBs are connected to fieldbus driver modules. In total, four fieldbus drivers are needed, since each module can control up to 16 LMBs. The fieldbus drivers are housed in a crate located in the USA15 counting room. This crate also contains a CPU and an additional fieldbus driver to control the Low Voltage system that, in case of need, can be switched off for a complete sector (see Section 11.1.14 below).

11.1.14 Low-voltage system

The issues of power distribution and power dissipation are important in the design of our system. The barrel trigger electronics is distributed over a large area (thousands of square metres). Since it is not allowed to dissipate a large amount of power in the ATLAS cavern, the system should be cooled. But cooling an electronic system spread over such a large area is a relevant problem. Air cooling is not allowed, because this could have some impact on the precision performances of the MDTs. The only solution would be a liquid cooling, but this implies the construction of a long and complex cooling pipe distribution system, and the circulation of a large amount of cooling liquid.

The approach that is taken, is to make a big effort to reduce the power dissipation of the on-detector electronics at a level that could be tolerated without cooling.

A very preliminary estimation of the power dissipated on the on-detector electronics has given a total of 28 kW. This value came from 30 mW / ch for the front-end and 40 mW / ch for the logic, for a total of 400000 channels. Now, after the introduction of some modification in the trigger design, this power dissipation will be reduced. A number of design decisions have been made that minimize the power dissipation: the use of LVDS signal instead of the ECL to transmit the signals from the front-end to the trigger logic, and the intensive use of deep submicron CMOS chips, powered with 3.3V or less, in which we implement almost all the function required by one board. It is clear that the final power consumption will only come when the final detailed design of every part of the apparatus is ready. However, it can already safely be said that the power dissipation will be < 20 mW / ch for the front-end and < 30 mW / ch for the trigger logic. Based on that, we have estimated the power consumption of the different parts of the apparatus illustrated in Table 11-13.

Table 11-13 Power consumption of the barrel level-1 muon trigger

Front-end channel	20 mW / ch
Trigger-logic channel	30 mW / ch
CM board	2.5 W / board
PL board	2.5 W / board
Front-end sector	120 W / sector
Trigger-logic sector	180 W / sector
Global dissipation	19.2 kW

The power distribution follows the segmentation of one power supply per sector. In total we will have 64 power supplies, located in 16 crates around the detector at $\eta = 0$. Each low voltage crate is controlled by means of the DCS system based on the use of the standard CAN fieldbus. Figure 11-21 shows how the low-voltage system is controlled. Each low-voltage crate has an interface connected to a fieldbus driver housed in the DCS crate, located in the USA15 counting room.

Table 11-14 Low-voltage specifications

Type	switching
Max power	300 W
Regulation	± 0.15%
Noise and ripple	< 10 mV pp, 0–40 MHz
Short-term stability	0.15% over 24 hours
Temperature coeff.	< 0.02% / 1°C

The trigger logic is supplied only with at least one voltage line per sector at 3.3V or less. Table 11-14 contains the most relevant specification for the trigger-logic low-voltage supplies.

11.1.15 Grounding and shielding

The grounding and shielding scheme of the level-1 muon barrel trigger follows some ATLAS ‘golden’ rules. The RPC system is isolated from the MTD system and from the support and mounting structure. The RPC chambers and the trigger electronics are contained in separated Faraday cages. The grounds of Faraday cages are connected together in groups by the same ground cable (one per sector, 64 in total), that is connected to ground in the USA15 counting room.

All the signals that penetrate the Faraday cage are in LVDS standard or optical. The cage is penetrated only by:

- LV power lines
- I2C lines
- DCS lines (CAN bus lines)
- JTAG lines
- TTC optical input
- LVDS trigger I/Os
- LVDS read-out I/Os
- LVDS PPS lines

Low-Voltage supplies are floating. In the present design, the low voltage is regulated down to the required value on each board. However the possibility to send the voltage already regulated and to monitor it through the DCS is being considered (this eliminates the power loss on the voltage regulator).

The system should provide some kind of ‘ground fault’ warning device, to monitor ground problems.

11.2 Software implementation of the system

In the preceding sections the hardware composition of the level-1 barrel trigger was illustrated, showing the different parts and their functionalities. This section illustrates how the different parts work together in a coherent way, how they interact with each other, how they are controlled, calibrated and monitored, and how test and diagnosis of the different parts of the apparatus are done: In other words it will describe what is called software implementation of the system [11-14][11-15].

The general strategy for the software implementation of the system follows some ATLAS implementation rules and some subdetector-specific ones [11-16]. There are some functions that in ATLAS should be performed in a standard way for all subsystems. Data are read out through the ROD–ROB chain, timing of the apparatus through the TTC system, and use of the standard DCS system for detector control. The other implementation rules are specific to each subsystem.

11.2.1 Data read-out

The barrel has an eight octant (0–7) azimuthal segmentation; each octant is further subdivided in two parts, covered by Barrel Large (BL) and Barrel Small (BS) units respectively; each of the two units is again subdivided in two sectors, covered by RPC chambers. We therefore have a total of four sectors (0–3) per half-barrel octant (see Figure 11-22).

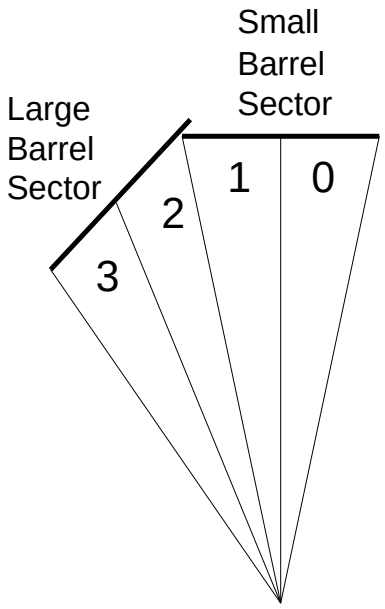

Figure 11-22 Azimuthal view of one half-barrel octant of the trigger segmentation.

The elementary hardware cell, from the read-out point of view is the Coincidence-matrix (CM). The strips of both the η and ϕ views of the RPC chambers are connected to CMs. Four CMs (two in η, two in ϕ) form the low- or high-p_T Pad. The superposition of one η CM and one ϕ CM within a Pad gives a RoI of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. The low-$p_T$ Pad performs the data read-out of the trigger chambers (RPC) in the middle station while the high-p_T one performs the data read out of the outer station.

Figure 11-23 and Figure 11-24 show the RPC chamber partitioning in Large and Small sectors for a half-barrel, including the mapping of the Pads and RoIs. Read-out data from both projections for a given Pad are sent through the same point-to-point copper link to the ROD as shown in Figure 11-25. One ROD combines data from 48 Pads, combining both half-barrels and all the trigger stations for the Large Chambers: another ROD does the same for the 56 Pads of the small chambers. This gives two RODs per octant, that is 16 RODs in total (Figure 11-26); the RODs are located in USA15 and each of them is connected to one ROB, sitting in the DAQ crate.

The RPC read-out is closely connected to the Monitored Drift Tubes (MDT) read-out organization and is based on the assumption that the RPC and MDT data from the same projective area have to be put in the same DAQ crate. This allows local pre-processing in the LVL2 trigger using both RPC and MDT data, the RPC data being used to guide a track search in the MDT.

The choice of the data to be sent to the read-out must allow the following tasks: monitoring of the functioning of the hardware processor, by reproducing the level-1 trigger algorithm with a software program and the same input data; provision of information in a form that gives fast guidance to the level-2 algorithm; compliance with the standard ROB input data format.

The data format is organized according to the hardware structure: CM, Pad and ROD. Zero suppression will be applied at the level of CMs and Pads. An empty structure will not be suppressed to ensure event synchronization at every level, but it will not be propagated along the read-out hierarchy. A further level of zero suppression will be applied within a chamber pattern, keeping only the addresses of the firing strips (this applies to both the input and the output data of each CM). A CM will be read out even if it did not give any trigger.

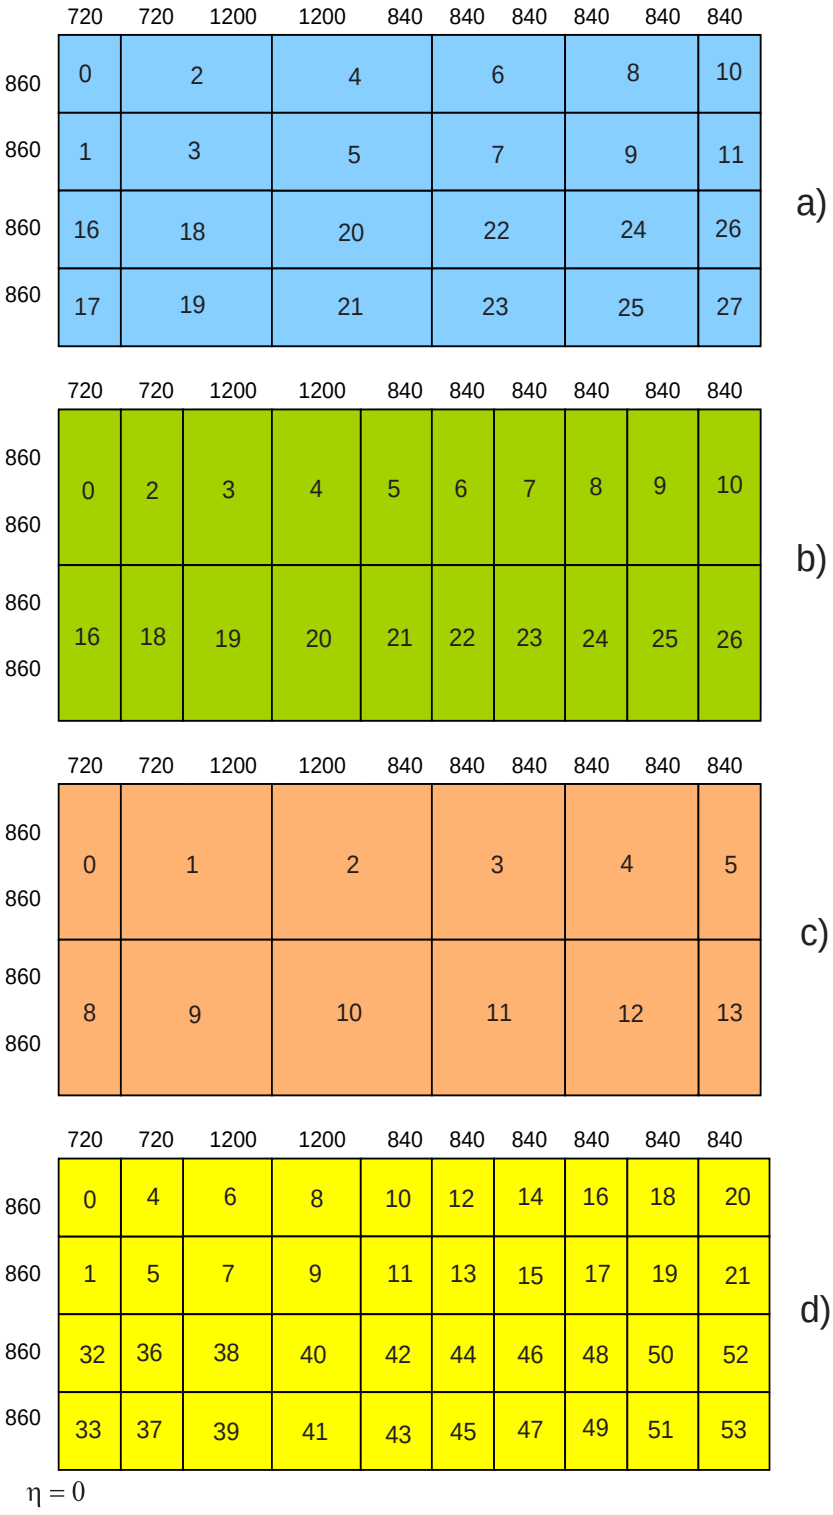

Figure 11-23 Partitioning and numbering conventions of a Large Sector: (a) ϕ strip and CM partitioning, (b) η strip and CM partitioning, (c) Pad partitioning, and (d) RoI partitioning. Chamber dimensions in the RPC1 station are indicated (in mm).

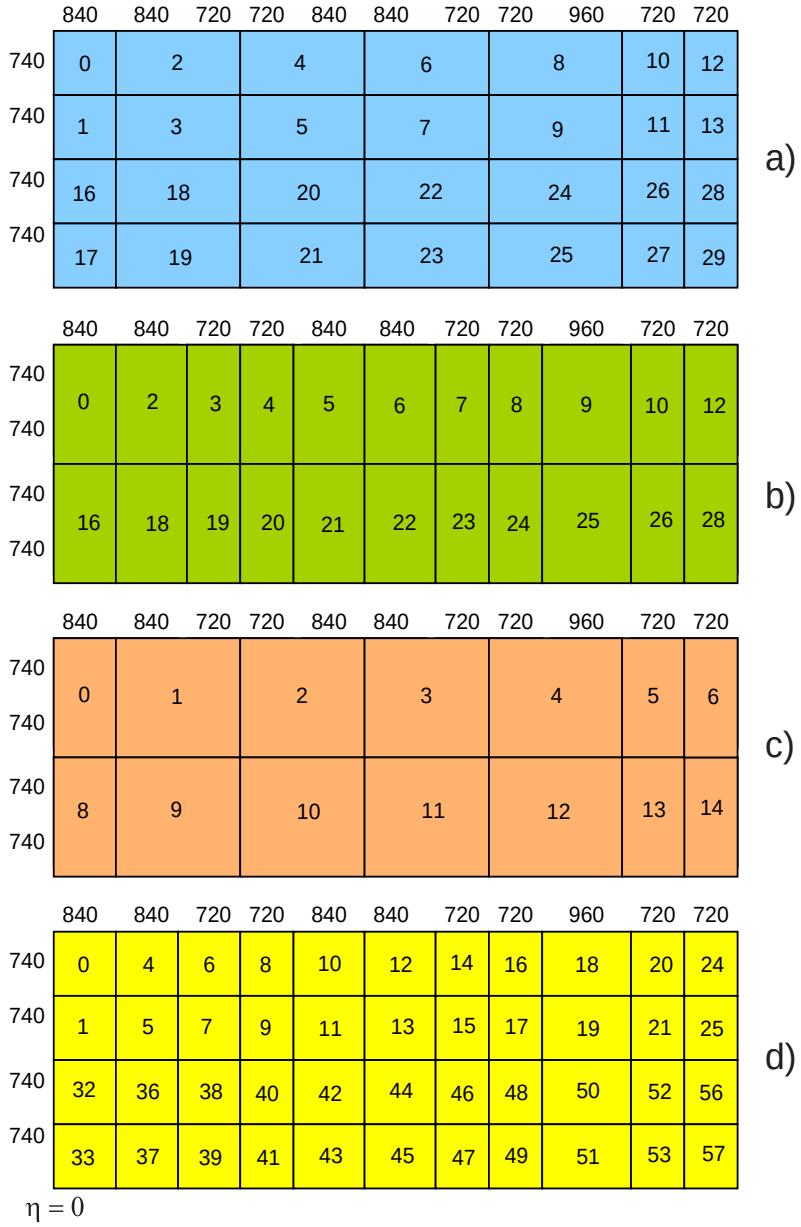

Figure 11-24 Partitioning of a Small Sector: (a) ϕ strip and CM partitioning, (b) η strip and CM partitioning, (c) Pad partitioning, and (d) RoI partitioning. Chamber dimensions in the RPC1 station are indicated (in mm).

Figure 11-27 shows the contents of the CM input and output patterns. For the low-p_T trigger, one input pattern is the 32-strip I0 plus the 32-strip I1 of the first RPC doublet. The second input is the 48-strip J0 plus the 48-strip J1 of the second RPC doublet. Low_out is the 32-bit pattern produced by the low-p_T CM plus eight more bits for local flags and thresholds information. For the high-p_T CM the first input pattern is the 32-strip I0H from the low-p_T CM output, while the second input pattern is the 48-strip J0H and the 48-strip J1H of the third RPC doublet. High_out is made of a 32-bit pattern plus eight bits of matrix information.

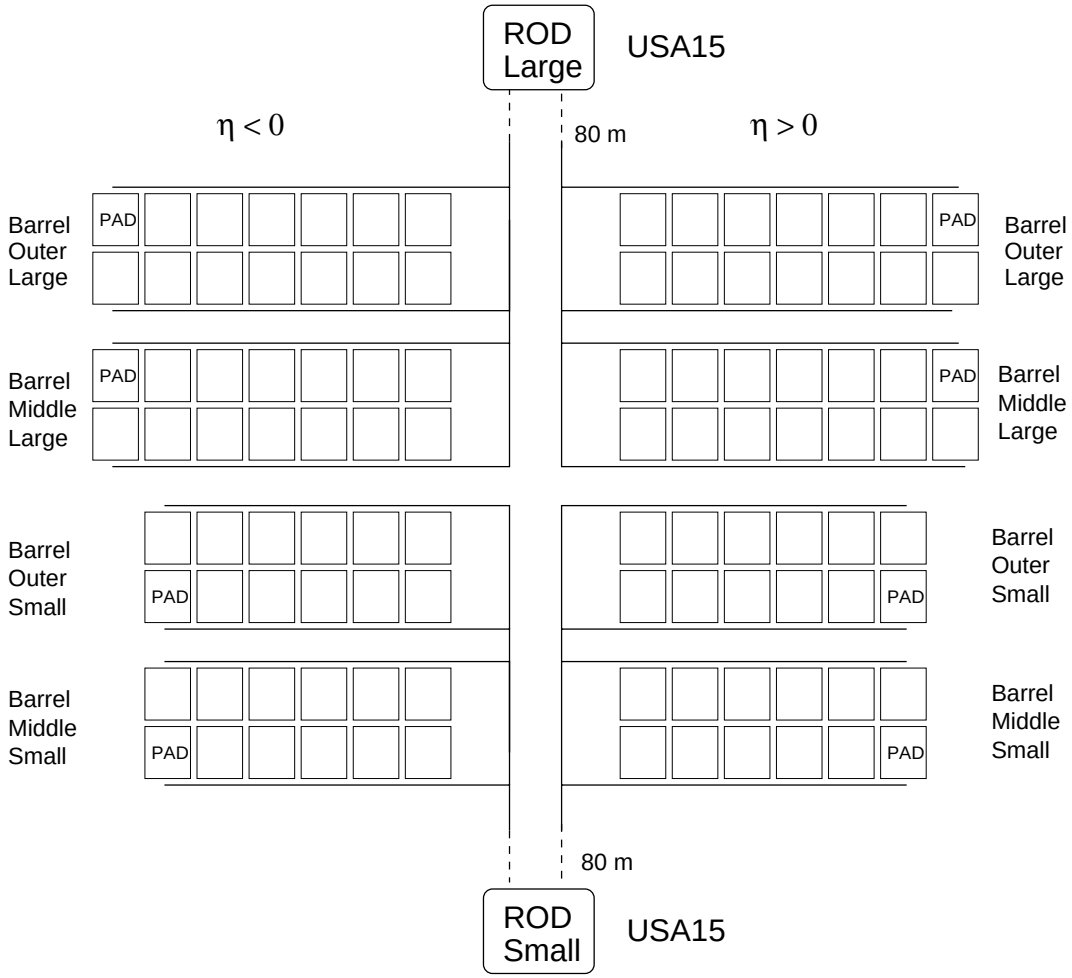

Figure 11-25 Read-out organization for RPC chambers.

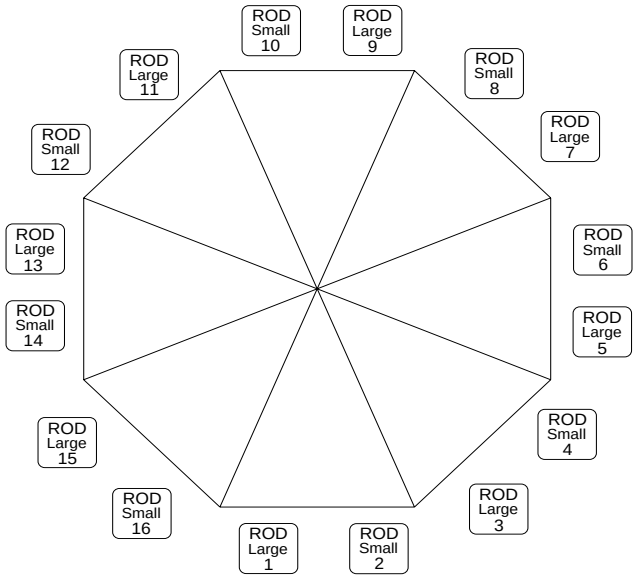

Figure 11-26 Read-Out Drivers of the RPC chambers. RPC ROD numbering follows the muon barrel layout numbering. RODs are physically located in USA15.

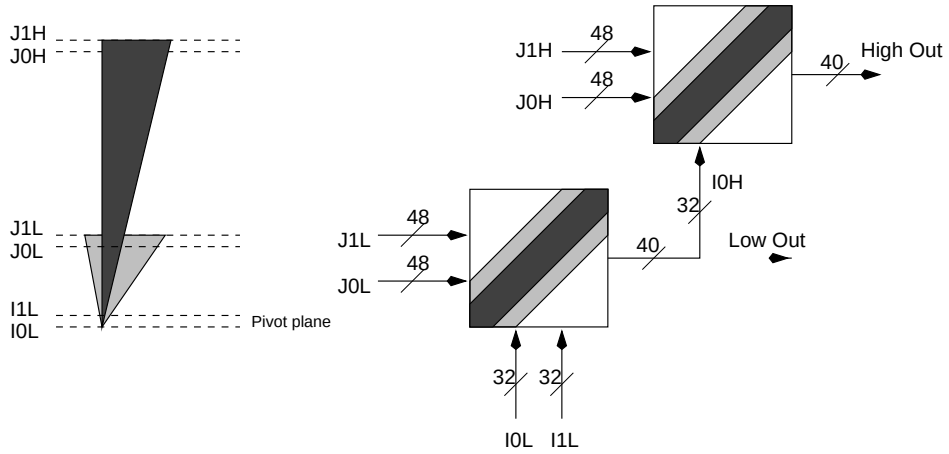

Figure 11-27 Coincidence-matrix input and output pattern.

For each CM matrix, on average 16 bytes are produced in the case of a low-p_T track candidate. Four clusters from the RPC1 and RPC2 doublets are expected. Assuming an average cluster size of 1.5 strips this accounts for six hits (12 bytes). The trigger output produces one hit (two bytes for the output pattern) and two bytes for the trigger threshold / overlap output.

Table 11-15 Read-out identifiers and numbering

ROD number	0–15	4 bit
SL number	0–63	6 bit
Pad number	0–55	6 bit
CM number	0–3	2 bit

In case of a high-p_T track 10 bytes have to be added to the previous low-p_T 16 bytes. These 10 bytes come from two clusters of the RPC3 doublets (six bytes accounting for the cluster size) and four bytes of the high-p_T trigger output.

In Table 11-15 the general read-out identifiers and numbering are shown.

As shown in Figure 11-28, each hit RPC strip will be encoded in two bytes, including the strip number and the time-interpolator measurement. In the same figure the overall event structure is shown, valid for both the CM-to-Pad link and the copper link going to the ROD. The header information is kept at minimum in order to minimize the data overhead in view of the low RPC occupancy.

Some extra information [Front-End-Bunch-Crossing Identifier (FEBCID, BCID), Front-End Level-1 Identifier (FEL1ID, L1ID)] is inserted at each level of the structure for synchronization checks. Structure identifiers, such as ROD number, SL number, etc. are generated in the hardware.

The CMs are not the only sources of read-out data. Each Pad produces data useful for monitoring the trigger system; Pad trigger patterns are sent via the copper link along with the CM data. Each SL produces, every bunch-crossing, a 32-bit trigger pattern whose content is illustrated in Table 11-4. Sixty-four such bit patterns are sent from the barrel trigger system to the MUCTPI through the fast synchronous 40 MHz trigger path. Each SL will also send read-out data, in order to be able to successfully monitor the trigger operation, through a dedicated copper link, to a corresponding SL ROD. In Figure 11-29 the ROD output data format is shown. It follows closely the required standard ATLAS event format for the I / O LDAQ modules.

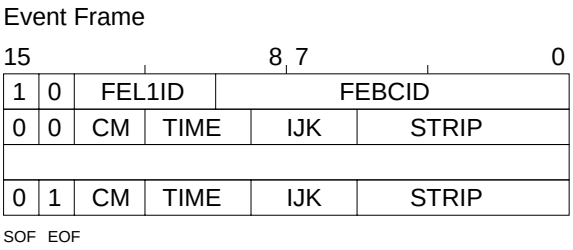

Figure 11-28 CM-to-Pad and copper-link data format.

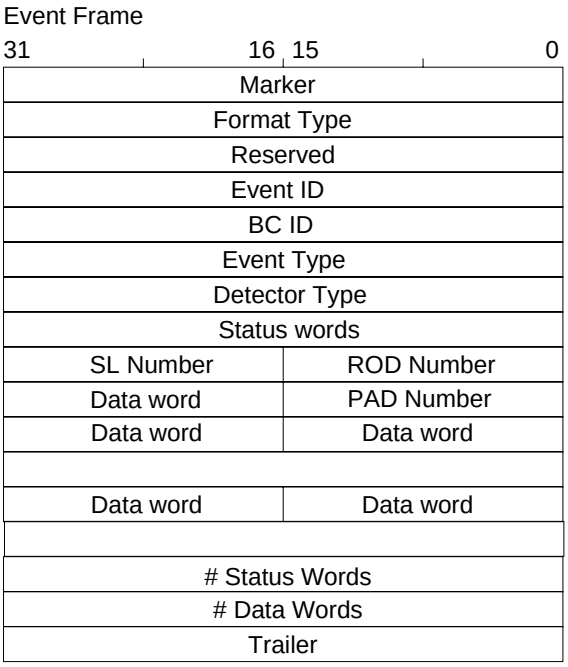

Figure 11-29 ROD data format.

11.2.1.1 Data-throughput calculation

A preliminary study on the data bandwidth requirements has been done using background calculations for the particle rates in the chamber. The input values used are shown in Table 11-16. The particle fluxes shown above give a total incoherent background flux of $10.3\,\mathrm{Hz/cm^2}$, and a coherent background of $0.7\,\mathrm{Hz/cm^2}$ (Incoherent background includes all the background sources that generate random clusters in the CMs, while the coherent background is defined as the one that generates correlated hits in the RPC doublets.)

Table 11-16 Background particle flux and RPC sensitivities used in the bandwidth calculation.

Particle type	Flux (kHz cm^{-2})	Sensitivity
e^+ e^-	2.3×10^{-3}	1
Charged hadrons	9.5×10^{-4}	1
μ	10^{-3}	1
γ	1.2	5.33×10^{-3}
n	3.3	10^{-4}

Table 11-17 shows a comparison of the bandwidth requirements. The calculations assume a safety factor of 10, to take into account uncertainties in the expected fluxes, a typical strip area of $360\,\mathrm{cm^2}$, an average cluster size of 1.5 hits, a trigger rate of 100 kHz and an encoding of hits in 16-bit words, in three cases:

- data encoding with no headers sent along the read-out path;
- data encoding with CM, PAD and ROD headers propagated along the read-out chain;
- no data encoding.

A read-out structure is assumed, where all components send data along their corresponding links in parallel, as foreseen in the final system. The final data-encoding scheme will require a data throughput which is between these two extreme cases. The data bandwidth that would be

Table 11-17 Average required bandwidths on CM, Pad and ROD links.

	Link bandwidth	Average bandwidth (16-bit encoding and no header)	Average bandwidth (16-bit encoding and full headers)	Average bandwidth (no encoding)
CM to Pad	10 Mbit / s	1.4 Mbit / s	3.0 Mbit / s	480 Mbit / s
Pad to ROD	20 Mbit / s	5.7 Mbit / s	13.7 Mbit / s	1920 Mbit / s
ROD to ROB	1 Gbit / s	0.30 Gbit / s	0.73 Gbit / s	100 Gbit / s

required, without hit encoding, is also shown. (24 time slices, corresponding to three consecutive bunch crossings, have to be transferred per strip each event.)

The current design seems adequate, but an accurate simulation of the system is necessary and will be made.

11.2.2 Control and monitoring

All control and monitoring functions of the trigger system are performed through five hardware parts: the read-out system, the Timing, Trigger and Control system (TTC), the I2C interface, the JTAG interface and the Detector Control System (DCS).

The read-out system is used for monitoring and checking data during data-taking runs, calibration runs and test and diagnosis operations. Data generated by the trigger system during data taking runs can be monitored by the run control system. Monitoring at the ROD level will be needed to allow independent running from DAQ. The main monitoring and control functions to be performed are outlined below:

- Input data quality: Hit maps are needed to check the working conditions of the trigger chambers.
- Online checks of the trigger algorithm are done using raw data and trigger data, comparing the results from processing in the LVL1 hardware with those from a software model.
- Monitoring of trigger rates, histogram filling, per sector.
- Monitoring of nearly full flags from the derandomizers has to be done to react and introduce dead time.
- Monitoring of ROB-BUSY signals, to force the level-1 CTP to introduce dead time when necessary, is needed.
- A test mode, where fake events are produced with known data, is foreseen to be able to test the read-out chain from the ROD or possibly from the derandomizers up to the ROB.
- Error-recovery procedures are foreseen, since the DAQ must not stop if a single failure is present in the read-out chain of a subdetector.

Calibration runs are used for timing calibration and momentum calibration. In both cases the appropriate set up of the trigger system parameters is made through the TTC, I2C and JTAG systems, and the data are read through the read-out system.

The TTC system performs the important functions of distributing the machine clock and level-1 accept signals and of setting up the right timing of the trigger system for different operations. A standard receiver chip, the TTCrx, is used on the detector to receive and decode the TTC signals. This chip is programmed through the use of the I2C standard fieldbus.

The functions of loading the appropriate parameters for different kinds of operations, and for test and diagnosis of the on-detector electronics, are performed through the JTAG or the I2C interface. These interfaces are part of all on-detector boards and are connected to the control system that is located in underground counting room. Since, in the present design, every on-detector electronics board is accessed by three buses (I2C, JTAG and DCS/CAN), to reduce the number of buses, the possibility of using the I2C bus and discarding the JTAG bus is being investigated.

The DCS system is used to monitor the voltage and the temperature of each board of the on-detector electronics and, if necessary, to switch off the board. It makes use of the standard CAN fieldbus system.

11.2.3 System initialization

The initialization procedure is accomplished by loading the on-detector electronics, the Sector Logic and the RODs with the appropriate parameters. The hardware used for the initialization are the TTC and the I2C and JTAG interfaces for the on-detector electronics and the VME control system for the Sector Logic and the ROD. Calibration data should be taken in order to properly initialize some of the trigger parameters. The initialization includes the following phases:

- Set the on-detector TTC circuits and the programmable delays between the TTC system clock and the coincidence-matrix, the Pad Logic and front-end boards.
- Set chamber front-end electronics thresholds.
- Set the mask-off pattern for channels that have been found noisy during previous monitoring of chambers.
- Initialize multiple trigger coincidence-matrix roads, to set up to six different p_{T} thresholds.
- Initialize coincidence-matrix input pipeline depths, one value per front-end signal cable (used to time align all input channels).
- Initialize coincidence-matrices mode registers. The mode registers are used to define the coincidence-matrix chip functionalities: low- or high-p_{T} algorithm data processing, run initialize mode, etc.
- Initialize the coincidence-matrix output pipeline depths (used to synchronize the various coincidence-matrices belonging to the same Sector Logic).
- Program digital shaping of coincidence-matrix chip inputs and outputs (used to optimize bunch-crossing identification efficiency and background rejection).
- Initialize Sector Logic input pipeline depths, to synchronize the Pad Logic belonging to the same sector.
- Initialize the Sector Logic look-up tables (used during local merging of muon candidates found in the RoIs defined in the bending and non-bending plane).

- Initialize overlap window sizes between adjacent sectors and between barrel and end-cap regions. (They are optimized to balance rejection to false dimuon triggers and to maximize dimuon trigger efficiency.)
- Set the mask for masking-off hot RoIs.

11.2.4 System calibration

Two kinds of calibration are foreseen: timing calibration and momentum calibration. The timing calibration establishes the correct set up of the timing system for beam data-taking runs, cosmic-ray runs, and test and calibration runs. The momentum calibration determines the correct set up of the of the p_T thresholds of the trigger system.

11.2.4.1 Timing calibration

Strategy for setting up timing with beam

In the following timing calibration strategies a coarse time alignment, using TTC and RPC cable propagation delay measurements, time-of-flight calculations and detector response is performed before starting the timing procedure.

Two different strategies are required, one for the local synchronization of front-end signals coming from different stations within a trigger tower (Pad), the other for the global synchronization of the trigger system.

Local time calibration: The front-end signals from the three stations enter the CMs with different delays due to time-of-flight, propagation along strips and cable length. Each CM can realign the input data using the programmable-depth input pipelines, which act independently per octet of discriminator signals.

Before calibration the input pipelines will be set to zero delay. In Figure 11-30, an example of the timing of signals before calibration is shown.

The CM processing is driven by a 320MHz clock, derived from the 40MHz BC clock coming from the TTCrx chip mounted on each Pad. In this way each bunch-crossing period is subdivided into eight time slices, and both front-end and trigger outputs are assigned to a time slice via the 3-bit time interpolator circuit. This facility will be used for monitoring both the detector and the trigger system.

The RPC front-end discriminator has a fixed shaping time of around 10ns. An edge detector will sample this signal at the input of the processor pipeline.

The time-slice information will be sent into the level-1 read-out pipeline, while a shaped output will go to the coincidence-matrix. Before calibration the shaping time of the input signals to the CM will be set to the maximum, eight time slices, to be able to identify muon candidates effectively even if the timing is not precisely set.

The same calibration procedure is applied to the two I and J doublets.

The matrix output will assign the muon candidate to a specific bunch crossing, but the detector and trigger data belonging to a muon track will be spread on a number of bunch crossings.

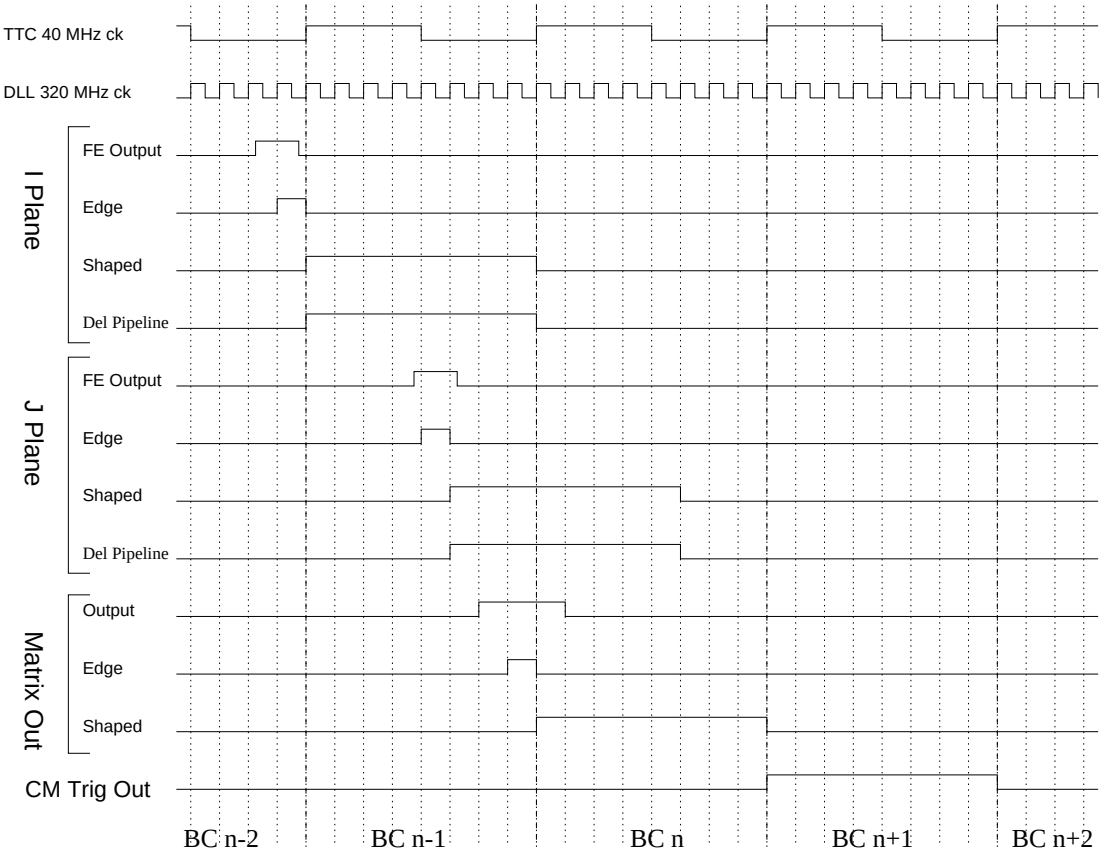

Figure 11-30 Timing of local processor signals before calibration.

For this reason the read-out part of the CM is able to send data belonging to up to five bunch crossings around the triggered one.

It will be possible to run uncalibrated with a single-muon level-1 trigger at the expense of a lower rejection to fake tracks and higher required read-out bandwidth.

An analysis of the data collected with this trigger will allow suitable delays to be determined that can then be programmed into the variable-depth input pipelines of the CMs, establishing temporal alignment of the different inputs. An example of the timing of signals after local calibration is shown in Figure 11-31.

Once the input delays have been correctly set up, the shaping time of the I and J signals can be shortened, reducing the rate of fake triggers from accidental coincidences. The shaping times will be set to the minimum values that give efficient track finding, taking into account the spread of signal arrival times due to variable time-of-flight, the detector resolution of 1.5 ns, the range of signal propagation times, the discriminator walk and the quantization error.

Final adjustment to be made are to adjust the clock phase from the TTCRx to its optimal value and to set the shaping of the CM output to match the 40 MHz external operation. This fine tuning is done by monitoring the CM trigger output, time tagged with a 3.1 ns least significant bit precision. The TTCRx phase must be adjusted to centre the time distribution inside the 25 ns bunch-crossing interval.

Since hits belonging to a muon track will belong to a combination of four out of eight CMs of a low- and high-p_T Pad, this operation will align in time locally the low- and high-p_T parts, and the bending and non-bending coordinates.

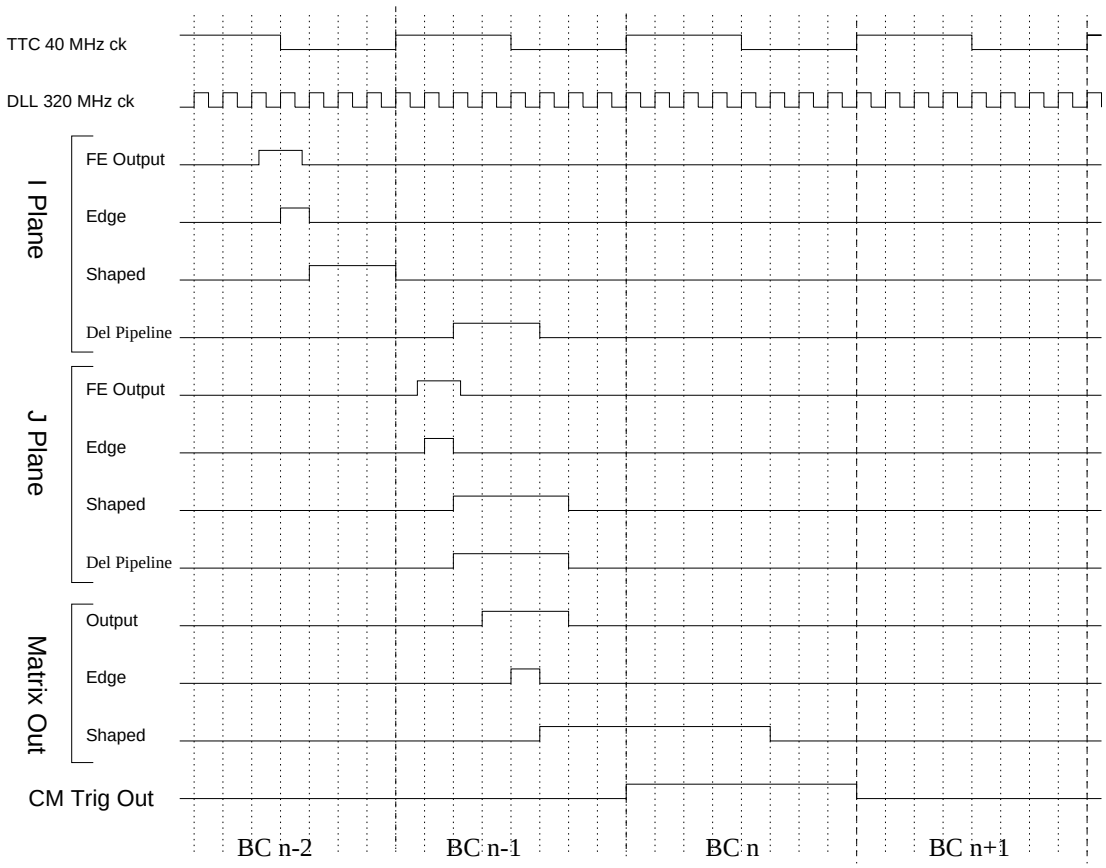

Figure 11-31 Timing of local processor signals after calibration.

This scheme has a powerful monitoring capability, because the CM inputs and outputs are sent to the read-out, time measured with a precision of 3.1ns least significant bit, which is much finer than the bunch-crossing interval of 25ns.

The monitoring needs of the detector system are also met.

Global time calibration: After the local time-calibration procedure described above, the different inputs and circuits within each pad will be correctly timed-in relative to each other. However, the timing calibration between the different pads, sectors and the rest of the LVL1 system still have to be adjusted.

The global synchronization procedure will first look for two muon candidates, triggering on single muons, and searching for a secondary muon trigger in the five BCs range around the triggered BC. The secondary muon can come from the same track that has been seen twice because of bad synchronization among different parts of the apparatus (low- and high-p_T CMs, overlap between Pads, overlap between Sectors, overlap between barrel and end-cap) or from an additional real muon belonging to the same event. The fake secondary muon can be easily recognized.

The relative phases of the TTCrx are then adjusted acting on the phase of the 40 MHz clock, until all muon tracks will assigned to the same BCID, either collapsing on a single muon track or generating a dimuon trigger.

Strategy for setting up timing with cosmics

Rate calculations are necessary to understand which fraction of the spectrometer can be calibrated with cosmics. The $\cos^2\theta/E^2$ dependence and the effect of the access pits have to be studied. Timing settings will be different for beam and for cosmics and a different calibration scheme has to be foreseen at least for the upper part of the experiment.

Local calibration with cosmics: In the case of tracks coming from the interaction region the propagation of signals along cables follows the particle flight, while for cosmic muons signal propagation and particle flight, in the upper part of the apparatus, are in opposite directions.

It will still be possible to align signals in time to make local coincidences, but the requirements on the CM design will be increased in one or both of the following ways:

- wider range of input delays (to allow for reversing the timing in the upper part of the apparatus);
- longer shaping for coarse coincidences before calibration.

The effect on the maximum time walk of the input signals due to different cable paths and lengths has to be carefully studied before freezing the CM design. The final layout values are being calculated at the time of writing.

Global calibration with cosmics: A cosmic trigger will look for muons pointing to the interaction region and, after following the same rules for the global calibration as with beam, a cosmic muon will fire the di-muon trigger at level-1 (Figure 11-32).

The lower part of the experiment uses the beam-run timing settings, while the upper arm is offset to earlier bunch crossings.

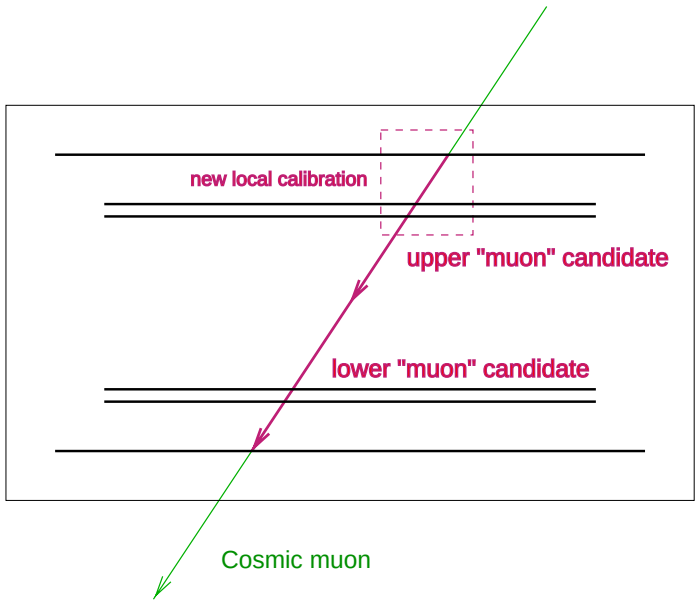

Figure 11-32 Level-1 cosmic-muon trigger.

Strategy for setting up timing in test and calibration runs

It is possible to make an initial set-up of the timing without using beam or cosmic runs. This requires a special command 'prepulse' (PPS) to be distributed by the TTC system, followed by an Level-1 Accept signal after a fixed time (see Sections 15.2.3.3 and 16.5.3).

The PPS command is received and decoded by the TTCrx circuit on the Pad board. The decoded command is transmitted to the four coincidence-matrix boards and to all front-end boards connected to the Pad. On the PPS command, the front-end boards set to 'one' the outputs of all the strips belonging to the board. The outputs of the front-end boards are connected to the inputs of the CM boards. Here the input bits of the CM chip can be masked individually to generate a trigger on a predefined pattern. The generated trigger is read-out by the Level-1 Accept signal that will follow the PPS command.

The timing settings are the same as for the beam data-taking procedure, already illustrated in the previous paragraphs. The only difference is in the time of flight that in this case does not need to be compensated since it is the same for all the trigger system.

The described procedure will also be used to test the full read-out chain, since it allows all the strips of the apparatus to be read simultaneously.

11.2.4.2 Momentum calibration

Momentum calibration is necessary for setting the p_T thresholds of the trigger system. The system must operate with three low-p_T and three high-p_T thresholds working concurrently.

At the beginning of data-taking, for each part of the apparatus, the setting of the thresholds is made by loading in the coincidence-matrix chips roads generated from the simulation programs. These roads should already be very close to the final ones. After a period of data-taking (the duration of this period has to be studied), when there is enough statistics, roads can be checked and if necessary updated by studying reconstructed muon tracks (p_T measured in the muon spectrometer and in the inner detector).

11.3 Prototype demonstrator programme

A long and intensive demonstrator programme was carried out in the last few years to demonstrate the feasibility of the barrel trigger system design [11-17]–[11-21]. The programme was dedicated to detector studies and associated electronics. The detector studies are described in Ref. [11-8]; here we report on the electronics studies for designing the level-1 muon barrel trigger processor. This part of the studies was carried out in the framework of the RD27 project [11-21].

The main goal of the demonstrator programme was to design and build, on a small scale, a trigger system that meets all the most important trigger requirements and that is conceptually very close to the final ATLAS trigger system.

The programme was mainly developed in three different phases and each phase was concluded with a series of tests and measurements, performed on a test-beam in the CERN North Area.

In the first phase a system was developed just to test the planned trigger algorithm, to implement and to demonstrate the possibility to identify the LHC machine bunch-crossing. Here the trigger processor was based on coincidence-matrix boards, connected to the RPC front-end and implemented through a series of commercial GaAs crossbar switches [11-22]. The results were very encouraging: the whole trigger system, made by the RPC and the processor electronics, performed very well the trigger algorithm and gave an excellent bunch-crossing identification capability, since the global time resolution, of the detector and of the electronics, was < 2 ns (excluding propagation delays on strips).

Based on the results of the first tests, something much closer to the final ATLAS trigger system was built, from both a detector and an electronics point of view. First a small-size trigger tower and then a full-size trigger tower were built. Both were equipped with an electronic processor based on a dedicated CM demonstrator ASIC. The systems were tested on the ATLAS H8 test-beam line.

11.3.1 The coincidence-matrix demonstrator chip

The most important building block in the design of the barrel trigger demonstrator system was the coincidence-matrix demonstrator ASIC [11-23], [11-24]. This ASIC contains, in a reduced scale, the most important functionalities that must be implemented in the level-1 barrel trigger.

The ASIC can work in fully combinatorial and fully pipelined operation or in any combination of the two. The dimension of the matrix is 8×24 and two thresholds can operate simultaneously. The chip can be programmed to implement the low- or the high-p_T algorithm. In the case of the low-p_T algorithm, a 3/4 majority coincidence of the four inputs coming from the two RPC doublets is performed. In the case of the high-p_T algorithm, a 2/3 majority coincidence of a high-p_T RPC triplet, plus the low-p_T output validation is performed. (In the previous ATLAS trigger design the high-p_T RPC station was made of a triplet instead of a doublet.) The ASIC contains the capability of adjusting the input cable delays, through a variable-length programmable input pipeline of four step length. Also, the possibility of masking noisy channels or missing planes is implemented. Figure 11-33 shows the block diagram of the ASIC.

The ASIC was implemented in a Fujitsu 0.5 micron CMOS 34K gate-array. From the tests performed on it, we have measured a maximum skew in the chip operation of 1.8 ns and a maximum pipeline working frequency of 120 MHz.

A number of 6U VME boards were made based on the ASIC (see Figure 11-34), to be used as trigger-processor elements.

The main element of each board is the coincidence-matrix ASIC, to which are connected the input signals of the RPC and the output signals of the trigger result. At the input, the signals are converted from ECL to TTL levels suitable for the 3.3V power supply of the CM. Additional input/output connections are the XIN and YIN serial inputs, and the THR0 and THR1 NIM trigger threshold outputs. A Xilinx FPGA circuit is used for the ASIC control and the VME interface.

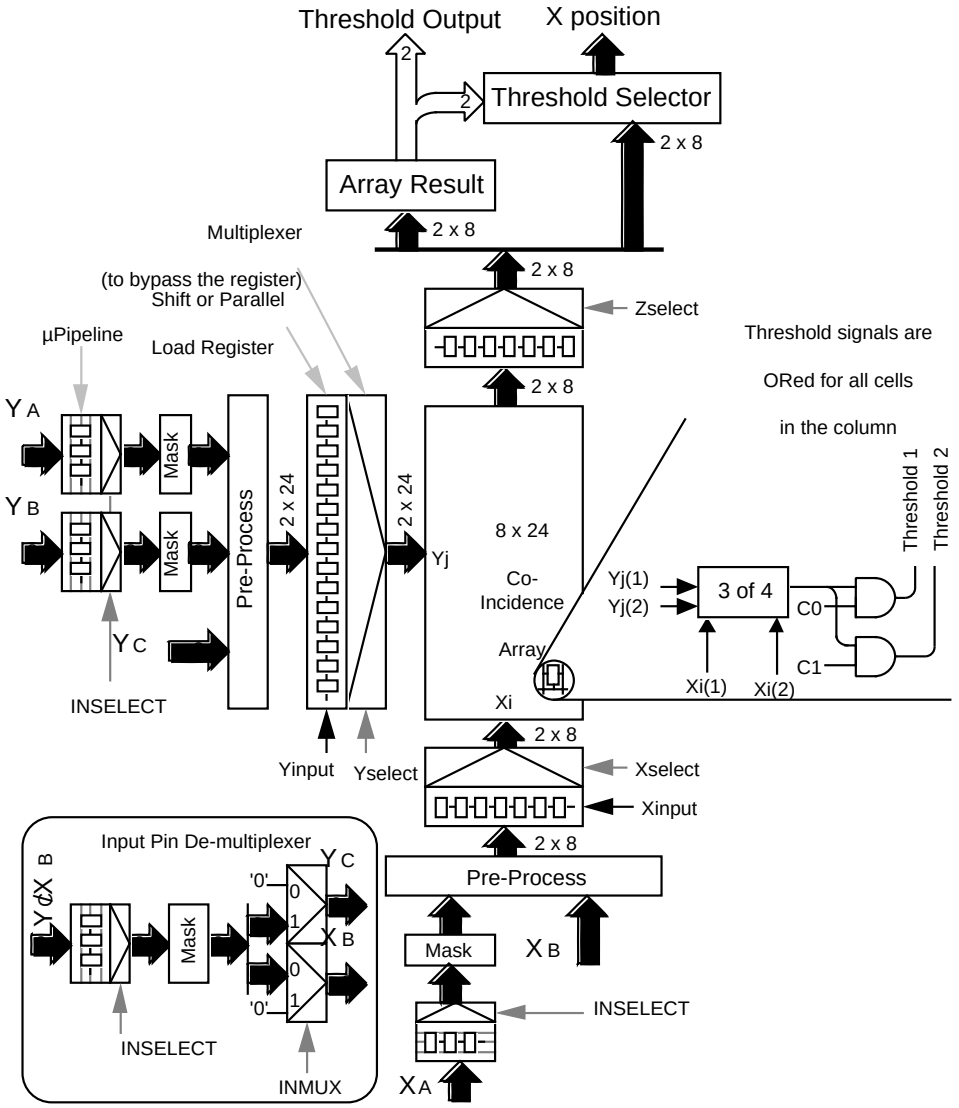

Figure 11-33 Block diagram of the coincidence-matrix demonstrator ASIC.

11.3.2 Trigger-tower demonstrator systems

Two trigger-tower demonstrator systems were built and tested in 1996 in the ATLAS H8 test-beam line [11-25]. The first demonstrator was a reduced-size trigger tower of $50 \times 50\text{cm}^2$ (Figure 11-35), composed of two RPC doublets with strips in both projections for the low-p_T trigger, and one RPC triplet with strips in both projections for the high-p_T trigger.

From the processor logic point of view, a 2×2 CM low-p_T Pad and a 2×2 CM high-p_T Pad were simulated. The whole system was equipped with eight CM VME boards, four in the η projection and four in the ϕ projection, each housing a CM demonstrator ASIC, for a total of 192 electronics channels.

A second full-size trigger tower ($90 \times 270\text{cm}^2$), made of two low-p_T RPC doublets and of one high-p_T RPC doublet, with strips in both projections, was constructed and tested in the same beam line.

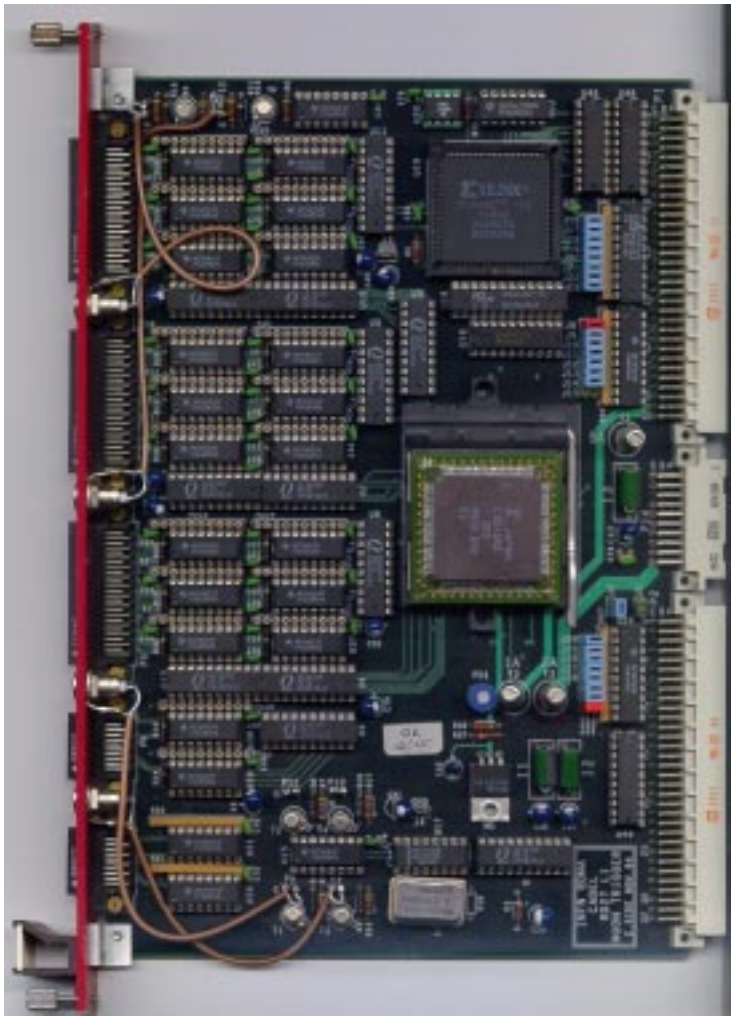

Figure 11-34 Coincidence-matrix board.

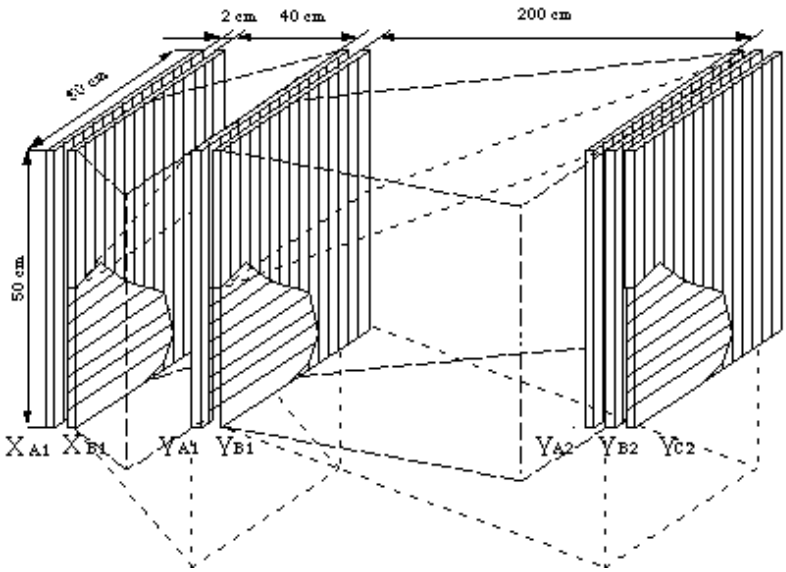

Figure 11-35 Reduced-size trigger tower.

Studies were performed for tracking, time resolution and efficiency. Data were taken with muons of 180 GeV, at 10, 50, 500 and 900 Hz/cm^2 beam fluxes. The LHC gamma background conditions were simulated through a 14 mCi ^{60}Co source, giving a 100 Hz/cm^2 background flux on the closest RPC [11-26], [11-27]. Results of the detector studies are documented in [11-8], Section 8.5.1. Concerning the trigger-processor studies, the results on tracking capability, trigger threshold behaviour and time resolution are in good agreement with the system design specifications. In particular, the excellent time performance of the system can be seen from Figure 11-36, where the processing time jitter of the level-1 trigger processor for low-p_T (dashed histogram) and high-p_T (full histogram) tracks is shown, for both X and Y coordinates. The average low-p_T processing time is subtracted so that the distribution is centred at t=0.

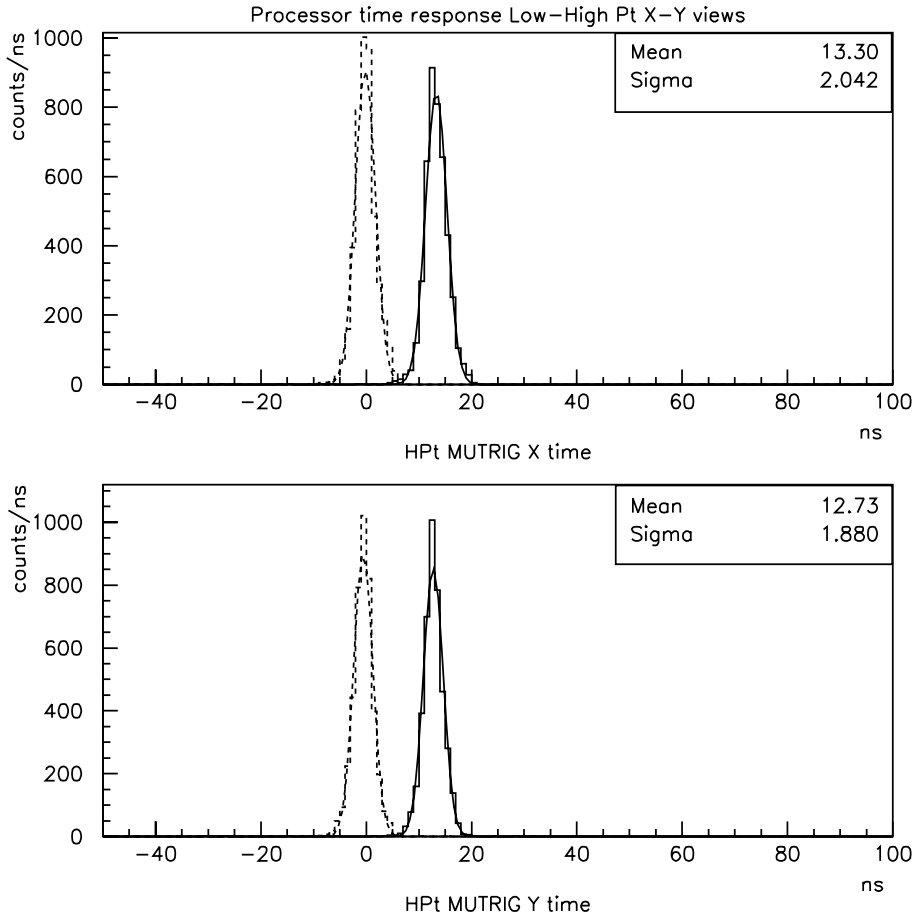

Figure 11-36 Processor time response for low- and high-p_T in both views.

11.4 Trigger system construction and assembly

The level-1 muon barrel trigger schedule has been worked out taking into account the necessity to follow the general schedule for construction, assembly, installation and integration of the ATLAS experiment. For each relevant component of the trigger system, a detailed schedule of the different phases and of the intermediate milestones was established on a quarterly basis (Figure 11-37). Following the schedule, the system should be ready to be pre-assembled on the chambers in the middle of 2002, allowing also for nine months of contingency.

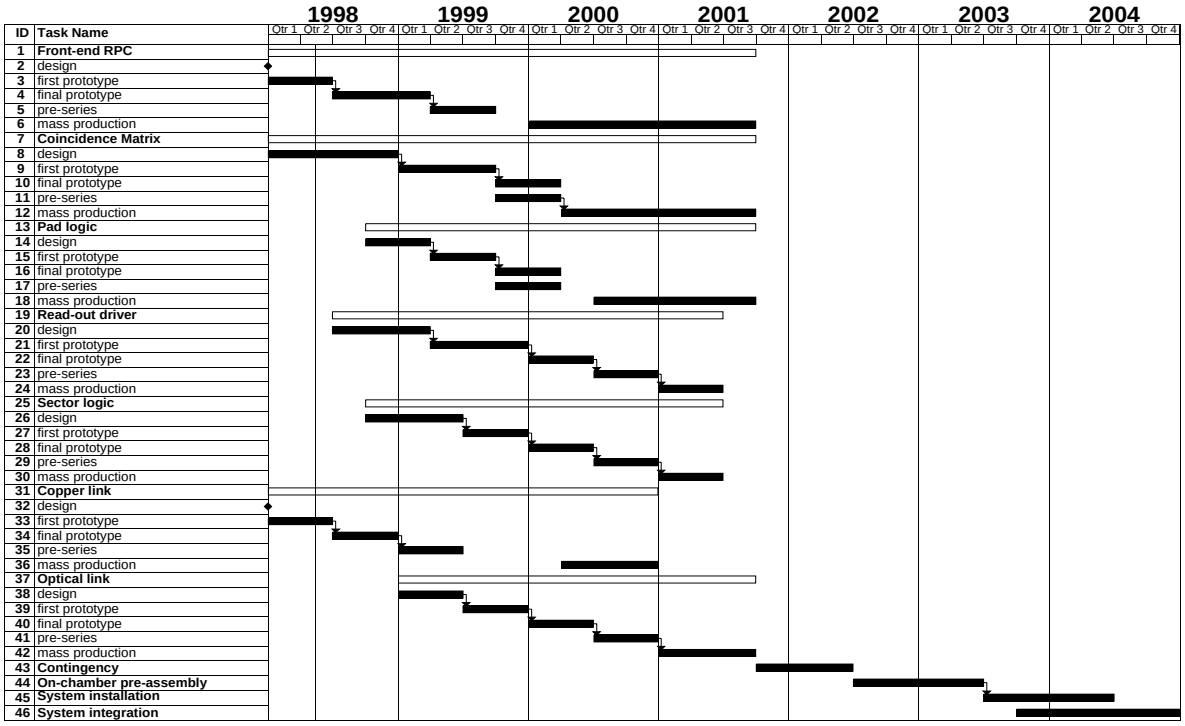

Figure 11-37 Level-1 muon barrel trigger schedule (The front-end electronics is part of the muon spectrometer.)

11.5 Quality assurance programme

The quality assurance programme is a relevant aspect in the implementation of our trigger system, since the system must work at least for 10 years in a very difficult environment, where the accessibility is very limited or in some cases even impossible. This means that a set of procedures and rules must be followed in the implementation of the trigger system to assure the required quality. The programme is summarized in Table 11-18.

As a first and most relevant aspect, it must be ensured that for the electronic system the requirements meet the physics goals. In this regard, we have produced a document, the Level-1 Muon Trigger User Requirements Document (URD), that is the basis for the design of the barrel trigger system.

Table 11-18 Quality Assurance programme summary

Requirements	URD document
Design	Preliminary Design Review followed by Interim Design Reviews if required
Development	Prototyping, iterate design, followed by Final Design Review
Production	Production Readiness Review (PRR)
Installation	Installation procedures and schedule set by ATLAS experiment
Quality and maintenance	Electronics Reliability Rules

The realization of the trigger system, that must be robust and adequately tested, should follow a set of rules (Table 11-18) that cover all aspects of the implementation chain.

Furthermore, we plan to use the following Electronics Reliability Rules:

- evaluation of allowable aggregate failure rate;
- evaluation of allowable replacement frequency for accessible components;
- suggested goal of <1% aggregate channel loss for a 10 year running cycle;
- in the design of ASICs, obtain from vendors data on defects / unit area for a given run process and use software-reliability tools to obtain the reliability of components;
- test all on-detector circuits to validate that they satisfy radiation-tolerance requirements;
- obtain for commercial components data from vendors and use Mean Time Before Failure software tools to evaluate the reliability of circuits;
- burn-in, at least for inaccessible components;
- ensure full availability of spare components and parts for a 10 year running cycle;
- define rules for access to fix problems and repair circuits in short- and long-term maintenance.

Before launching the construction of the major parts of the experiment, ATLAS has established to organize a Production Readiness Review (PRR). This review will cover design, integration, logistics, quality and safety issues, other than organizational matters. The items of the level-1 muon barrel trigger reviewed by a PRR and the corresponding dates at which the reviews will take place are shown in Table 11-19.

Table 11-19 Production Readiness Review schedule

Coincidence-matrix	April 2000
Pad Logic	July 2000
ROD and Sector Logic	February 2001
Optical link	January 2001

11.6 References

11-1 C.Bacci et al., *Test of a resistive plate chamber operating with low gas amplification at high intensity beams*, Nucl. Instrum. Methods A344 (1994) 194–198.

11-2 F.Ceradini et al., *A fast tracking muon trigger for high luminosity colliders using resistive plate chambers*, Large Hadron Collider Workshop, CERN / 90–10, ECFA / 90–133, eds. G.Jarlskog and D. Rhein, vol. III, p. 99.

11-3 C.Bacci et al., *Muon trigger and identification at the LHC*, Proc. Large Hadron Collider Workshop, CERN / 90–10, ECFA / 90-133, eds. G.Jarlskog and D.Rhein, vol.III, p. 467.

11-4 M.Della Negra et al., *Study of muon triggers and momentum reconstruction in a strong magnetic field for a muon detector at LHC*, CERN / DRDC / 90–36, DRDC / P7, 1990.

11-5 Gy.L.Bencze et al., *Study of resistive plate chambers for muon detection at hadron colliders*, Nucl. Instrum. Methods A340 (1994) 466–473.

11-6 S. Veneziano et al., *The first level muon trigger in the central toroid of the ATLAS experiment*, Vienna Wire Chamber Conference 1994, Nucl. Instrum. Methods A367 (1995) 428–431.

11-7 E. Petrolo et al., *The First Level Muon Trigger of ATLAS in the Barrel Region*, Third Workshop on Electronics for LHC, London, UK, September 1997.

11-8 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

11-9 *LVL1 Muon Trigger User Requirements Document (Draft version 1.4),* ATLAS working document, ATL–DA–ES–0002,

11-10 P. Farthouat, *Current Understanding of the Muon Level-1 System*, version 2, ATLAS note DAQ–Note–89, May 1998.

11-11 B*arrel LVL1 Muon Trigger Coincidence Matrix ASIC: (User) Requirement Document*, version 0.1, ATLAS working document, November 1997.
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/lvl1mu_bckTDR.html

11-12 *Timing, Trigger and Control System (TTC) (User) Requirements Document (Draft version 0.2),* ATLAS working document, ATL–DA–ES–0004, February 1998.

11-13 *Study of the Radiation Tolerance of ICs for LHC*, LHCC / CERN / 97–63, LEB Status Report / RD49, December 1997.

11-14 E. Barberio et al., *Implementation of the first level muon trigger*, ATLAS note DAQ–NO–39. 1994.

11-15 R.Cardarelli et al., *Trigger logic of the ATLAS first level muon trigger*, CERN / LHCC 95–56 First Workshop on Electronics for LHC Experiments, Lisbon, Portugal, September 1995.

11-16 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0),* ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

11-17 N. Ellis et al., *First Level Trigger Systems for LHC Experiments*, CERN / DRDC / 92–17, DRDC / P37, 1992.

11-18 I. Brawn et al., *First Level Trigger Systems for LHC Experiments,* CERN / DRDC / 93–32, 1993.

11-19 N. Ellis et al., *First-level trigger systems for LHC experiments*, CERN / DRDC / 92–46, DRDC / M20, 1992.

11-20 N. Ellis et al., *First Level Trigger Systems for LHC Experiments,* CERN / DRDC / 94–52, 1994, 36 pp.

11-21 N. Ellis et al., CERN / LHCC 96–4, LERB Status Report / RD27, 1995.

11-22 E. Petrolo and S. Veneziano, *Use of GaAs circuits for first level muon triggering at LHC*, 5th Topical Seminar on Apparatus for High Energy Physics and Astrophysics, S. Miniato Conference, 26 April 1993, Nucl. Instrum. Methods A344 (1994) 194–198.

11-23 J. Dowdell et al., *A coincidence array demonstrator ASIC for the RD27 muon trigger*, RD27 Note 30, 1994.

11-24 R.Bindra et al., *The RD27 Muon Trigger Co-incidence Array Demonstrator ASIC*, 1995 IEEE Nuclear Science Symposium, IEEE Trans. Nucl. Sci. 43(3) (1996).

11-25 S. Veneziano, *A test of the ATLAS first level muon trigger logic*, Proc. Third International Workshop on the Resistive Plate Chambers and Related Detectors – RPC '95 Pavia, Italy, 11–12 October 1995, Publ. in: eds. S. P. Ratti and M. Merlo, INFN, Pavia, 1996 Scientifica Acta, v.11, no.1, 1996 (177–186). 1996

11-26 M.Angelone et al., *Test of a resistive plate chamber under irradiation of photons and neutrons*, Nuclear Instruments and Methods in Physics Research A355 (1995) 399–405.

11-27 M.Abbrescia et al., *Resistive plate chambers working at high rate*, Proc. Fourth International Conference on Advanced Technology and Particle Physics – ICATPP 4 Como, Italy; 3–7 Oct 1994. Publ. in: eds. E. Borchi, S. Majewski, J. Huston, A. Penzo and P.G. Rancoita, Nucl. Phys. B, Proc. Suppl. 44 (1995) (218–221) 1995.

12 Muon trigger for the end-cap

12.1 Trigger chamber layout

Thin Gap Chambers (TGCs) provide the muon trigger in the end-cap region. The TGCs are arranged in seven layers in each end-cap at $|z| \sim 14$m (Figure 12-1). See also Chapter 10. They are grouped in three planes in z, one of triplet units and two of doublet units. The doublet forming the plane farthest from the interaction point in each end-cap is referred to as the pivot plane, and its chamber layout and electronics are arranged such that, to a good approximation, there are no overlaps or holes in this plane. For triggering, the TGCs cover a pseudorapidity range $1.05 < |\eta| < 2.4$.

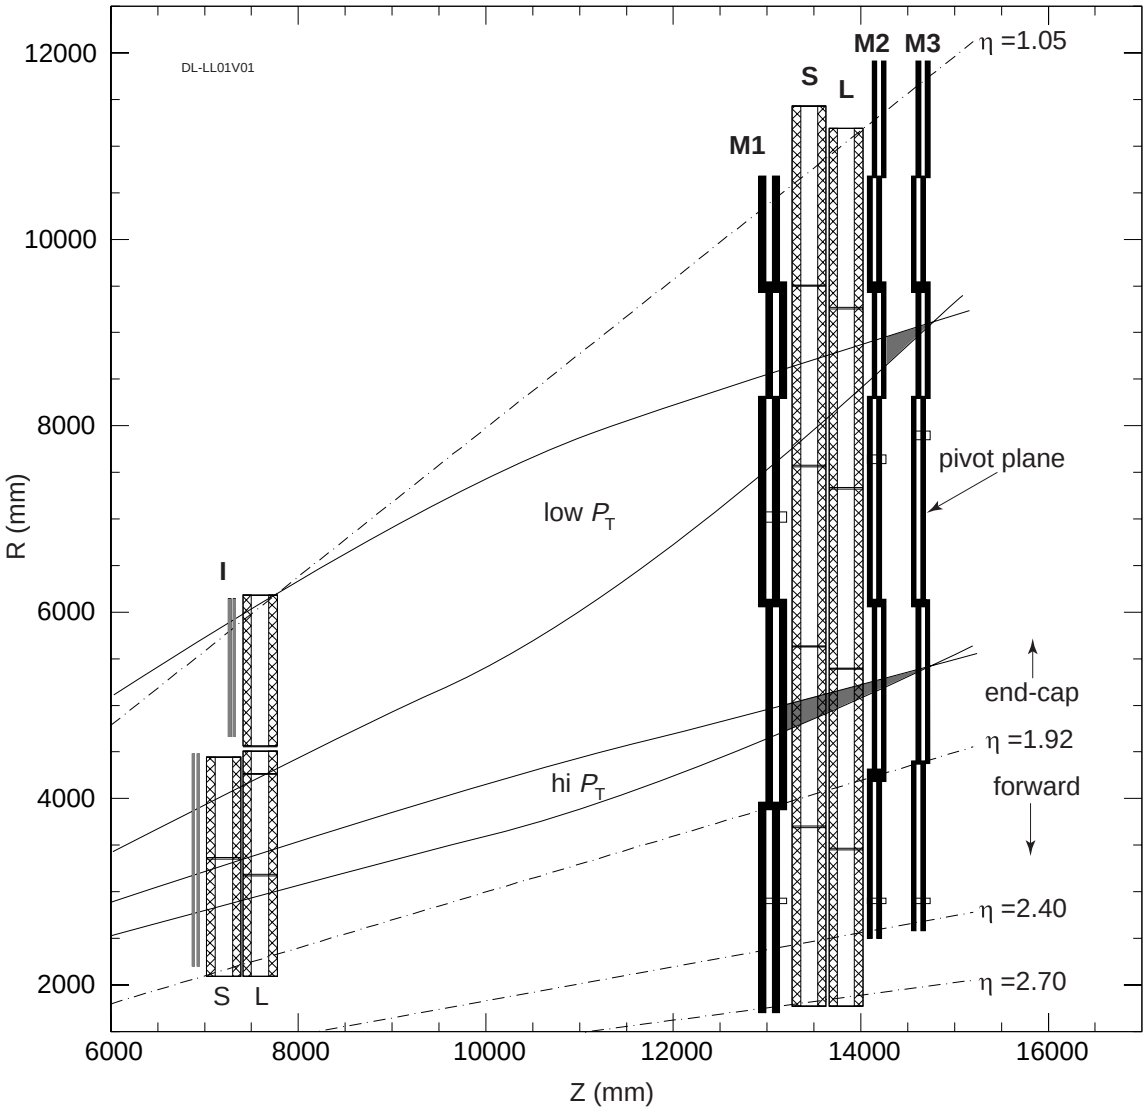

Figure 12-1 Longitudinal view of the TGC system, showing trajectories of high- and low-p_T muons and their trigger windows. M1 is the TGC triplet, M2, M3 are the TGC doublets, S and L are the small and large MDTs, I is the inner-station TGC (not used in the trigger). To allow overlapping of the physical chambers, doublets and triplets at adjacent ϕ are at slightly different z; the bold lines show a doublet/triplet at two adjacent ϕ's.

Each trigger plane of TGCs consists of a ‘wheel’ of eight octants of chambers symmetric in ϕ. Each octant is divided radially into the ‘Forward region’ and the ‘End-cap region’, Figure 12-1. The Forward Region layers each contain three sets of units, where a set consists of one doublet or triplet unit. The End-cap Region layers also each contain three sets of units, where a set consists of 2×4 triplets, or 2×5 doublets, as shown for the triplet and pivot planes in Figure 12-2. Anode wires of TGCs are arranged in the azimuthal direction and provide signals for r information, while read-out strips orthogonal to these wires provide signals for ϕ information. Both wire and strip signals are used for the muon trigger. Signals from two wire-planes and two strip-planes are read out from the doublet chambers, and signals of three wire-planes but only two strip-planes are read out from the triplet chambers.

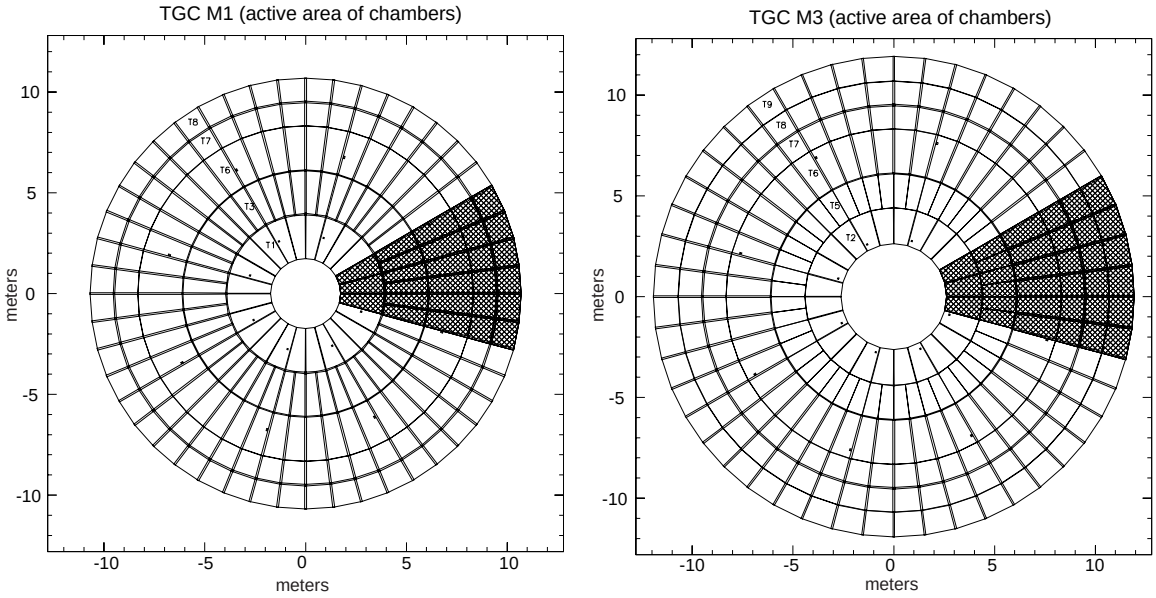

Figure 12-2 Left: R-ϕ view of the first trigger plane (triplet). Right R-ϕ view of the third trigger plane (pivot plane doublet). Each cell shown represents a single TGC unit (doublet or triplet). Chamber boundaries are shown. A single octant has been shaded and the chamber type at each radius indicated.

Several anode wires are grouped and fed to a common read-out channel for input to the trigger electronics. The number of wires per group varies from six to twenty wires, according to the desired granularity as a function of pseudorapidity, resulting in wire-group widths in the range 10.8 to 36 mm. Each chamber has 32 radial strips and thus the width of a strip is 4 mrad (8 mrad for the chambers in the Forward region). The alignment of wire-groups in consecutive layers is staggered relative to the direction pointing to the interaction point in order to achieve good position resolution whilst minimizing the number of electronics channels. Between the two planes of the doublet this staggering is by half a wire-group, and in the triplet it is by a third of a wire-group between each of the three planes.

Figure 12-1 shows the level-1 muon trigger scheme in the end-cap region. The trigger algorithm uses pivot plane hits and extrapolates to the interaction point to construct the apparent infinite-momentum path of the track. The deviation from this path of hits found in the preceding ‘confirming’ trigger planes is related to the track momentum. A window is constructed for each trigger region in the r and ϕ directions around the infinite momentum path. A coincidence is signalled if there is a hit in the window corresponding to the hit location in the pivot plane. Independent signals are generated for r and ϕ. The low-p_T trigger uses information from only the two doublets, whilst the high-p_T trigger utilizes information from all three planes,

both doublets and the triplet. Typical values for the low-p_T and high-p_T trigger thresholds are 6 GeV and 20 GeV, respectively, and these are the values for which the layout has been optimized. In order to have a good trigger efficiency and an efficient background reduction, a 3-out-of-4 coincidence is required for the doublet pair planes, for both wires and strips, a 2-out-of-3 coincidence for the triplet wire planes, and a 1-out-of-2 coincidence for the triplet strip planes is required. Finally the trigger windows are formed in r-ϕ space, with the wire signals determining the r-coordinate and strips the ϕ. The p_T threshold applied is determined by the size of the two-dimensional window, and for any single threshold this window is optimized to provide 90% efficiency. The trigger-logic can concurrently generate three different high-p_T plus three different low-p_T threshold triggers by being configured with three different high-p_T plus three different low-p_T sets of window sizes. Tracks are flagged according to the highest threshold they cross.

In almost all cases the time jitter distribution of TGC signals, Figure 10-12, is smaller than the LHC bunch-crossing spacing of 25 ns, and they thus provide bunch-crossing identification with an efficiency of 99% per gap for wires.

The level-1 system provides the level-1 trigger with information on candidate muons, defining regions-of-interest (RoIs). The RoI is a cell in r-ϕ space in which a track has been found that has crossed at least one of the six configured thresholds. For the end-cap muon trigger an RoI is a single trigger subsector as shown in Figure 12-6 in Section 12.3.

12.1.1 Chamber boundaries

In order to attain as complete a geometric acceptance as possible, gaps between chambers are avoided by overlapping adjacent chambers. More overlap may be added to cover the bending due to the magnetic field. For example, in making a wire coincidence, the wire hit needed to confirm a wire hit in the pivot plane, due to bending in ϕ, may be at a ϕ outside the area of the pivot chamber projected onto the confirming plane. The wire length in the confirming plane must be made longer by the amount of ϕ bending. An oppositely charged track in the adjacent pivot plane chamber will similarly require extending the wires in its confirming plane, resulting in overlapping wires. Making a coincidence for both pivot planes with the OR of the wires in the two confirming planes solves the problem, but doubles the effective wire length and so doubles the random coincidences. This can be avoided by overlapping only enough to cover the bending for a 6 GeV particle for the doublet confirming planes and for a 20 GeV particle for the triplet. In the η direction, wire-groups are simply overlapped to cover gaps. Similarly the strips need to be either extended by OR'ing or by extending their overlap in the η direction to cope with bending in η. In the ϕ direction, strips are simply overlapped to cover gaps.

Without care, these overlaps can lead to substantial double-counting of single muons which could swamp the real dimuon trigger. The techniques used to minimize double counting for the different cases are described below. A key principle is that for any given area in R-ϕ defined by overlapping chambers, it is not necessary to prevent double counting in both R and ϕ. If, for example, the double counting in R is prevented, but not in ϕ, a second ϕ-hit will not have a coincidence partner in the R-ϕ coincidence matrix; only one track will be found. With the same reasoning, double counting needs to be eliminated only in the pivot plane.

A simulation of the double counting rates can be found in Section 14.6.3.

12.1.1.1 Overlaps within the End-cap region

Overlap in the radial direction

Wires: In the pivot plane, wires-groups that were overlapped to cover gaps are OR'ed together. The result is a single virtual chamber in *R*; a track results in a unique hit in the pivot plane.

Strips: In the pivot plane, the strips are overlapped to cover gaps. There is no double counting since a track yields only one hit in the corresponding area of the wire-groups. Since the pivot chamber boundaries are not projected onto the confirming planes along infinite momentum tracks, particles of different momenta passing through the same pivot point may hit either of two adjacent confirming chambers. This is handled by sending the OR of pairs of confirming plane strips, one strip from each chamber, to the coincidence matrices.

Overlap in the ϕ direction

Strips: In the pivot plane, the chambers are overlapped, but some edge strips are masked from being sent to the coincidence matrices, as shown in Figure 12-3. This avoids routing signals across ϕ boundaries. The masking is programmable and can be optimized to minimize the inefficiency and the double counting. Some double counting remains.

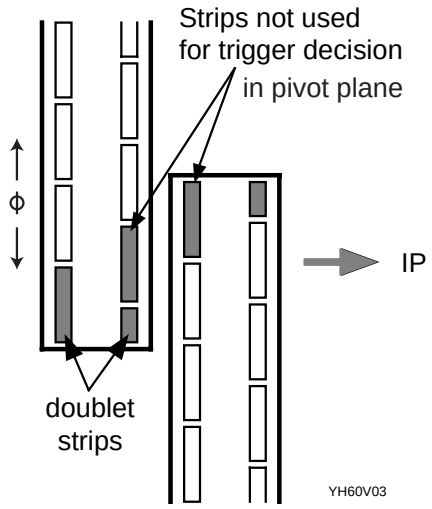

Figure 12-3 The overlapped region in ϕ for the pivot layer showing the masked strips. The effective overlap is half a strip.

Wires: In the pivot plane the wires are overlapped to cover gaps. The area where there is double counting is defined by the strips and is much less than the physical overlap of the wires. In the confirming planes, the wire overlaps are extended in ϕ to cover the bending, as described above.

12.1.1.2 Overlaps within the Forward region

A single chamber at each ϕ spans the complete radial dimension of the Forward region. In the ϕ direction, the situation is the same as the ϕ direction in the End-cap.

12.1.1.3 Overlaps between Forward, End-cap and Barrel regions

Both the boundary between Forward and End-cap and that between End-cap and Barrel are overlapped to cover the gaps and the bending. The Forward–End-cap division is projective along the infinite momentum path. The overlaps of the low-p_T planes are defined by the 6GeV particle trajectories and those for the hi-p_T planes by 20GeV trajectories. If double counting in this region is excessive, the size of the overlap can be reduced at the expense of less acceptance by programming masks to disable wire-groups. Double counting in the End-cap–Barrel overlap is removed by the MUCTPI, described in Chapter 14.

12.2 Overview of the level-1 trigger electronics system

The signals generated by the TGCs are amplified, discriminated and shaped on the detector. Detector-mounted electronics first identify the bunch-crossing and then find coincidences independently in r and ϕ roads. The coincidences provide r, δr and ϕ, $\delta\phi$, where r and ϕ are the track co-ordinates in the pivot plane and δr and $\delta\phi$ are the track's deviation from the infinite momentum track.

Electronics situated outside the ATLAS cavern combines the measurements of r, δr and ϕ, $\delta\phi$ to make a trigger decision which is then passed to the Muon Interface to the Central Trigger Processor, MUCTPI (see Chapter 13). The chamber hits and the intermediate r, δr and ϕ, $\delta\phi$ values are read out by the on and off-detector electronics.

Figure 12-4 shows an overview of the TGC level-1 trigger electronics scheme and Figure 12-5 shows the placement of these electronics relative to the trigger chambers. The wire and strip signals emerging from the TGCs are fed into a two-stage amplifier in an Amplifier Shaper Discriminator (ASD) circuit. Four such circuits are built into a single ASD chip and four ASD chips are incorporated into an ASD Board; hence each ASD Board handles 16 channels of signals. The ASD Board is physically attached to the edge of a TGC and enclosed inside the TGC electrical shielding.

Signals from the ASD Boards are sent to a Patch-Panel, Figure 12-7, which houses receivers for the ASD outputs, TTC receivers, Bunch-Crossing Identification (BCID) circuits, logic to take care of physical overlap in the TGCs and fan-outs. Bunch-crossing identification is performed at this stage, prior to the signal passing through any trigger-logic. Up to 512 channels are handled by each Patch-Panel. The Patch-Panel also routes signals to and from the Detector Control System (DCS) and the muon end-cap trigger control and monitoring system and supplies power to Slave Boards.

Outputs from the Patch-Panel are sent to corresponding Slave Boards where the coincidence and read-out circuits are placed. There are four different types of Slave Board; wire and strip boards for each of the triplet and doublet. They differ in their number of inputs, the kind of coincidence made and the maximum window width.

Slave Boards are placed on the accessible outer surfaces of the TGC wheels. Thus, electronics for the two doublets are mounted on the outside of the outer doublet wheel and those for the triplets on the inner surface of the triplet wheel. (See Figure 12-5.) Information from the Slave Boards is encoded to produce more compact signals and the encoded coincidence information is passed to a High-p_T coincidence Board located near the outer rim of the triplet wheel. Signals from the doublet and triplet Slave Boards are combined here to find high-p_T track candidates . Wire (r-coordinate) and strip (ϕ-coordinate) information is treated separately.

Signals from the High-p_T Boards are sent to Sector Logic Boards containing an R-ϕ coincidence unit and track selectors, to select the highest-p_T coincidences. These are located outside the main ATLAS cavern in USA15. The resulting trigger information is sent to the Muon Central Trigger Processor Interface (MUCTPI) in a standard format [12-1]. Full-information data sets are read-out through the DAQ system in parallel with the primary trigger-logic. For read-out purposes the Slave Boards of one or more trigger sectors are grouped into Local DAQ Blocks. Each Slave Board is connected to the so-called Star Switch which manages the data collection for a Local DAQ Block. The transfer, via optical fiber, of data from the Star Switch to the Read Out Drivers (RODs) in USA15 is managed by the Local DAQ Master.

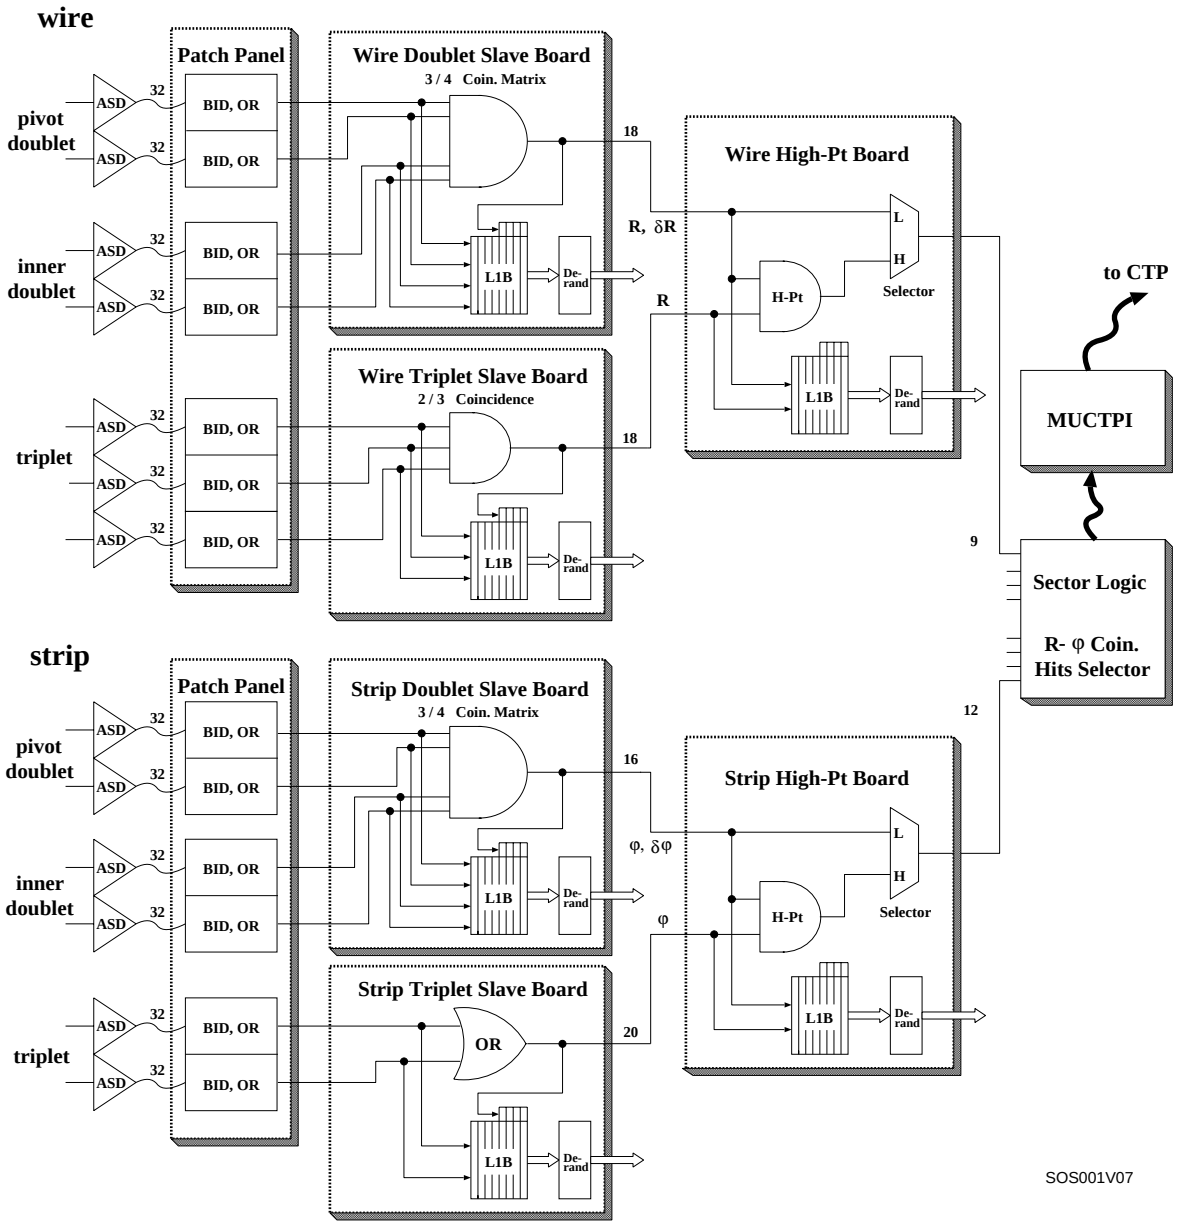

Figure 12-4 Overview of the TGC level-1 trigger scheme. Prior to the Sector Logic, wire and strip information are treated independently. The original number of signals from the ASD boards are reduced by over an order of magnitude before they are sent to Sector Logic, located outside the ATLAS cavern, via optical transmission lines. The numbers on the connecting lines indicate the data width.

The total latency of the system, from the bunch-crossing in which the interaction occurs until the delivery of the level-1 track candidates to the MUCTPI is 1.15μs. See Section 12.9, for details.

Each electronics board may communicate with one or more of three information paths in addition to the read-out data and trigger signal paths. They are the Detector Control System (DCS), the Timing, Trigger and Control system (TTC), and the muon end-cap trigger control and monitoring system. DCS is a slow path primarily for monitoring the environment, power, gas flow, etc. It is ATLAS-wide. TTC is the ATLAS uni-directional synchronous system for distributing the 40MHz clock, the level-1 accept and other timing signals. These signals are routed to the Slave Boards via the Patch-Panel.

The environment in which the electronics modules sit is considered in the technology choice made for each component of the trigger system. Those items mounted on the TGCs in the ATLAS cavern must be radiation tolerant and able to operate in the ambient magnetic field.

The latest values of various dimensions, channel and board counts, etc. can be found in the TGC Excel workbook [12-2], which is the master database for all TGC parameters.

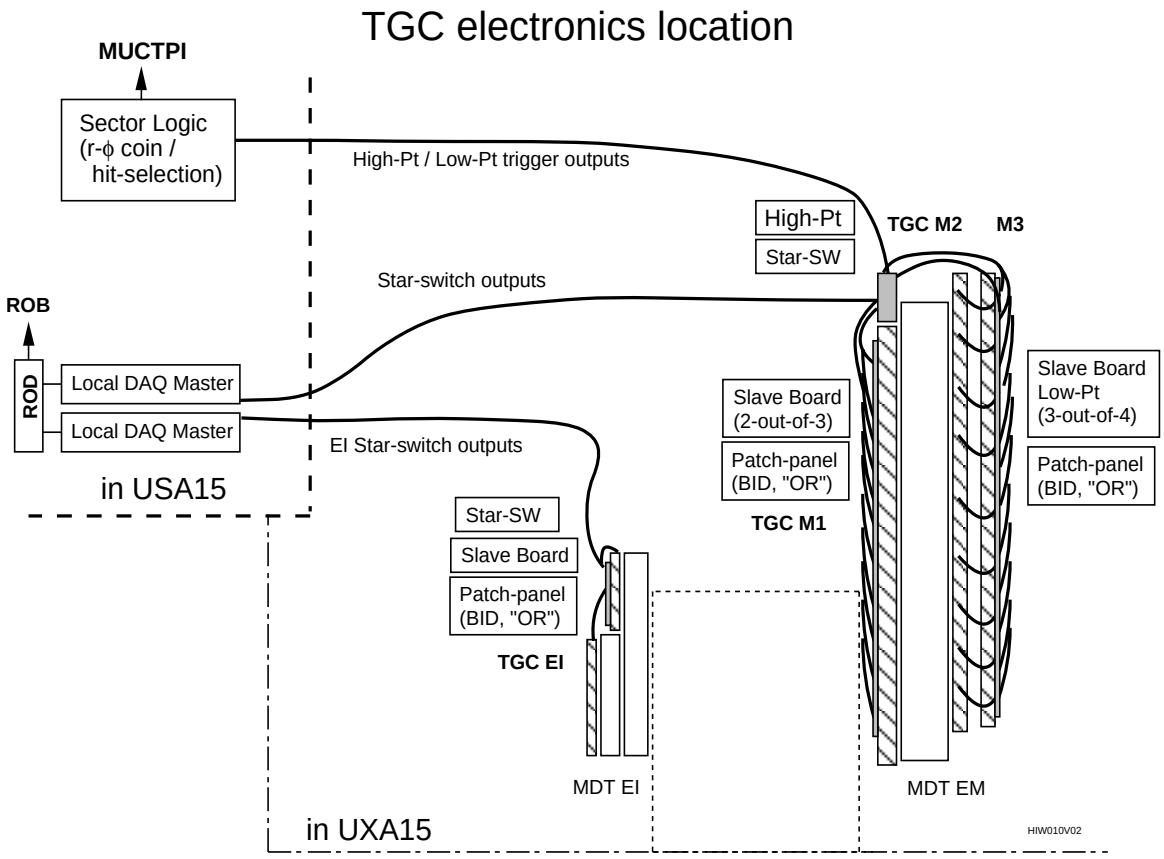

Figure 12-5 TGC electronics placement. Electronics prior to the high-p_T logic are placed on the TGC wheel surfaces. High-p_T logic and Star Switch are placed on the rim of the TGC triplet wheel. Sector logic and electronics beyond this point in the trigger and DAQ chain are placed in USA15.

12.3 System segmentation

Figure 12-6 shows the pivot plane octant formed by the TGC doublet plane furthest from the interaction point. The pivot plane is divided into two regions, End-cap ($|\eta| < 1.9$) and Forward ($|\eta| > 1.9$). The End-cap region of each octant is divided into six trigger sectors in ϕ, where a trigger sector is a logical unit that is treated independently in the trigger. Similarly the Forward region of each octant is divided into three trigger sectors. Thus in each end-cap of TGCs there are 48 End-cap trigger sectors and 24 Forward sectors.

The smallest regions shown in Figure 12-6 are trigger subsectors which correspond to the smallest unit area of the trigger segmentation. A trigger subsector corresponds to eight channels of wire-groups and eight channels of read-out strips. An End-cap trigger sector contains 49 η rows by 4 ϕ columns of trigger subsectors, a total of 196 trigger subsectors. A Forward trigger

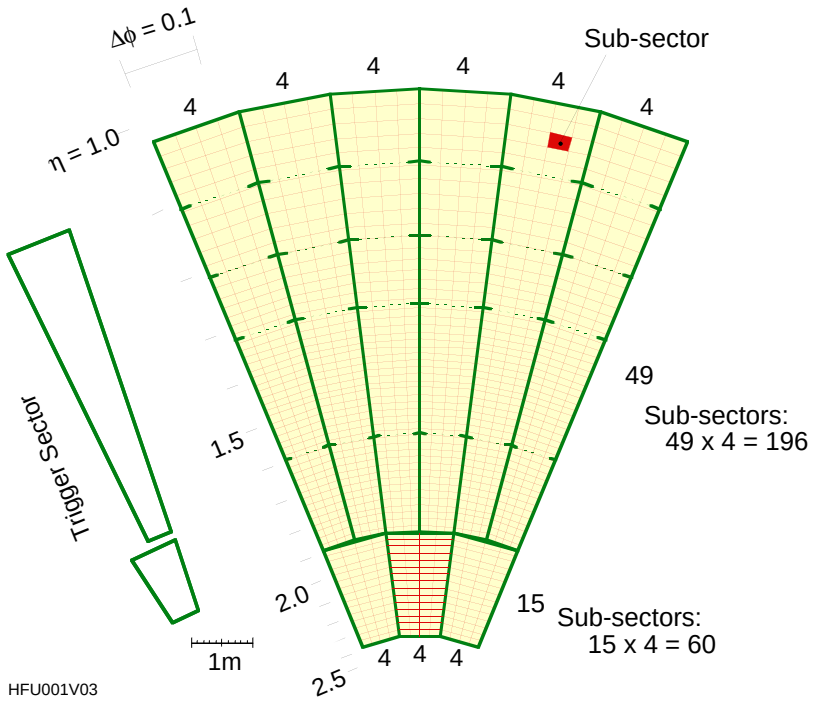

Figure 12-6 TGC level-1 trigger segmentation for an octant. One octant wheel is divided into six End-Cap sectors and three Forward sectors. Bold lines in the figure indicate individual trigger sectors. They are further subdivided into trigger subsectors.

sector contains 15 η rows by 4 ϕ columns of trigger subsectors, a total of 60 trigger subsectors. Each η–ϕ trigger subsector corresponds to one Region of Interest (RoI). Each subsector is treated independently in the trigger so that the δR and δϕ inputs that determine the p_T condition applied can be set separately for each subsector.

In each trigger sector the two highest-p_T track candidates are selected and sent to the MUCTPI.

12.4 Trigger system boards

12.4.1 ASD board

TGC signals from wire-groups and strips pass through an Amplifier Shaper Discriminator IC (ASD-IC) placed on ASD Boards which are housed inside the TGC shielding cage. The ASD-IC uses SONY Analog Master Slice technology (bipolar). Each ASD-IC contains four channels and there are four ASD-ICs on each ASD Board, giving 16 channels per ASD Board. Outputs of the ASD-ICs are Low Voltage Differential Signalling, LVDS [12-3], signals which are sent to a Patch-Panel Board. An analog amplifier output from one channel on each ASD board is used to monitor the TGC performance. The ASD Board receives the threshold voltage (V_{th}), analog test-pulse signal (differential signals) and ±3V power from the Patch-Panel. The analog test-pulse is fanned out to all ASD inputs through on-board capacitors. The amount of charge in the test-pulse and its timing are controlled in the Patch-Panel. This function is required for timing calibration as well as verifying the adequacy of threshold levels. The ASD Boards and

ASD chips are part of the muon system front-end electronics and are described in detail in the ATLAS Muon Spectrometer TDR [12-4].

12.4.2 Patch-Panel

The Patch-Panel, Figure 12-7, receives TGC wire-group and strip signals from the ASD Boards, performs bunch-crossing identification and allows the correct treatment of physical overlaps in the TGCs by OR'ing signals from overlapping regions as discussed in Section 12.1.1 and below. Inputs to the Patch-Panel from the ASD boards are LVDS signals; outputs to the Slave Boards are single-ended signals. The Patch-Panel provides the analog test-pulse differential signals and the ±3V power to the ASD Boards. A DAC, loaded via DCS, generates the ASD threshold voltage (V_{th}). The Patch-Panel also distributes the Timing and Trigger Control (TTC) and other control signals as well as DC power to the Slave Boards, and provides the interface between the Detector Control System (DCS) and the trigger electronics. The Patch-Panel Board is placed on the TGC wheel surface and each board receives signals from up to 32 ASD Boards and additional signals from neighbouring channels, yielding a maximum of 512 output channels per board. Each Patch-Panel has a unique electronic ID.

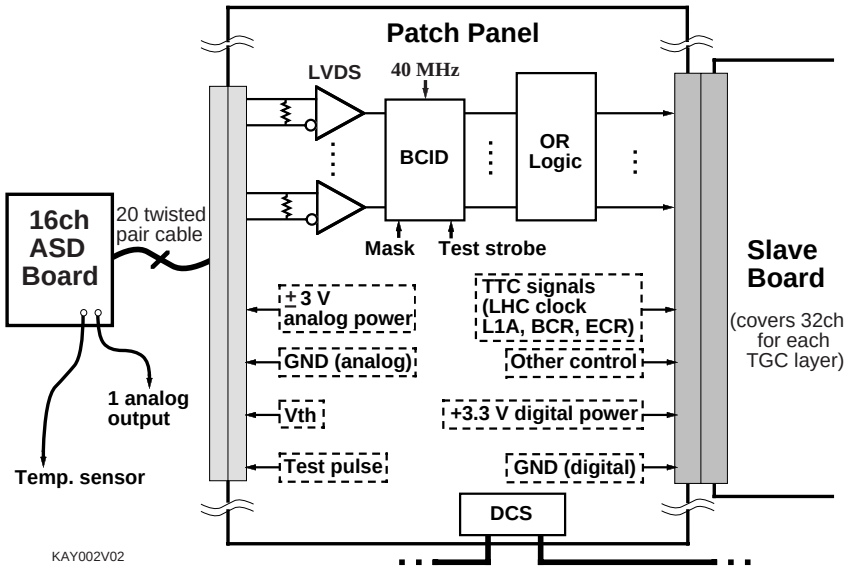

Figure 12-7 Functional structure of a Patch-Panel.

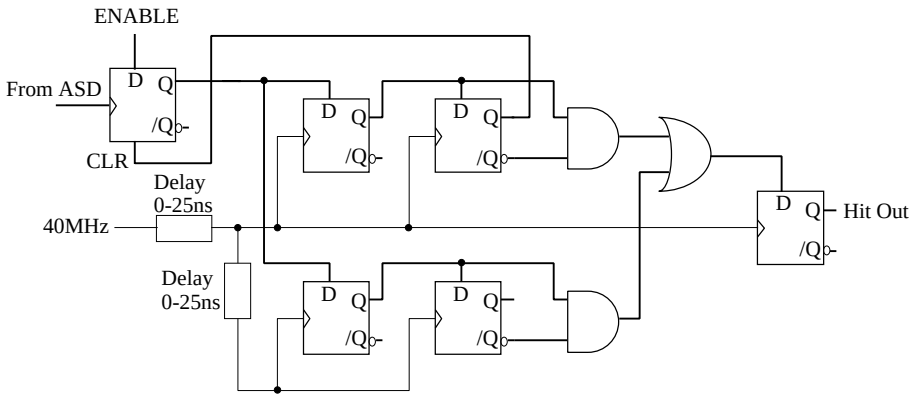

Figure 12-8 The Bunch-Crossing-ID circuit, located in an IC package on a Patch-Panel. The first delay unit is to adjust for signal arrival times and the second delay unit (in the lower-half section) allows adjusting the gate width to be between 25ns and 50ns.

Bunch-crossing identification (BCID) is performed in a dedicated circuit on the Patch-Panel. Figure 12-8 shows the conceptual design of the BCID circuit. It contains two adjustable delays. The first delay is used to adjust for the phase difference between the 40 MHz clock on the board and the earliest arrival times of signals from the ASDs. The second delay is used to adjust the effective gate width to be between 25ns and 50ns [12-5]. Both delays can be adjusted with sub-nanosecond precision. Set values for each delay are common to a group of signals that come from the same ASD Board. The BCID circuit also contains system diagnosis functions as well as the ability to mask individual channels. Beyond the BCID, all processing is controlled by a synchronous 40MHz clock.

In order to prevent inefficiency, neighbouring TGC units are overlapped in both the *r* and ϕ directions which allows the possibility for a single muon passing through the overlap region to make two triggers. In the *r* direction, fake double triggers are prevented using logic in the Patch-Panel, where an OR is made of the channels of wire-groups of the detectors in the overlap region. As a result, four (triplet case) or five (doublet case) TGCs at the same ϕ can be treated as a single continuous chamber. This eliminates the double counting of a track in an overlap region.

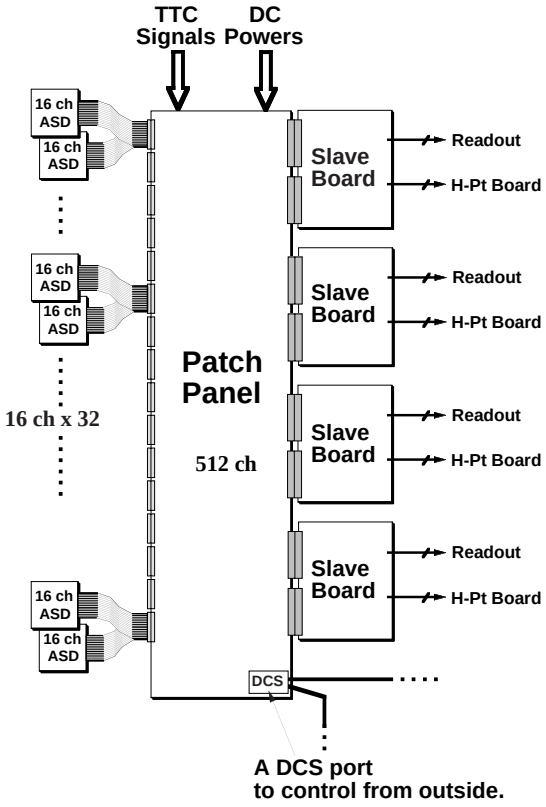

Figure 12-9 TGC P-S Pack (Patch-Panel and Slave Board package). A Patch-Panel is directly connected to four Slave Boards forming a unit (P-S Pack) including power/signal distribution and cooling/shielding.

The printed-circuit pattern of each Patch-Panel reduces the complexity of the signal routing to enable Slave Boards to be of the same design regardless of their placement. This, however, demands several different Patch-Panel patterns.

Each Patch-Panel has a unique electronic ID which can be read by the Slave Board, combined with the Slave Board's geographic address on the Patch-Panel and incorporated into the data stream. This unique ID is implemented by a laser-programmed ROM. The Patch-Panel receives TTC signals via a TTCrx chip and distributes them (40MHz clock, Level-1 Accept (L1A), Bunch-crossing Counter Reset (BCR), Event Counter Reset (ECR)) to the Slave Boards.

Four Slave Boards are attached to the side of each Patch-Panel, constituting a compact, planar package, known as the PS-pack. Below this is a cooling panel and above it a shielding cage. The PS-pack is shown in Figure 12-9. Cooling is described in Section 12.12.3.

The major Patch-Panel functions such as LVDS receivers, DAC, test-pulse generator, BCID and OR logic are covered by 4 (or perhaps 8) CMOS ASIC ICs per Patch-Panel. The Patch-Panel dimension is approximately 130cm by 20cm and the Slave Board dimensions are that of a 6U-VME Board.

12.4.3 Slave Boards

Slave Boards receive Bunch-crossing-ID assigned wire and strip signals from the Patch-Panel and perform coincidence operations on the input signals from TGC planes. There are four different types of Slave Board, two for the triplets and two for the pair of doublets, where wire and strip signals have different Slave Boards. Each Doublet Slave Board (DSB) receives signals from a pair of TGC doublets and performs 3-out-of-4 coincidence operations on them, to form low-p_T triggers. Each Triplet Slave Board (TSB) receives signals from a triplet plane of TGCs and performs coincidence operations on them. This is 2-out-of-3 coincidence for the wires and 1-out-of-2 for the strips. Outputs from both the DSB and TSB are sent to the High-p_T Board where high-p_T trigger signals are formed.

All Slave Boards contain a Level-1 Buffer and Derandomizer in addition to the trigger-logic circuits to enable data read-out at a later stage. The raw data, as well as BCID and L1ID, are read out. Further details are given in Section 12.6.

All Slave board coincidence matrices are built in a sub-micron CMOS gate array ASIC chip. It is hoped that a single ASIC will be able to include the functions needed by all four types of Slave Board. The dimensions for the Slave Boards are the same as a 6U-VME Board.

12.4.3.1 Slave Board for the doublet pair (DSB)

The Doublet Slave Board (DSB) services the four layers of the two sets of TGC doublets. Signals from the Patch-Panel first pass through a phase-adjust circuit (adjustment range 0–50ns in steps of 12.5ns) to compensate for transmission-delay differences between signals and then continue through a 3-out-of-4 coincidence matrix. A block diagram of the Doublet Slave Board is shown in Figure 12-10.

Each DSB covers a 32-channel per layer width of either wire or strip signals, giving 128 channels of inputs to a DSB from the four TGC layers. For wires, a track segment found by the coincidence matrix is encoded as an 18-bit word containing *r*, its pivot plane position, and δR,

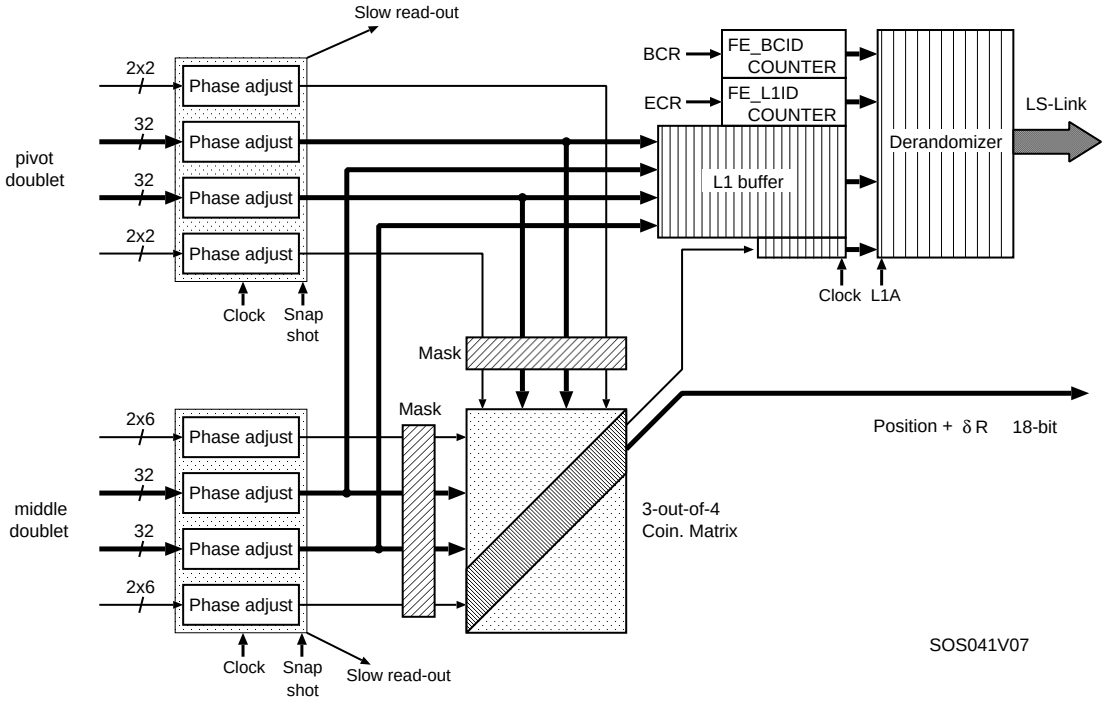

Figure 12-10 The functional structure of a Slave Board for the TGC doublet pairs. Shown is the wire Slave Board; the strip Slave Board has 16-bit output instead.

its deviation from the infinite momentum path. Similarly, for strips, a 16-bit word is encoded containing ϕ and $\delta\phi$ information. If high background rates produce a large number of accidental coincidences, the 3-out-of-4 coincidence matrix can be programmed to provide a 4-out-of-4 coincidence offering a more robust trigger. In this mode, however, the acceptance is lower, since those chamber regions with internal supports or frames can have at most 3-out-of-4 hits. Just prior to the 3-out-of-4 coincidence matrix are masking circuits to remove unwanted signal lines from input to the trigger-logic. (Such masked channels will however still be read out.) This masking facility is used to eliminate noisy channels, to force on dead channels, or to select particular channels for testing.

Figure 12-10 also shows the data read-out components that are placed on the Slave Board. These are described in Section 12.6.

Figure 12-11 shows the 3-out-of-4 low-p_T coincidence matrix structure, where the two sets of inputs are 32-channel wide patterns from each of the two sets of doublets. For both sets of doublets some additional bits are shared between adjacent Slave Boards:

- A track passing through the pivot plane near the edge of a Slave Board may be bent into the area of the middle doublet handled by the adjacent Slave Board. Therefore, the six adjacent channels of the neighbouring Slave Board are provided to allow such tracks to make a coincidence with the pivot plane hits.
- Two adjacent pivot plane channels are also provided to neighbouring Slave Boards. They enable a Slave Board to know that a cluster is spanning a Slave Board boundary and that only one of the Slave Boards should use the pivot plane hits near the edge. This avoids double counting due to clusters spanning boards. See Figure 12-13 and the text below for more on declustering.

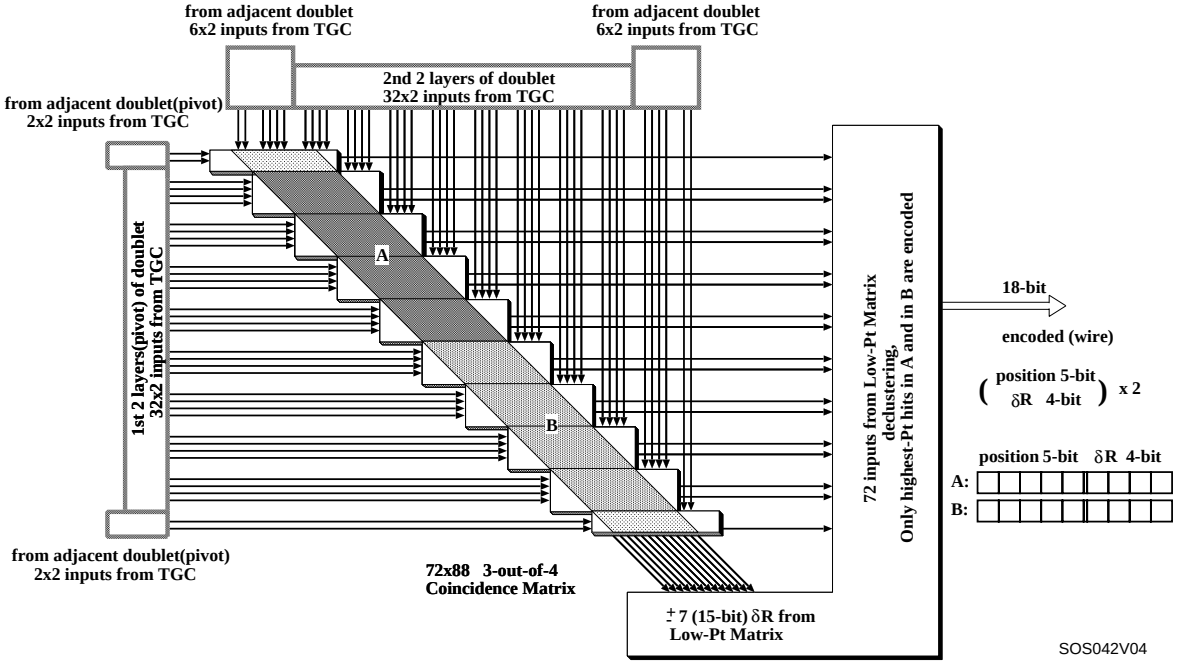

Figure 12-11 Structure of a 3-out-of-4 Doublet Slave Board Low-p_T matrix and the encoded output information. The wire (*r*) Slave Board is shown.The strip Doublet Slave Board is similar, but with $\delta\phi=\pm 3$.

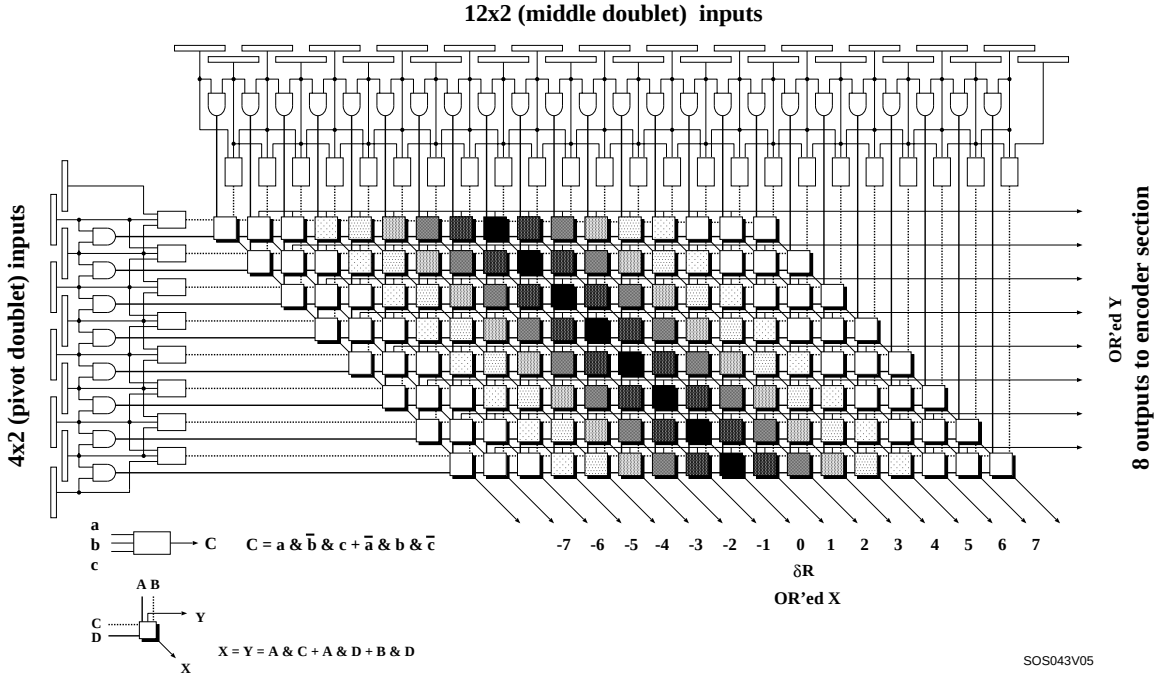

Figure 12-12 Detailed structure of the 3-out-of-4 Low-p_T matrix for wires. Also shown is the function of a matrix element. Cells with the same shading have the same δr. The matrix for ϕ is the same, except $\delta\phi$ goes from -3 to +3. (Actually X and Y in the figure are the same.)

Thus for the middle doublet 32 channels plus six more from each adjacent board are included, giving a 44-channel width in two layers, or a total of 88 bits. For the pivot plane doublet, two channels from each of the adjacent boards are included, and so, in total, signals for a 36-channel width in two layers, giving 72 bits, are received.

Figure 12-12 shows in more detail the structure of the 3-out-of-4 low-p_T coincidence matrix and the logic equation of the coincidence element. Signals first pass through circuits to account for the half-cell staggering between layers in order to improve the effective resolution. The 3-out-of-4 coincidence is then performed in each coincidence element. When a coincidence is found, a hit signal and δR (or $\delta\phi$) signal are generated. Hit signals in the elements of each row are OR'ed, while the δR (or $\delta\phi$) signals are OR'ed along the diagonal elements. For R, the δR output pattern is 15 bits wide; for ϕ, the $\delta\phi$ output pattern is seven bits wide.

In Figure 12-11, the coincidence matrix is shown divided into two halves, A and B, (each half services signals from a 16-channel per layer width) and the highest-p_T track candidate found in each half is selected. Hence the OR'ing mentioned above is done for each half independently. To select the highest-p_T track, δR (or $\delta\phi$) outputs are examined first and then the selected hits with the smallest $|\delta R|$ from the 3-out-of-4 coincidence are passed through declustering logic to produce a single track. Figure 12-13 shows the declustering rules. Even though a real particle almost always makes at most two hits, a cluster is treated as a single track, regardless of its size. If more than one cluster is found in any half-matrix, only the highest-p_T track candidate is kept. For wires, the 15-bit δR pattern ($\delta R = -7$ to $\delta R = +7$) is encoded in four bits and position information is encoded in five bits so that the R information for each track is nine bits wide. This encoding reduces the required number of cables and hence the line-driver power dramatically. Since each of the two halves of the Slave Board can produce a track candidate, the output from this circuit is 18 bits wide. For strips, the seven bit wide $\delta\phi$ pattern ($\delta\phi = -3$ to $\delta\phi = 3$) is encoded in three bits. Position in ϕ is encoded in five bits so that the ϕ information for each track is eight bits wide and for the two highest-p_T tracks, 16 bits. Despite this attention to multi-muon encoding, the occurrence of more than one muon in the area covered by a Slave Board is rare.

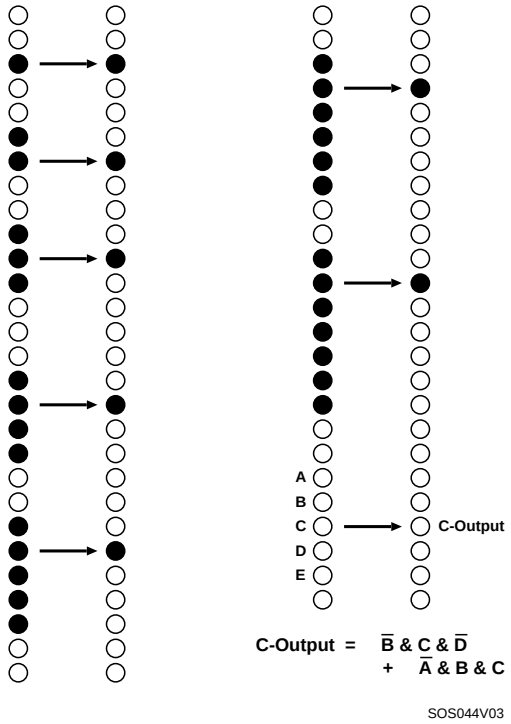

Figure 12-13 Declustering rules. When a hit pattern wider than two is found, the rules shown assign a position to the track.

12.4.3.2 Triplet Slave Board (TSB)

The Triplet Slave Board (TSB) (Figures 12-14 and 12-16) performs coincidence operations on signals from the three layers of the triplet. The TSB is similar to the Doublet Slave Board. The wire coincidence operation uses 2-out-of-3, and the strip 1-out-of-2, logic. Figures 12-14 and 12-16 show the TSB trigger-logic and read-out components for wires and strips respectively. For details of the read-out components on these Slave Boards, see Section 12.6.

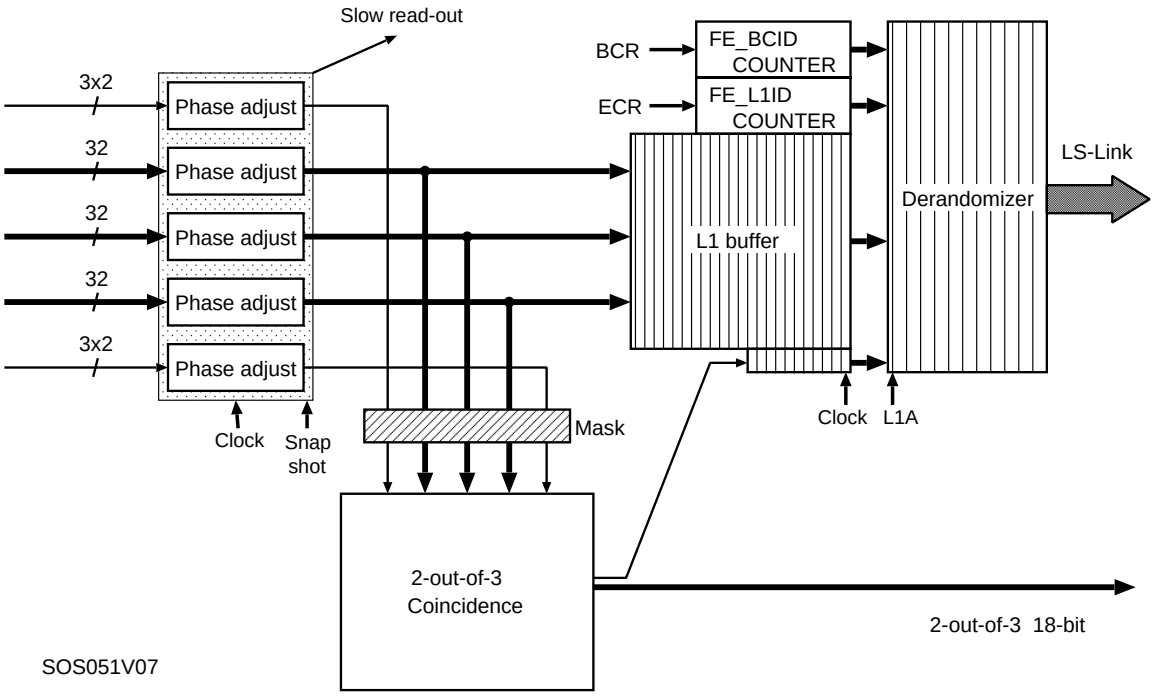

Figure 12-14 Slave Board for TGC Triplets (for wire signals).

Wire Triplet Slave Board Each wire TSB receives from each layer of TGCs a 36-channel wide wire pattern. It also includes two bits from each of the two adjacent boards to take care of boundaries between regions covered by adjacent Slave Boards, resulting in a 108-bit input pattern. The coincidence matrix is subdivided into three 32-channel sections. The structure of the 2-out-of-3 coincidence for triplet wire signals and the logic equation are shown in Figure 12-15. Signals first pass through circuits to account for the cell-staggering of 1/3 between successive layers before they go through the coincidence circuits. A 3-out-of-3 coincidence can also be required should the background rate be so high that the accidental coincidences force a more robust approach. The output is encoded to form a 5-bit hit position plus a 1-bit hit-found flag giving a total of six bits for each hit. The hit at highest *R* is selected from each section giving three hits per TSB, an 18-bit wide result after encoding (see Figure 12-15).

Strip Triplet Slave Board The middle TGC of the triplet does not have strips so there are only two 32-bit patterns. No channels from adjacent boards are used in this logic. A single strip TSB is able to serve **two** TGC triplets, and so each strip TSB receives 128 channels of strip signals. The strips are staggered by 1/2. Figure 12-16 shows the block diagram of the strip TSB. The inner structure is similar to that for the wires. However, as shown in Figure 12-17, the coincidence operation performed for the strips is 1-out-of-2. The full logic equation is given in the figure. Again, this coincidence section can be configured to provide a 2-out-of-2 coincidence should the background rate become too high. The 32-bit wide per layer input pattern is split into four sections, each eight channels per layer wide. The declustering logic, as described in Section 12.4.3.1, selects one hit per section. Each hit is encoded as a 4-bit position plus a 1 bit valid hit flag. The four sections thus give a 20-bit wide output per triplet as shown in Figure 12-17, resulting in a 40-bit output per strips TSB.

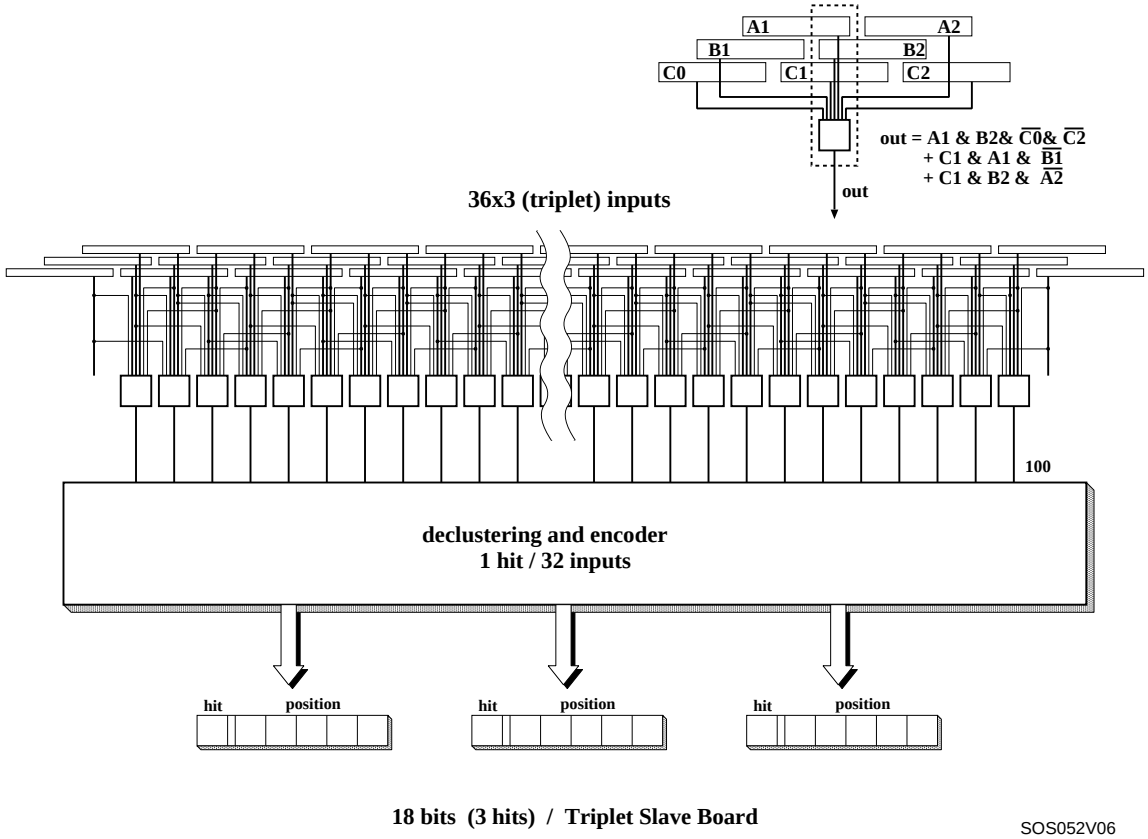

Figure 12-15 Structure of the Triplet logic for wire signals. Logic to deal with the staggering of triplet layers and the output format after encoding are also shown.

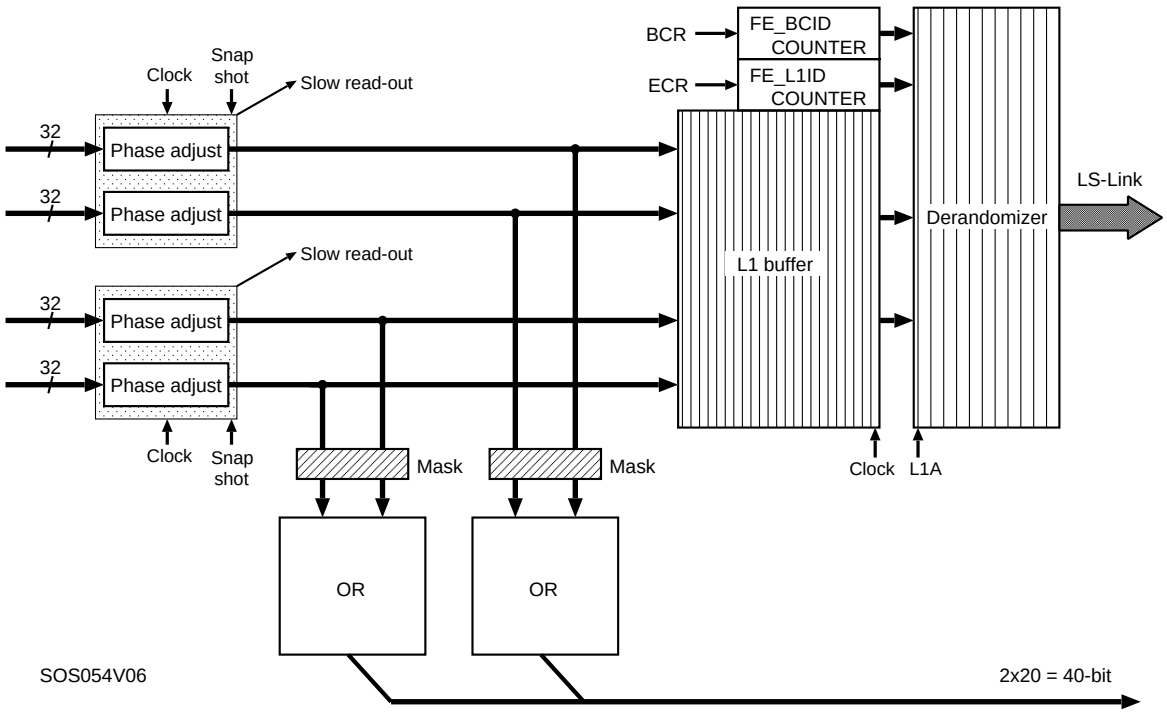

Figure 12-16 Functional structure of a Slave Board for TGC Triplet strip signals. There are only two active layers of strips per ‘triplet’. One chip handles up to two triplet units independently.

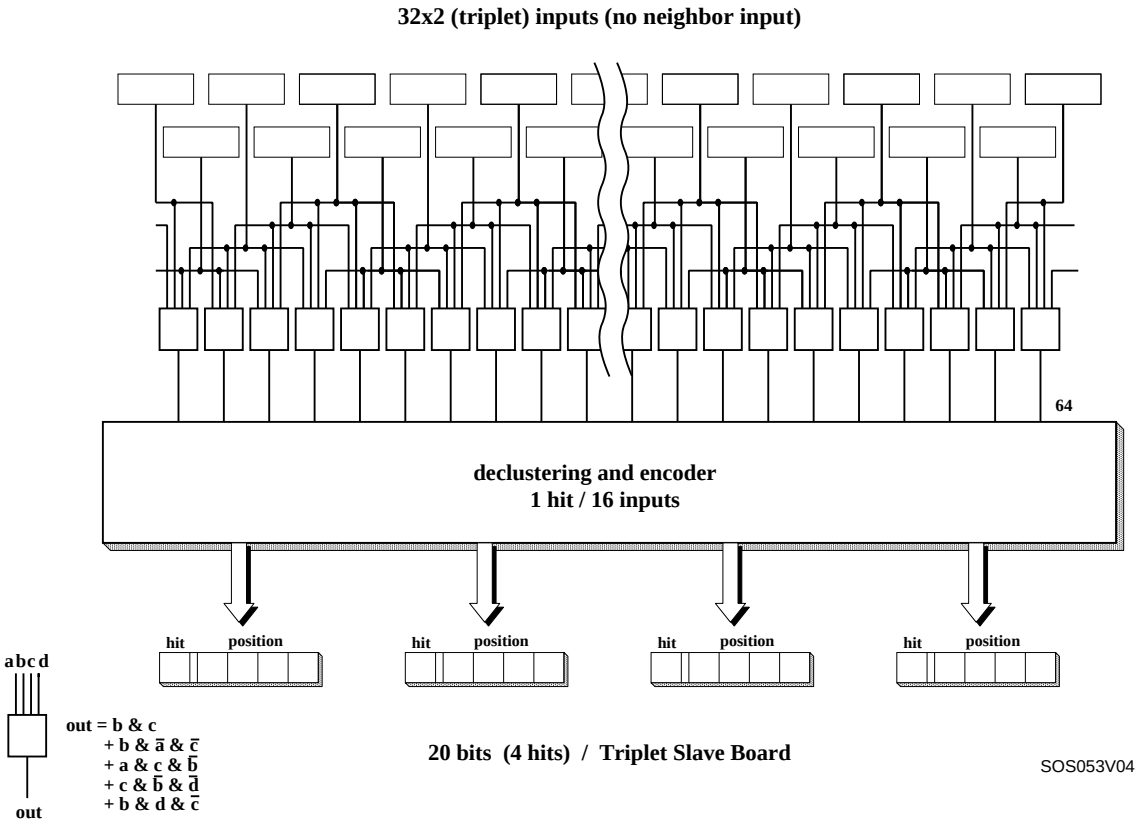

Figure 12-17 The functional structure of the triplet Slave Board for strip signals. Logic to deal with the staggering of triplet layers and output format after the encoder section are also shown.

12.4.4 High-p_T Board

The High-p_T Board, shown in Figure 12-18, receives the low-p_T signals from sets of four Doublet Slave Boards and hits from a projectively-matching set of Triplet Slave Boards, and performs a two-fold coincidence on these input signals to form high-p_T trigger signals. The board then merges the high-p_T output with the low-p_T signals, giving priority to the high-p_T signals. It has a similar structure to that of a Slave Board. Signals from the DSBs and TSBs are received by phase-adjust sections of the High-p_T Board. The coincidence operation is performed by a 256×288 two-fold coincidence matrix that combines the doublet and triplet results to produce high-p_T trigger signals. There are separate High-p_T Boards for R and ϕ. The R, δR and ϕ, $\delta\phi$ information from the High-p_T Boards is passed to the Sector Logic. The read-out of the High-p_T Board input signals is analogous to the Slave Board read-out of the chamber hit patterns. The High-p_T Board input from the Slave Boards thus becomes part of the read-out data flow in the same way as the chamber data. See Section 12.6, "Data read-out".

In Figure 12-19, one of the four identical slices of the high-p_T coincidence matrix is shown for wire signals. Since the outputs from the Slave Boards are encoded, they pass through decoders before entering the two-fold coincidence section. Figure 12-20 shows in more detail the internal structure of the eight blocks in Figure 12-19. Hits in each coincidence unit are OR'ed diagonally, resulting in 31-bit wide δR information (δR in the range -15 to +15). Four of the sections of the type shown in Figure 12-20 are grouped (corresponding to 'AABB' or 'CCDD' of Figure 12-19) and the track with the smallest $|\delta R|$ is selected. If there is no track candidate and a low-p_T track

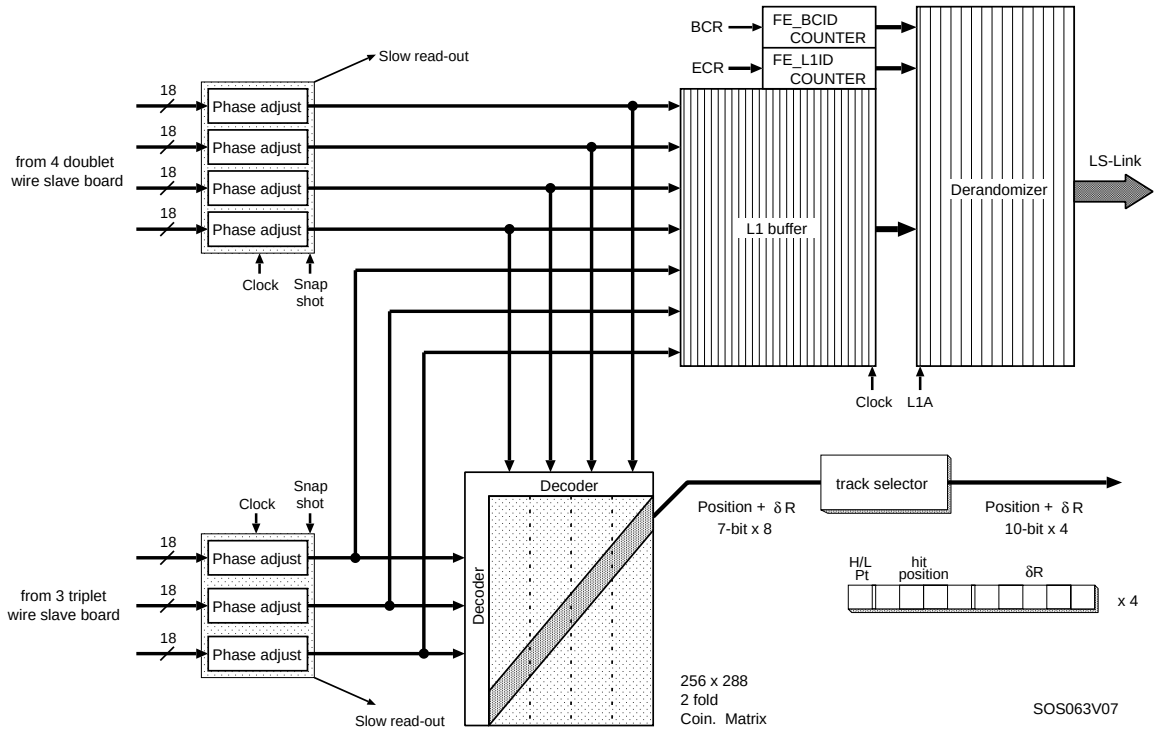

Figure 12-18 Functional structure of a High-p_T Board. The High-p_T Board has a similar structure to other Slave Boards. Note that the coincidence matrix is divided into four slices, each handling a quarter of the range of the pivot plane coordinate. Each slice corresponds to a Doublet Slave Board and can provide up to two track candidates (encoded format shown) to the track selector.

exists, then the low-p_T track is passed. If both exist, the high-p_T track is passed but its δR is checked to be consistent with the low-p_T δR, by requiring low-p_T $\delta R < 1/2$ the hi-p_T δR. These tracks are encoded as a 5-bit δR, a 1-bit hit position and 1-bit high/low-p_T indicator, giving a total of seven bits per track. The resulting output is thus 56 bits wide and is sent to the Track Selector unit (as shown in Figure 12-18).

The Track Selector receives up to eight tracks. There is a very low probability for more than four tracks in the area serviced by a High-p_T Board. Should this occur, the tracks at the highest radius are selected because the real muon signal is roughly constant in R but the backgrounds fall with increasing R. The up to four track candidates are sent to the Sector Logic in USA15 via an optical fiber. Forty bits are transmitted in a 25ns interval with the data format shown in Figure 12-18.

The High-p_T Board for strips is similar to that for wires. However, since the detector area covered by each matrix is large, there is no Track Selector and up to 8 tracks can be sent to the Sector Logic. Hits in each ϕ coincidence unit are OR'ed diagonally, resulting in a 15-bit δϕ value (δϕ in the range −7 to +7) (c.f. for wires 31-bits and ±15). Each track candidate consists of a 4-bit δϕ, a 1-bit hit position and a 1-bit high/low-p_T indicator, giving a total of six bits per hit. The resulting output, 48 bits wide, is sent to the Sector Logic in USA15 via an optical fiber; 48 bits (24 bits × 2) are transmitted in a 25ns interval.

The high-p_T circuit is built in a CMOS gate array ASIC chip. The board's dimensions are the same as for a 6U-VME Board.

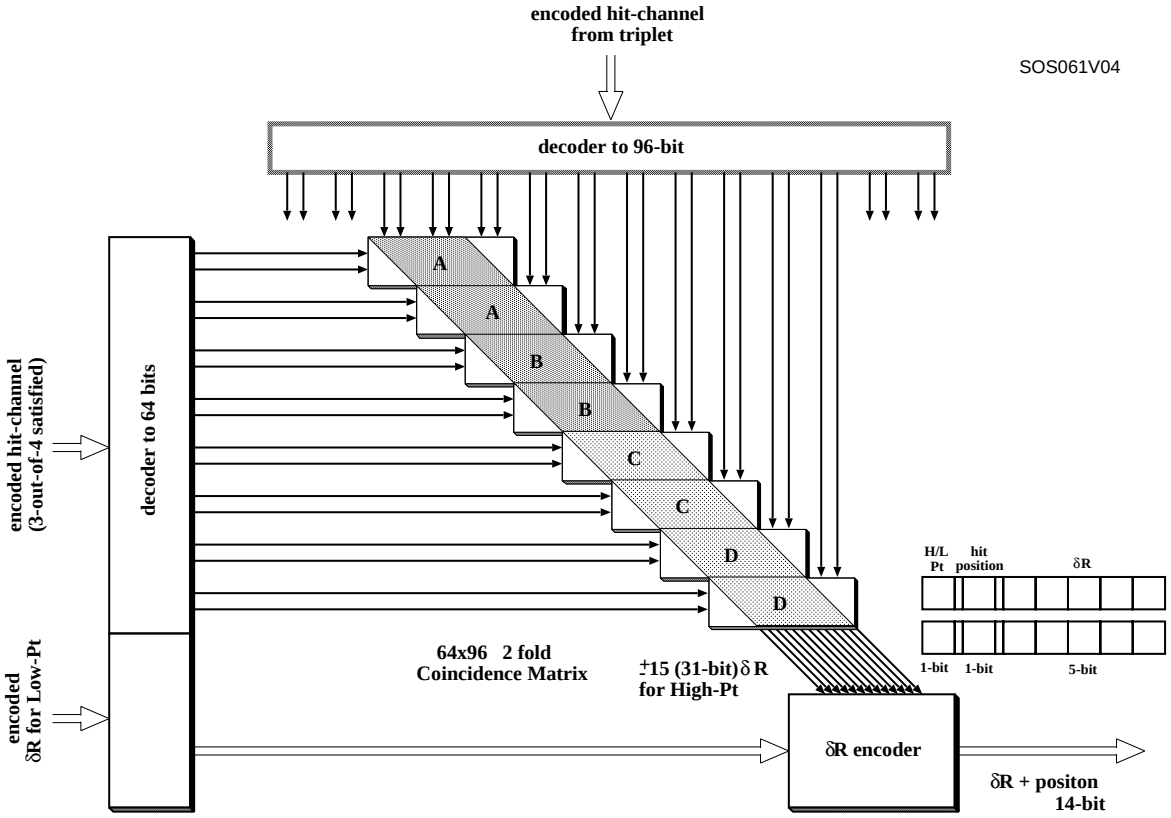

Figure 12-19 High-p_T matrix structure for wires. Shown is one of four identical vertical slices (here shown horizontally) of the high-p_T coincidence matrix shown in Figure 12-18, together with its output format after encoding. A similar structure exists for strips. The two blocks, 'AABB' and 'CCDD', can each produce up to one track candidate. If they do not find a candidate, they can pass through a low-p_T track.

12.4.5 Sector logic

The Sector Logic, Figure 12-21, combines the δR information from the wires with the δϕ information from the strips, using the outputs of the High-p_T Boards. Each Sector Logic Board handles the trigger data from a single trigger sector, as shown in Figure 12-4. The Sector Logic Board contains a track selector which selects the two highest-p_T tracks in each sector. This Sector Logic consists of *r*-ϕ coincidence matrices, a Track-Preselector and a Track Selector. Hit information for *r* and ϕ is encoded so decoding is necessary immediately after the receivers. Combining the δϕ and δr information through the *r*-ϕ coincidence matrix makes the track information more complete. Figure 12-22 shows the *r*-ϕ coincidence matrix element which receives δR (Low / High) and δϕ (Low / High) information and outputs six levels of p_T as shown, namely three ranges for each of Low- and High-p_T . Each level can make full use of both high- and low-p_T input information.

Since there is no radiation problem in USA15 where the Sector Logic is located, the *r*-ϕ coincidence can be implemented with FPGAs. This allows the coincidence element to be reconfigured as needed to match the running conditions.

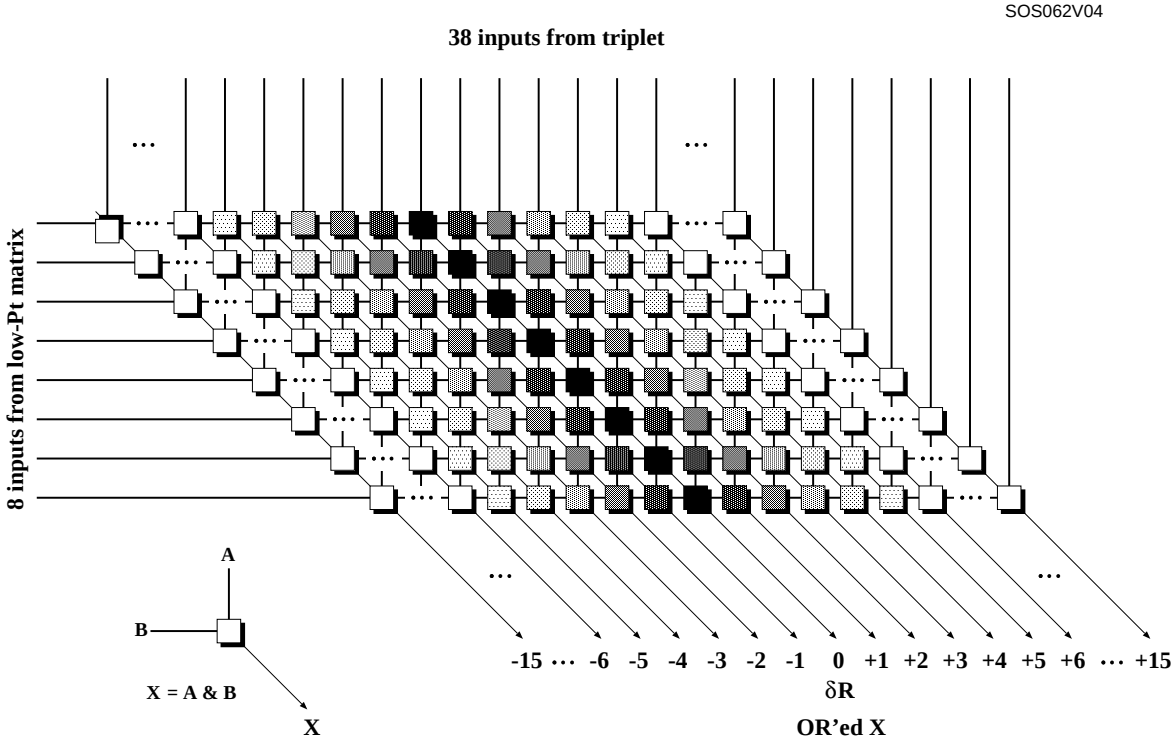

Figure 12-20 The detailed structure of the high-p_T coincidence matrix for wires. Shown is one of the eight blocks shown in Figure 12-19 together with the logic in an element. A similar structure exists for strips.

One coincidence element covers a single trigger subsector (RoI unit) as shown in Figure 12-6. The conditions for coincidence can be set independently for each trigger subsector. Furthermore, each matrix is fully programmable so that a range of threshold values can be provided. Outputs from the r-ϕ coincidence matrices are sent to the Preselector section. The Preselector selects the two highest-p_T tracks per pivot TGC. The output from the Preselector is a combination of position information (RoI number) and the six levels of p_T information to be sent to the Selector section. The Selector section then picks up the two highest-p_T tracks among the inputs.

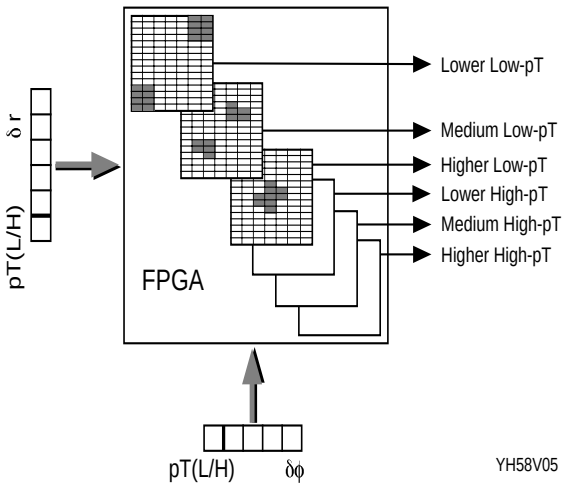

Figure 12-22 Detail of a matrix element in the Sector Logic shown together with the input format it receives.

In the End-cap region there are five pivot TGCs per End-cap sector and in the Forward region only one per sector. Thus up to ten tracks per End-cap and two tracks per Forward region are examined by the Track Selector. Two tracks from the possible ten are selected as the final two highest-p_T tracks in the End-cap, and in the Forward sector no reduction is made. Position information and the p_T value are combined and sent to the MUCTPI in a standard format [12-1].

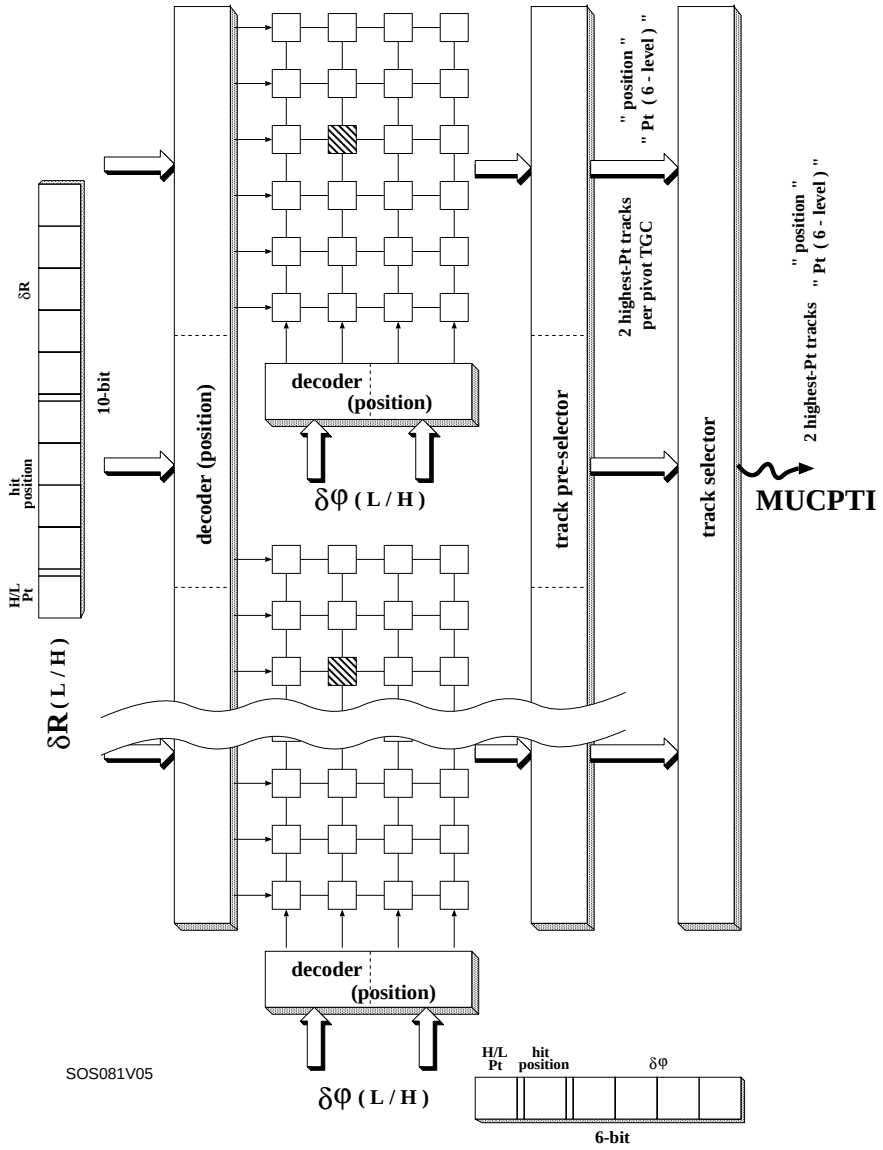

Figure 12-21 The block diagram of a Sector Logic board shown together with its input format from the High-p_T boards. Details of a matrix element, a shaded square in this figure, are shown in Figure 12-22.

12.4.6 Timing calibration

Between the bunch-crossing at the interaction point and the input to the coincidence matrices there are several signal / particle paths whose delays can be different for different detector channels. Signals from the chambers are synchronized to a bunch-crossing by the BC-ID circuit on the Patch-Panel. As described in Section 12.4.2 the phase of the clock can be adjusted with sub-nanosecond precision for each group of 16 channels from an ASD board (16 channels share the same cable bundle). This phase adjustment compensates for these differences so that all the signals originating from a given bunch-crossing are correctly assigned to it. This section describes how these delays are set. The calibration procedure is done after chambers have been mounted with their electronics on their frame to form a ‘set’, i.e. one third of an octant, in the work area on the surface prior to being mounted on the TGC support wheel in the cavern.

Setting the test-pulse timing: The goal of this first step is to produce a set of simultaneous output test-pulses from the ASDs that can be used to calibrate the trigger-logic. Each Patch-Panel has a test-pulse trigger input and a distribution tree with programmable delays to deliver this pulse to the ASD boards. Each ASD board distributes this test-pulse to all its channels. All 16 digital channels respond as if it were a chamber pulse; one of the 16 channels is equipped with an additional analog output. The first step of the calibration is to adjust the Patch-Panel test-pulse delays so that the digital outputs from all the ASD boards in the set are simultaneous. A test system generates a test-pulse trigger for all the Patch-Panels in the set under test, measures the analog output pulse timings by means of discriminators and TDCs, and sets the delays in the distribution tree until all analog outputs are in time. The in-time digital outputs can now be used to calibrate the down-stream logic timing.

Setting the bunch-crossing ID clock phase: This step compensates for the different cable lengths from ASD Board to the Patch-Panel and the different routing lengths on the Patch-Panel to the BCID circuit. By setting the mask registers in the Slave Board the ASDs for one chamber layer at a time can be read out. For each chamber layer, the test-pulse timing is stepped (in sub-nanosecond increments) through a range of times relative to the 40MHz bunch-crossing clock distributed to the BC-ID circuits. The BCID outputs are read via the Slave Boards for each step. For each ASD board there will be particular step where it comes in-time with respect to the 40MHz clock and its BCID outputs will go from zero to one. The time difference between this step and the clock edge is used to set the BCID clock phase for that ASD board. In this way all ASD signals will be clocked when they arrive.

The BCID clock phase is adjusted further to compensate for the time difference due to the variation in time-of-flight with radius. This time difference is calculated from the location of the wires and strips. In general the signal routings from chamber to ASD chip are equal length, but the BCID clock phase can be further adjusted for any known differences. On entry to the Slave Board, prior to the coincidence matrix and pipeline, there is a coarse delay for adjusting the timing in units of 12.5ns. This ensures reliable coincidences in the matrix and capture in the pipeline.

There remains one global time offset for the *whole set of chambers* – the phase of the local 40MHz clock with respect to the bunch-crossing at the interaction point. This can be set with tracks from collisions in the special LHC operating mode where after each populated bunch-crossing there are several empty ones so that it is clear which bunch-crossing gave rise to the tracks.

12.5 Constraints and rules for the trigger

Several rules are applied to the level-1 muon trigger for the TGCs with regard to the multiplicity of hits and tracks allowed in the system at any particular stage. In practice, these limitations have a negligible effect on the trigger efficiency and purity. They can be summarized by the following constraints, which are shown graphically in Figure 12-23:

1. Not more than two tracks per doublet Slave Board for wires are permitted. This limits to two the number of tracks allowed per 32 channels of wire-groups. Only the highest-p_T candidate in each 16-channel block is kept.
2. Not more than two tracks per doublet Slave Board for strips are permitted. This limits to two the number of tracks allowed per 32 channels of strips. Only the highest-p_T candidate in each 16-channel block is kept. (This is the same constraint as for the wires above.)

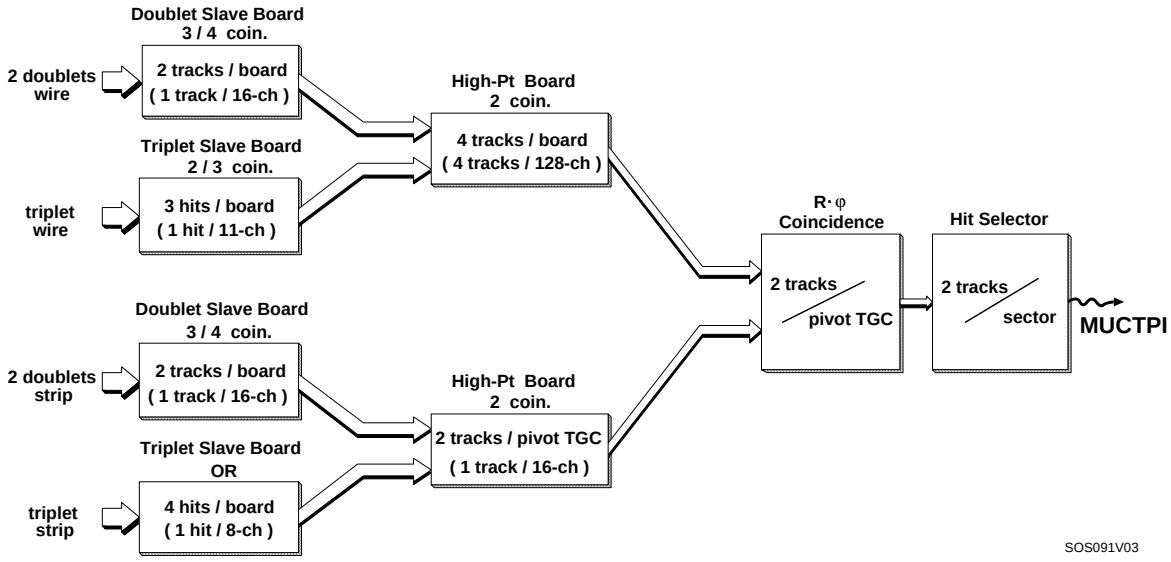

Figure 12-23 Imposed constraints on each section of the trigger-logic.

3. Not more than three hits per triplet Slave Board for wires are permitted, that is one hit per 11 channels.
4. Not more than four hits per triplet Slave Board for strips are permitted, that is one hit per eight channels.
5. In the encoding procedure following the centre-finding / declustering only one hit position per hit cluster is encoded.
6. Not more than four tracks are permitted per High-p_T board for wires. Thus four candidates per 128 channels are selected from a maximum of eight tracks. For strips, up to eight tracks are selected.
7. In the Sector Logic track preselector the two highest-p_T tracks per pivot plane are preselected from a maximum of 16 candidates.
8. In the Sector Logic track selector (second stage) the two highest-p_T tracks per sector are selected from a maximum of ten tracks.

12.6 Data read-out

All chamber hits and intermediate trigger logic results are read out for each trigger. Each of the 16 octants of the TGC has an independent read-out system from its Local Slave and High-p_T Boards through to its single Read Out Buffer (ROB) in the surface counting room. Each ROB receives data via an optical link from a single Read Out Driver (ROD). For read-out, the Local Slave Boards, which are mounted on the chambers, are grouped into Local DAQ Blocks, LDBs. Each LDB is comprised of a Star Switch near the detector that routes data from a number of Slave or High-p_T Boards, via an optical link to an intelligent I / O port (Local DAQ Master) in the ROD crate in USA15. This scheme is shown in Figure 12-24. Since the TGC occupancy is very low, the grouping of Slave boards into Local DAQ Blocks reduces the number of optical links.

The data read-out system described here conforms to the requirements set forth in Ref. [12-6].

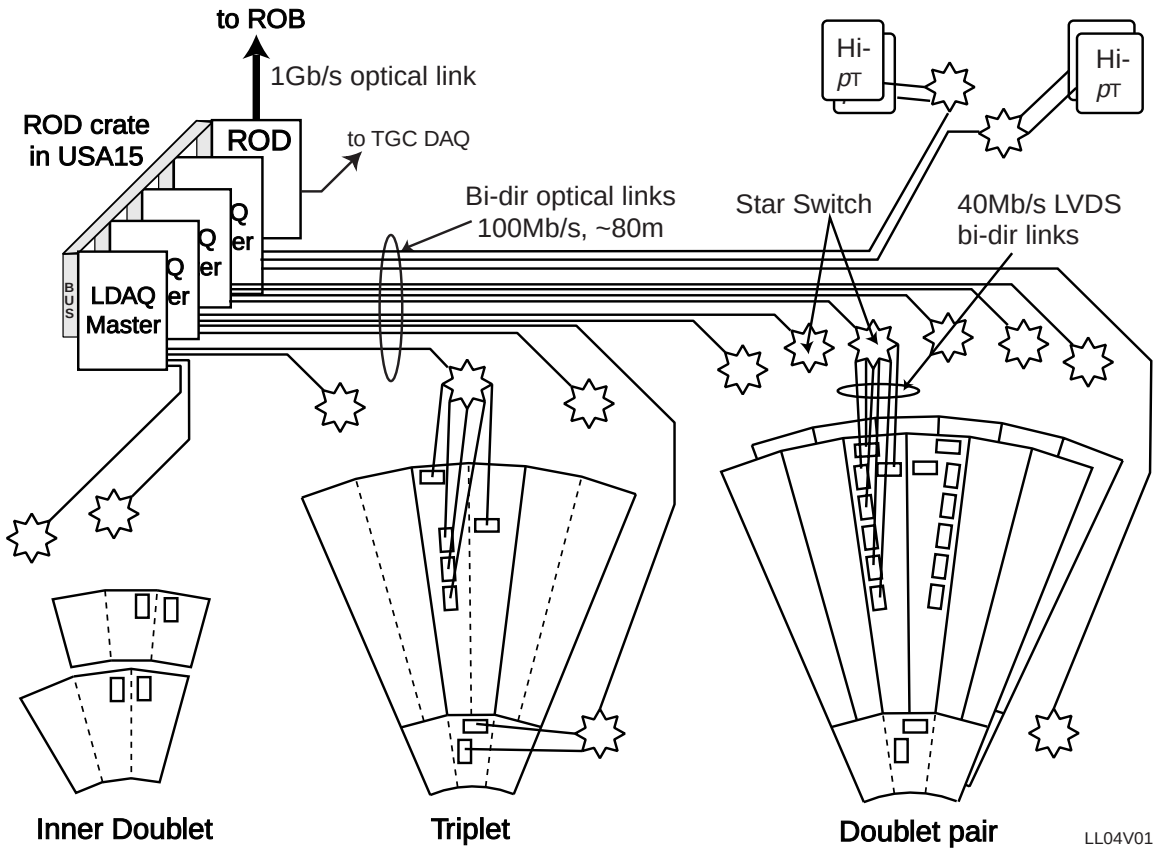

Figure 12-24 Overview of end-cap muon trigger chamber read-out for an octant – from chambers to the ROB. Chambers are grouped into Local DAQ Blocks shown by the solid lines. Only a few of the Slave Boards and High-p_T Boards are shown. Each Star Switch collects data from up to 32 LVDS links in its LDB.There is one ROD crate per octant. More than one LDAQ Master is expected per crate slot. (The dotted lines show, for the doublet and triplet, the projection of the trigger sectors on the octant; for the inner doublet, the chamber boundaries.)

LDBs are defined so that they contain Slave Boards from only one read-out layer of TGCs (i.e. triplet, doublet pair, or inner doublet) and so that LDB boundaries fall on one or more trigger sector boundaries. An LDB contains both wire and strip Slave Boards. There are two additional LDBs in an octant: the End-cap and the Forward High-p_T Board LDBs. Figure 12-24 shows the division of the triplet and doublet pair layers into four and seven Local Data Blocks respectively per octant. There are 15 LDBs of eight different types in an octant. Table 12-1 shows the number of slave boards and channels, and the expected data rates for each type of LDB. The number of LDBs, slave boards and channels and the data rates for each of the stations in an octant is shown in Table 12-2. The expected data rates are based on the ‘TP32’ background study [12-7] (See Section 14.7 and Figure 14-34). Note that after zero suppression (see Section 12.6.1) the expected data rates are modest, even when substantial safety factors are included.

Rather than a direct connection from each Slave Board to the ROD crate, the architecture presented here is hierarchical with a star switch as an intermediate data concentrator. This concentration enables more efficient use of optical links since the read-out is very sparse and the slaves are spread over a large area. The star architecture is preferred over a token ring or bus architecture because:

- It is not affected by the failure of any one link. This is important since the Slave Boards are not conveniently accessible.

Table 12-1 Characteristics of the various types of Local Data Blocks (LDBs) in an octant. Data rates are for a 100kHz level-1 trigger rate. Hits per event are uncorrelated background based on 'TP32' [12-7] plus measured electronic noise. Encoding is 48 bits/hit, as described in Section 12.6.1.

			Per Local Data Block				
LDB type		**Nº. per octant**	**Slave boards**	**Nº. Chan.**	**Raw (MB/s)**	**Hits/ event**	**encoded (KB/s)**
Doublet	**F**	1	21	2652	33.2	0.11	65.0
	E	6	19	2274	28.4	0.37	223.6
Triplet	**F**	1	18	1461	18.3	0.07	41.0
	E	3	24	2204	27.6	0.19	114.8
Inner	**F**	1	6	288	3.6	0.20	120.2
	E	1	3	252	3.2	0.13	75.8
Hi-p_T	**F**	1	9				
	E	1	30				

Table 12-2 Data flow rates and counts of LDBs, boards, and channels for the various stations in an octant. Data rates are for a 100kHz level-1 trigger rate. Hits per event are uncorrelated background based on 'TP32' [12-7] plus measured electronic noise. Encoding is 48 bits/hit, as described in Section 12.6.1.

		Per octant					
Station		**LDBs**	**Slave boards**	**Nº. Chan.**	**Raw (MB/s)**	**Hits/ event**	**Encoded (KB/s)**
Doublet	**F**	1	21	2652	33.2	0.11	65.0
	E	6	114	13644	170.6	2.24	1342.0
Triplet	**F**	1	18	1461	18.3	0.07	41.0
	E	3	72	6612	82.7	0.57	344.2
Inner	**F**	1	6	288	3.6	0.20	120.2
	E	1	3	252	3.2	0.13	75.8
Hi-p_T	**F**	1	(9)				
	E	1	(30)				
Total 1 octant		15	234	24909	311.4	3.31	1988.2
Total 2 sides		240	3744	398544	4981.8	53.02	31812.0

- Its point-to-point link logic is the easiest to implement in the Slave Board ASIC.
- The required band width can be handled by a conventional serial communication protocol.

The read-out design presented here is thought to be adequate for the expected data and message rates and to deliver the data to the ROB within 500ms after a level-1 trigger. A simulation of the data flow from front-end to ROB will be done to confirm this.

The read-out system consists of several component parts, each is described in the following sections:

1. Slave Board (Level-1 buffer, derandomizer, zero-suppression)
2. LS-Link
3. Star Switch
4. Local DAQ Master and Read-out Driver (ROD)

12.6.1 Level-1 buffer, derandomizer and encoding

The Slave Boards and High-p_T Boards all contain read-out functions. As shown in Figures 12-10,12-14,12-16 and 12-18, read-out sections for all Slave Boards and High-p_T Boards have basically the same structure. All chamber hits from TGCs (to DSB and TSB) or trigger data from DSB and TSB (to High-p_T Board) enters, after an adjustable delay circuit, the Level-1 Buffer of each board. The buffer is a pipeline clocked at 40MHz and is 2.5μs long. The FE_BCID counter is an 8-bit counter that counts clocks, i.e. bunch-crossings, and is reset by BCR (Bunch-crossing Counter Reset) once per LHC orbit. The FE-L1ID counter is an 8-bit counter that counts L1A (Level-1 Accept) and is reset by ECR (Event Counter Reset) at the beginning of a run. The data with 8-bits of FE_BCID plus 8-bits of FE-L1ID are stored in the derandomizer when the corresponding L1A signal arrives. For Slave Boards (DSB, TSB), the result from the trigger-logic is also stored in the Level-1 Buffer. For calibration, one L1A signal can initiate storing the data corresponding to up to three bunches (programmable). The full FE_BCID (Front-end Bunch-crossing ID), 12 bits, and FE-L1ID (Front-end Level-1 ID), 24 bits, are restored by the ROD. The low order bits stored here are used by the ROD as a synchronization check.

The derandomizer holds data corresponding to up to three bunch-crossings, that is, one before and one after the BC for which the L1A was received. The required depth of the derandomizer is determined by the maximum L1A rate of 100kHz, the maximum read-out time to the ROD, the values of M and T in the CTP flow-control constraint of no more than M L1As in T μsec (see Section 15.2.5) and the maximum tolerated data loss, as described in the Front-end Requirements Document [12-6].

Data compression is performed in the derandomizer. Without data compression the data rate would be about 3.3GB/sec per octant. Given the extremely low occupancy (see Tables 12-1 and 12-2), zero-suppression is desirable and several compression schemes are possible. For occupancies as low as presently expected, the simple encoding of the address of each hit is optimal. For higher occupancies with local correlations, better compression and simpler hardware can be attained by grouping wires or strips in cells (e.g. 8-bits) and, if the cell is non-zero, encoding the cell address and the cell's hit pattern. The optimal cell definition depends on backgrounds and is under study. Full read-out without compression will also be possible for debugging. Tables 12-1 and 12-2 show the expected data rates after encoding each hit in 48 bits (Slave Board ID: 12 bits, cell ID: 4, hit pattern: 8, BCID: 8, L1ID: 8, terminator: 8). This assumes the worst case of one hit per cell. Occupancies are calculated from the backgrounds simulated in 'TP32' [12-7]. The read-out system will be designed to cope with a data rate at least 10 times that shown. Figure 12-25 shows an example of a derandomizer circuit

The read-out data is transmitted from the Slave Board derandomizer via a serial cable link, called the Local–Slave Link (LS-Link), to the Star Switch concentrator. Each slave sees the switch as a simple point-to-point connection via the LS-Link. On receiving a read-out command

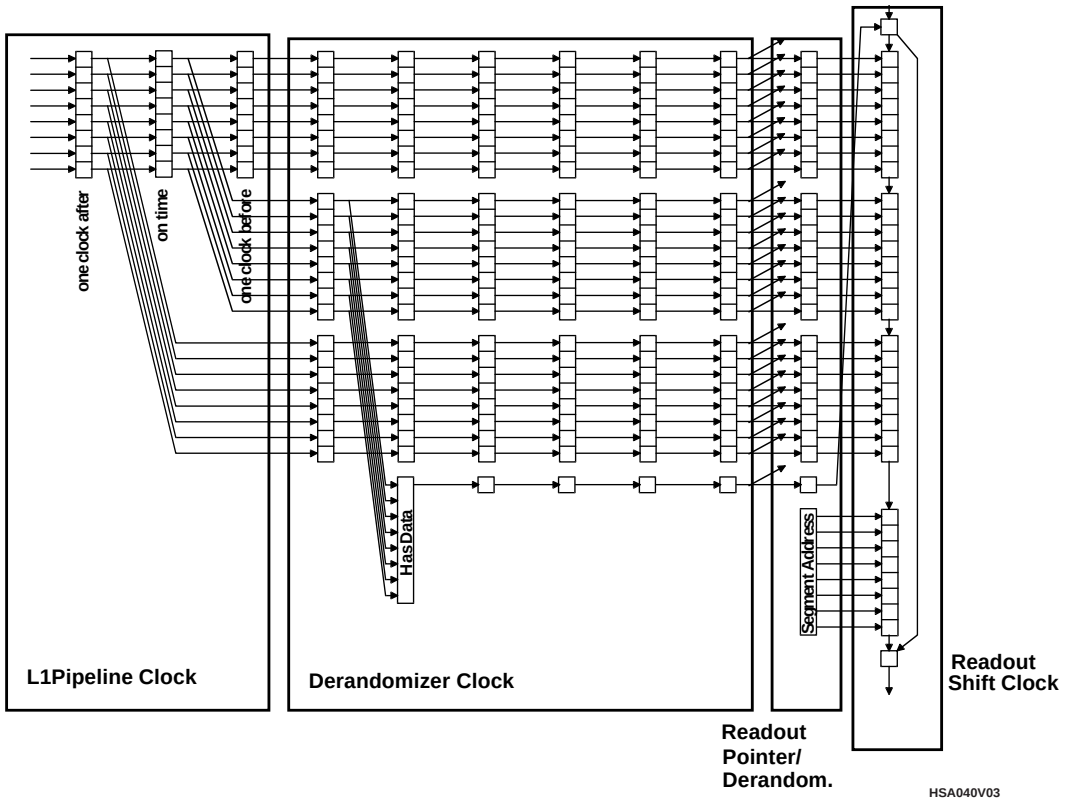

Figure 12-25 A possible derandomizer circuit. The white boxes are flip-flops. Each enclosed area is fired by a different clock. The left-most area is a part of the LVL1 buffer and is clocked at 40MHz. The hit patterns for three consecutive clocks can be strobed into the derandomizer block. In the first stage of the derandomizer block, existence of non-zero data can be checked by a simple OR operation. The result of the check determines whether this word of data should be skipped or transferred to the read-out shift register.

from the master, each slave starts its data transmission by synchronizing the data transfer with a clock which is generated by the slave itself. On completion of transmission, a termination code is sent. If a slave has no data, it will transmit only a no-data signal.

12.6.2 Local Slave Link

The local slave link (LS-Link) is implemented as a pair of uni-directional digital connections using the LVDS standard. Each one-way connection consists of three signal lines: CLOCK, DATA and SYNC. The SYNC signal initializes the sequence of the interface circuit. By indicating the start and stop bits of a packet, the frame of the packet can be confirmed and, as a result, both the reliability of the transfer is improved and the transmitter and receiver logic become much simpler. For example, by detecting SYNC=0 and DATA=0 (Start bit) at the clock edge, the receiver logic resets the sequence, and then, by detecting SYNC=0 and DATA=1 (Stop bit) at the clock edge, the receiver logic starts execution of the received command. The LS-link packet frame is shown in Figure 12-26. The clock is generated from the 40MHz LHC clock. These kinds of on-chip control signals can be generated by simple logic operations on hardware signal lines.

Each slave has two connections, one for receiving commands from the Local DAQ Master, and the other for sending data to the Master. A 14-pin header connector is used for the physical connection. The size of the data depends on the contents of the message. There are commands

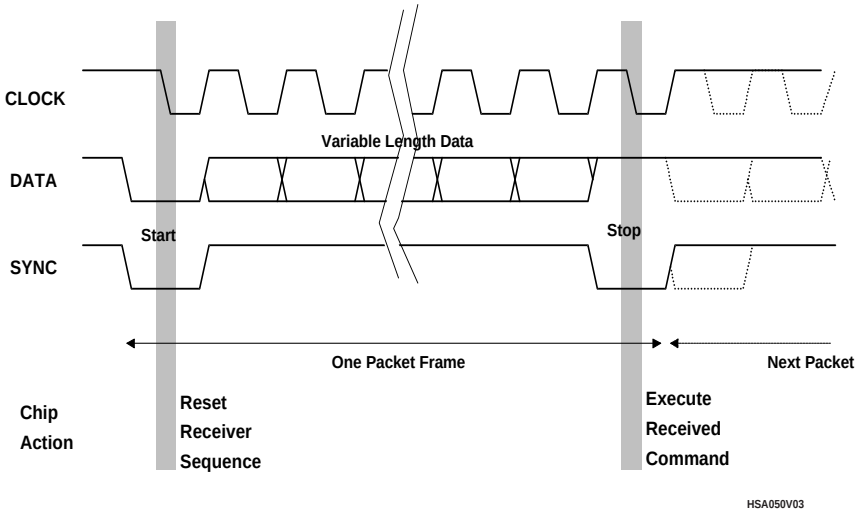

Figure 12-26 Serial communication signal frame for the TGC Local Slave link. Three signal lines form a unidirectional connection. The CLOCK line gives the decoding timing for receivers. The SYNC and DATA line combination offers a data frame which can be decoded by very simple receiver logic: START=(DATA=0 & SYNC=0) and STOP=(DATA=1 & SYNC=0).

to start or stop the transfer and extended commands to download parameters, to verify downloaded data, or to read-out 'snapshots' implemented in chips for diagnostics. As the LS-Link is connected to all the trigger chips, this link can be used for initialization, test and diagnostics when the system is not in data-taking mode. The format of the Slave Board event data can be seen in Figure 12-29 as the 'Module Data Format'.

12.6.3 Star Switch

Each Local DAQ Block of up to 32 Slave Boards is serviced by its own Star Switch situated near the LDB. There are 15 Star Switches (15 LDBs) per octant. Each is connected to a Local DAQ Master (described below) in the ROD crate in USA15 via a bi-directional optic link. Figure 12-27 shows a block diagram of the Star Switch. The command path from the Local DAQ Master is used to control both the switch and the slaves. The switch contains a FIFO which allows headers and trailers to be easily added and to change the protocol from that used on the LS-Link.

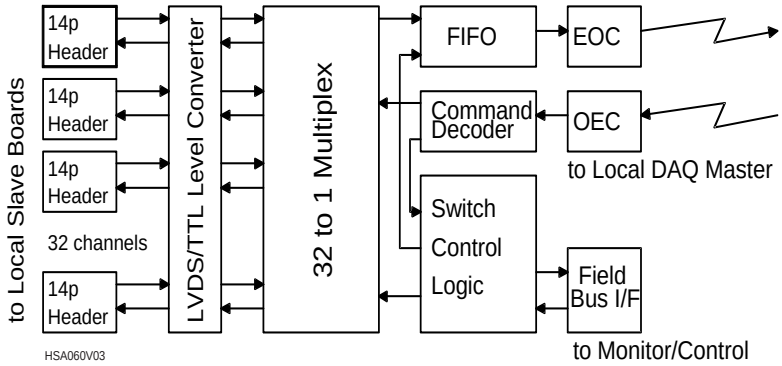

Figure 12-27 Block diagram of the Star Switch. Up to 32 Slave Boards can be connected to one switch module. Data from slaves are stored in the FIFO and then transmitted to the local DAQ master via a fiber link. Messages from the master will be partly decoded and used for control of the switch.

On average, at the 100kHz level-1 trigger rate, the switch must poll up to 32 Slave Boards and transfer two 48-bit hits (the worst case in Table 12-1 times a safety factor of five) in 10μsec. This can be done with 40MHz clocking – provided that the Star Switch controls the sequencing, and not the Local DAQ Master for which the cable delay is too long.

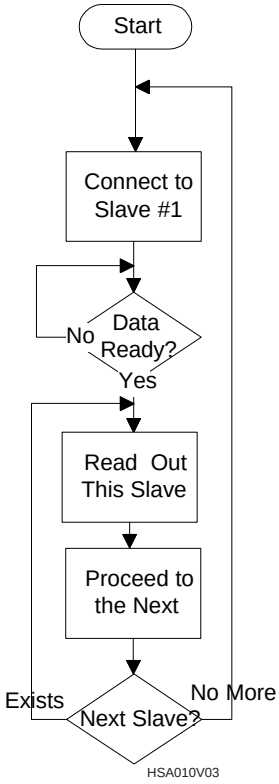

Figure 12-28 Flow chart of the read-out sequence executed by the Star Switch. The logic has a map register which indicates which slaves should be read. Slave #1 means the first slave indicated in the map.

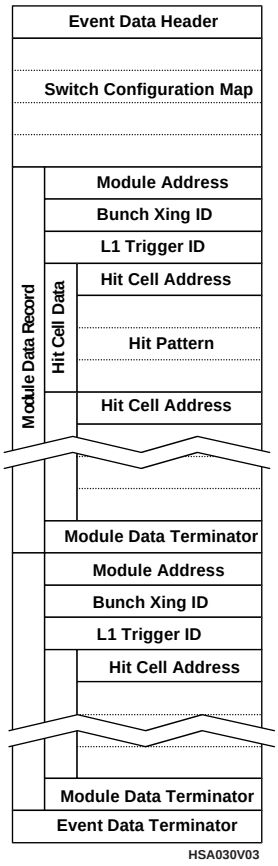

Figure 12-29 An example format of data sent from a Star Switch to its Local DAQ Master. Each Slave Board is a 'module'. The switch configuration map is checked for consistency with the module data records.

The Star Switch sequence for slave module read-out is shown in Figure 12-28. The sequencer implemented in the switch waits for the first slave. When the data set is ready, it will be transmitted by the slave to the switch. On receiving the termination code from the first slave, the switch disconnects and reconnects to the next slave and requests the data transfer. By repeating this procedure, all the data in a single event is transferred to the switch. A register in the switch is configured with a map that indicates which slaves should be connected. Slaves which are broken, unnecessary or that have timed-out are removed from this map. The corresponding map is inserted into the output stream for each Level-1 Accept. By checking the consistency between the map and the slave addresses in the data and with its own expected map, the ROD confirms that the data transfer is being performed properly. Figure 12-29 shows the data format transmitted to the Local DAQ Master.

The Star Switches will either be situated in the Hi-p_T crates on the periphery of TGC wheel M1, as shown in Figure 12-4, or mounted like a Slave Board on Patch-Panels. It will be implemented as a 6U VME board which can be made compatible with either, or perhaps both, options. For the

small number required, it may be more convenient to implement the Star Switch with a radiation tolerant anti-fuse FPGA instead of an ASIC.

12.6.4 Local DAQ Masters and ROD crate

The primary function of the ROD crate complex (shown in part of Figure 12-24) is to collect the data from the Local DAQ Blocks so that a data record for a complete octant can be built and sent to the ROB. Even though the data volume is modest, a difficult requirement is to handle up to 100kHz of messages on each of 15 links. This requires intelligent IO processors, called Local DAQ Masters (LDMs) in this system. Each LDM collects the data from one Local DAQ Block via its Star Switch. As part of this collection process the following operations must be performed before sending the data to the ROB:

- Initiate the event transfer from the Star Switch for each Level-1 trigger;
- Verify the record structure transmitted, warning when expected Slave Boards are missing;
- Verify the proper sequence of the 8-bit Front-end event IDs and BCIDs and replace them with the full event ID and BCID;
- Use the event ID and BCID to match event fragments to be built into an event;
- Translate the mixed encoding of cell address plus cell hit patterns to full hit addresses;
- Collect statistics on data rates and occupancies;
- Store the data for the ROB in a transmit buffer shared with the other LDMs, reformating it into the standard ROB input format.

Such processing and buffer management at these rates is suitable for FPGA-based processors and it is planned to implement the Local DAQ Master as an FPGA processor dedicated to a bi-directional IO link. The ROD itself would be a standard commercial processor whose function would be to control the Local DAQ Masters and the ROB link, to gather statistics, to communicate with Run Control and to handle exceptional situations.

The input link to the LDM from the Star Switch is an optical fiber link. Our bandwidth requirements are modest, but the cavern end of the link must be radiation tolerant. We will use, if possible, the same link as used elsewhere in ATLAS in order to share radiation tolerance certification, spare parts, and the incorporation of the link hardware into a form meeting a standard such as CERN's S-link. It may be possible to take advantage of the fact that although our link is bi-directional, the bandwidth needed from the LDM is much less than that into the LDM.

The events are built in shared memory, accessed over a bus, by LDMs and the ROB output link. The expected zero-suppressed data rates for an octant (see Section 12.6) can be handled by a backplane bus such as PCI or VME.

Since the muon precision chamber (MDT) read-out will also probably be segmented in octants, an end-cap trigger ROB can be put in the same ROB crate as the several MDT ROBs containing data corresponding to the muon end-cap trigger. Having corresponding MDT and TGC data in the same ROB crate should simplify level-2 trigger operations. Also the one ROD per octant matches the ROD segmentation of the Muon Barrel Trigger chambers.

12.7 System interconnections

Table 12-3 summarizes the interconnections between the parts of the TGC trigger electronics.

The 24,528 ASD Boards sit on the TGC, inside the shielding cage, and the Patch-Panels sit on the surface of the TGC wheel. The connection is thus both short and straightforward in routing.

The 624 High-p_T Board to Sector Logic connections are made by optical fibers running from the rim of the M1 TGC wheel (triplet station) to USA15. The fiber's routing will be optimized to minimize the trigger latency. A preliminary estimate by the ATLAS Technical Coordination Group is 80m.

12.8 Timing and trigger distribution

The distribution of the timing and trigger signals around the ATLAS detector is done through a standard Timing, Trigger and Control (TTC) system [12-8], using optical fiber signals. Signals from the optical system are decoded for local use by the standard TTCrx chip. For the level-1 trigger electronics mounted on the TGCs, the following four signals from the TTCrx will be used:

- 40MHz Bunch-crossing Signal (Clock)
- Level-1 Trigger Accept (L1A)
- Bunch-crossing Counter Reset (BCR)
- Event Counter Reset (ECR)

In addition to these four fundamental signals, the following are used by the ROD:

- Bunch-crossing number (12 bit)
- Event number (24 bit)
- Broadcast commands

The low-order byte of the BCID and event numbers are compared with those counted independently in each Slave Board and the high-order bits are added by the ROD. This checks that all data coming into a ROD have the appropriate event number and BCID.

In order to effectively distribute signals over all on-detector electronics (Patch-Panel, Slave Board and High-p_T Board), plus the ROD and Sector Logic, a large number of TTCrx chips are required. There will probably be two TTC partitions, one per end-cap. In total 1496 TTCrx chips will be used as follows.

- TTCrx on each Patch-Panel (total 1376): Decoded TTC signals are distributed to all ASICs on Patch-Panels and to Slave Boards through equal-length paths.
- TTCrx in each High-p_T crate (total 96): Decoded TTC signals are distributed to all High-p_T Boards in a crate through equal-length paths.
- TTCrx in Sector Logic rack (total 8): Decoded TTC signals are distributed through equal-length paths.
- TTCrx in ROD crate (total 16): A single TTCrx board in each ROD crate will be used.

Table 12-3 Details of the system interconnections of the TGC trigger system. 'n TP' is a cable of n twisted pairs.Note that up-to-date numbers for this table are found in [12-2].

Source	Destination	Destination location	Connection type	Number of connections
I. Trigger data flow				
ASD board	Patch-Panel	chamber surface	20TP– AWG28	
– two doublets				32 bundles per PS-pack
– triplet (wire)				24 bundles per PS-pack
– triplet (strip)				12 bundles per PS-pack
Patch-Panel	Slave Board	PS-pack	DIN41612 connector	Two per Slave Board
Slave Board	High-p_T board	Outer rim of triplet wheel		
– Doublet wire			18TP – AWG28	102 bundles per octant
– Doublet strip			18TP – AWG28	33 bundles per octant
– Triplet wire			16TP – AWG28	75 bundles per octant
– Triplet strip			20TP – AWG28	15 bundles per octant
High-p_T board	Sector Logic	Crate in USA15		
– wires			Optical fiber ~80m	27 links per octant, each 40-bits × 40MHz
– strips			Optical fiber ~80m	12 links per octant, each 48-bits × 40MHz
Sector Logic	MUCTPI	Crate in USA15	Electrical cable	One cable per trigger sector, 9 connections per octant
II. Data acquisition data flow				
Slave Board	Star Switch	High-p_T crate on rim of wheel	8TP – AWG28	13 Star Switches per octant 234 connections per octant
High-p_T board	Star Switch	High-p_T crate on rim of wheel	8TP – AWG28	39 High-p_T boards per octant, 2 Star Switches per octant 39 connections per octant
Star Switch	Local DAQ Master	USA15	Optical link	2 optical links per connection 15 Star Switches per octant
Local DAQ Master	ROD	In same crate	Backplane	1 crate per octant for a set of Local DAQ Masters and a ROD
III. TTC signal distribution				
TTC crate	TTCrx on Patch-Panel		Optical link	86 optical links per octant, 1376 in total
TTC crate	TTCrx in High-p_T crate		Optical link	6 optical links per octant, 96 in total

Table 12-3 Details of the system interconnections of the TGC trigger system. ‘n TP’ is a cable of n twisted pairs.Note that up-to-date numbers for this table are found in [12-2].

Source	Destination	Destination location	Connection type	Number of connections
V. Low-voltage power supply line				
Doublet	MDT rim to Patch-Panel		AWG0 cable	48 cables per octant
Triplet	MDT rim to Patch-Panel		AWG0 cable	48 cables per octant
	MDT rim to High-p_T crate		AWG0 cable	12 cables per octant

12.9 Latency

The total estimated latency of the end-cap muon trigger system from interaction point through to the input to the MUCTPI is 1.15μs. The contributions to this latency are listed in Table 12-4. Adding the additional 12 bunch crossings for the MUCTPI and CTP, plus 20 bunch crossings for the 80m cable back to the TGC front-end, gives a total latency of 1.95μs. This meets the requirement of <2μs and provides the full 0.5μs contingency allocated within the 2.5μs absolute maximum latency. See Chapter 18 for a further discussion of the trigger latency.

The contributions for the electronics components are based on the number of pipeline stages in the block level design. The cable delays are from an evaluation of cable lengths by the ATLAS Technical Coordination Group.

12.10 Environmental issues

In the design of the TGC trigger electronics the difficult radiation and magnetic-field environments in the ATLAS cavern must be considered to produce a trigger design that will survive and perform for the required ten or more years of LHC running.

One solution to radiation or field intolerance is to move electronics out of the cavern, and this has been done where possible. All the electronics boards prior to the Sector Logic, however, are placed in the vicinity of the TGC wheels, where both radiation and magnetic-field levels are high. A similar situation is faced by the low-voltage power supplies although they will be at the outer-most radius of the TGC wheels. The electronics of the Sector Logic and the high-voltage power supplies are housed in the USA15 area where the radiation and field environment is benign, as is all read-out electronics beyond the Star Switch.

Bipolar and sub-micron CMOS gate arrays will be used in those circuits placed within the ATLAS cavern, along with commercial off-the-shelf components that are certified for use in such an environment. No radiation or field consideration is necessary in the USA15 area.

Table 12-4 Contributions to the total estimated TGC latency, in bunch-crossings (25ns each).

Process/transmission	Time required (# of bunch-crossings)	Accumulated time
TOF to TGC	3	3
TGC response	1	4
ASD	1	5
Cable to Patch-Panel	1	6
Bunch-crossing ID and OR	2	8
Cable to Slave Board	1	9
Delay Adjust	1	10
3/4 or 2/3 coincidence	3	13
Cable to high-p_T coincidence	3	16
Delay adjustment	1	17
High-p_T coincidence matrix	4	21
Cable to USA15 (80m)	16	37
Sector Logic processing	8	45
Cable to MUCTPI (5m)	1	46
Total delay sum	46BC	1.15μs

12.10.1 Radiation tolerance

It is clear that much of the volume of the ATLAS detector will be in an extremely harsh radiation environment. The region of the TGC wheels where the end-cap muon trigger electronics come closest to the interaction point is not the most extreme region, but still offers challenges that must be met to ensure the long-term performance of the trigger.

Using data from a simulation of the expected radiation environment in ATLAS [12-9], radiation levels at the TGC wheels are summarized in Table 12-5.

Table 12-5 Typical radiation levels per year in the regions where TGC electronics will be located. The two locations correspond to the most and the least extreme in terms of radiation environment.

	Worst Location	Best Location
Neutrons, n/cm^2	9.7×10^{10}	3.1×10^{10}
1 MeV Neutron equivalent/cm^2	2.0×10^{10}	3.6×10^{9}
Dose , Gy/yr	6.2×10^{-1}	2.1×10^{-1}

A safety factor of four is required [12-9], so it is intended that all electronics placed on the TGC wheels will be tested to such a radiation level. In view of their higher sensitivity to radiation damage, for bipolar transistors an additional factor of 1.5 will be demanded for the neutron flux and a factor of five for the ionizing rate.

Since the trigger is required to perform for a decade in such conditions an additional factor of ten will be demanded on the total dose that can be tolerated.

The radiation tolerance criteria to which the wheel-mounted electronics should conform, for 10 years running, are shown in Table 12-6.

Table 12-6 Radiation tolerance criteria for chamber-mounted electronics

	CMOS	Bipolar
Neutrons, n / cm^2	3.9×10^{12}	5.9×10^{12}
1 MeV Neutron equivalent / cm^2	8.1×10^{11}	1.2×10^{12}
Dose, Gy	2.5×10^{1}	1.3×10^{2}

Bipolar technology will be used for the ASD chip and sub-micron CMOS technology for the trigger and read-out gate arrays. Such technologies have already been investigated for radiation tolerance [12-10]. Special radiation tolerant processes will not be utilized due to problems with both availability and cost. The standard-process ICs that we plan to use will be checked for radiation tolerance level, particularly the bipolar ASD chip. Some commercial components will be used since the TGC electronics cannot be designed without them. However all components will be confirmed to conform to the above ATLAS specified radiation tolerance criteria, either by direct radiation tolerance tests or by using components certified for the appropriate dose level.

12.10.2 Magnetic field

The magnetic field can impose constraints on the electronics that can be chosen for those components that sit close to the toroids, notably those mounted on the chambers themselves. In the region of the TGC planes this field is strong enough to require careful design and selection of electronics components. Simulations suggest that the field strength in the region of the TGC wheels will span the range 100 to 1000 Gauss [12-4].

Since the ASD Boards are part of the chamber package, their location is constrained. However, for other electronics, we can choose the location of the electronics where magnetic field strength is of the order of a few hundred Gauss. One area where problems might be expected is with power supplies containing transformers or chokes. These will be positioned very close to the cavern walls, in the cavern regions with the lowest field. Tests have been performed and low-voltage power supplies selected that operate without problem in fields of up to 500Gauss. Given the field strength expected, and the possibility of building additional magnetic shielding for these power supplies, it is unlikely that magnetic field will affect the performance of such supplies. Additionally we are following the CERN-wide effort that has been initiated to study the performance of various power supplies in strong magnetic fields.

Where it is unavoidable that electronics will be subject to strong magnetic fields (which is the case for some of the chamber-mounted boards) we will avoid using components that are known to be affected adversely by the magnetic field. For example, choke coils for noise filtering will be replaced by alternative means of filtering. Tests will be performed on all electronics intended for the strong field region to ensure their operation in such magnetic field.

12.11 Control, configuration and monitoring

12.11.1 Detector control system

Throughout the ATLAS experiment the Detector Control System (DCS) is used to monitor and control the detectors. The TGC Detector Control System (TGC DCS) is a subsystem of the ATLAS DCS, and as such complies with its requirements and specifications. It is based on the ATLAS standard CAN field bus as a backbone and a distributed architecture of Local CAN nodes. An overview of the general structure of the TGC DCS is presented in Figure 12-30. A Local CAN node (LCANN) provides the interface between a Local Control Monitor Board (LMB), situated in each Local DAQ Block, and the CAN field bus. The LMB performs all the local functions related to the monitoring and control of the TGC counters as well as the ASD Board, Patch-Panel and Slave Boards.

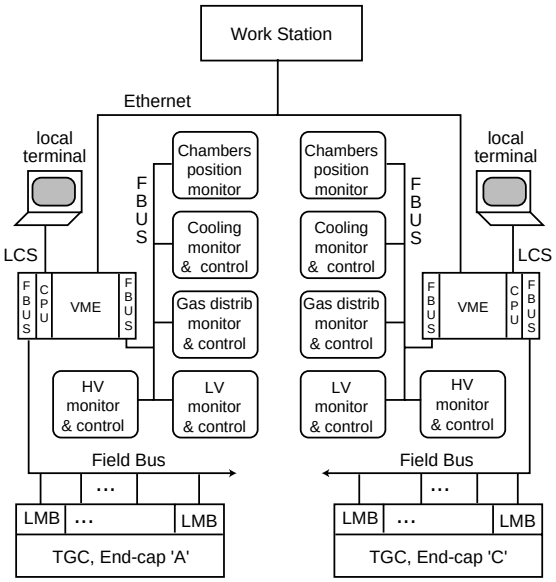

Figure 12-30 The general structure of the TGC Detector Control System.

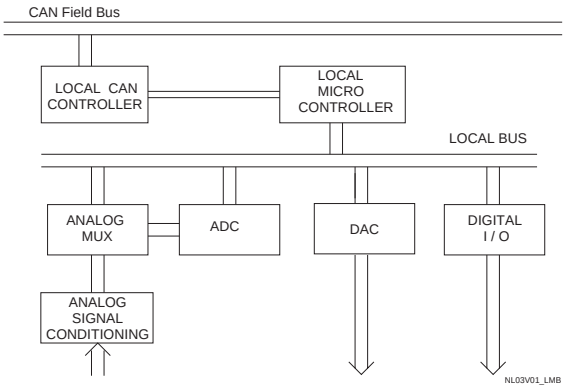

Figure 12-31 A typical example of the architecture of an ATLAS DCS Local Control Monitoring Board (LMB).

The monitoring and control functions comprise: temperature monitoring at the ASD Board and PS-pack level, DC supply voltages and current monitoring, generation of threshold voltages for the discriminators on ASD chips, monitoring of gas flow and pressure for the TGC detectors, relative displacement monitoring of doublet and triplet units, and monitoring of TGC analog signals. To perform these tasks, the LMB architecture comprises: an Analog Multiplexer (AMUX), Analog-to-Digital Converters (ADC), Digital-to-Analog Converters (DAC), Digital Registers (DREG), and microcontrollers, as well as specialised sensors.

The LMB is powered independently from the rest of the electronics on the TGC in order to be able to perform monitoring during shutdowns of the TGC detector. A typical ATLAS DCS LMB is shown in Figure 12-31. Table 12-7 shows the total number of TGC DCS channels.

Table 12-7 Operating and environmental parameters monitored by the TGC DCS.

Item	Multiplicity	Total channels
Chamber		
ASD temperature	1 / ASD board	25968
Faraday cage temp.	2 / unit	3168
Patch Panel		
-3.0V power OK	1 / Patch-Panel	1376
+3.0V power OK	1 / Patch-Panel	1376
+3.3V power OK	1 / Patch-Panel	1376
Temperature	2 / Patch-Panel	2752
Slave Board		
Temperature	1 / Slave board	3936
Hi-p_T Crate		
Temperature	3 / Hi-p_T crate, 3 crates / octant	144

Item	Multiplicity	Total channels
HV		
HV voltage	1 / chamber	3600
HV current	1 / chamber	3600
HV trip recovery	1 / chamber	3600
Gas		
Diff. gas pressure	1 / gas channel	288
Mass flow	2 / gas channel	576
Gas pressure	2 / gas rack	16
Mass flow	2 / gas rack	16
Mixture monitor	2 / gas rack	2
Oxygen monitor	1 / end-cap	2
Alignment		
Linear potentiometer	2 / unit	3168
Position measure	9 / octant	144

12.11.2 Control and monitoring of the trigger electronics

Table 12-8 lists the data items that are used to control, configure and monitor the trigger system. If the CAN field bus can provide the needed bandwidth (under investigation), it will be used to implement the trigger control and monitoring system. In this way connection hardware, controllers, and software can be the same as those for the DCS system. Interfaces will be purchased or produced to interface between CAN and JTAG, I2C and any other serial protocol needed locally by the trigger electronics. Using the same CAN technology also would enable individual functions to be moved from the TGC trigger control network to the DCS network depending on the final scope of the DCS system. Each octant is separately controlled from the ROD crate. It is not known yet whether the ROD crate–Star Switch link will carry the fieldbus protocol to a CAN controller near the Star Switch or if the controller will be in the ROD crate with a dedicated link per octant to the cavern.

Table 12-8 Control, configuration and monitor registers.

	channels	purpose	R/W
ASD chip			
Discriminator threshold set	7200	configure	RW
Analog test-pulse pulse-height	3600	calibrate, diagnose	RW
Analog test-pulse trigger	3600	calibrate, diagnose	W
Wire analog output from ASD	3600	calibrate, diagnose	R
BCID chip			
Mask register	11008	configure	RW
Test strobe	11008	diagnose	W
BCID phase	~25000	configure	RW
BCID gate width	~25000	configure	RW
Patch Panel			
Patch panel ID	1376	set up	R
Analog test-pulse trigger for ASD	1376	calibrate, diagnose	W
TTCrx chip	1376	configure	RW
Slave Boards			
Mask / set register	3744	configure	RW
Read-out span: 1 or trigger ±1BC	3744	configure	RW
Slave and Hi-pT Boards			
Input delay length	29808	calibrate	RW
Matrix input snap shot registers	3744	monitor, diagnose	RW
Matrix output snap shot registers	3744	monitor, diagnose	RW
Snap shot mode: load / OR / AND	3744	configure	RW
LS-link output enable	3744	configure	RW
L1 data pipeline length register	3744	set up	RW
LVL1 matrix output pipeline offset from LVL1 data pipeline	3744	set up	RW
Star Switch concentrator			
Port enable register	240	set up, diagnose	RW
Other concentrator registers	240	set up, diagnose	

12.12 Power, grounding and cooling

12.12.1 Low-voltage system

The total power requirement for all on-chamber TGC electronics is estimated in Table 12-9. The total consumption is 110 mW per channel. Power supplies for electronics in USA15 (Sector Logic, Local DAQ Master, ROD, and so on) will be treated separately, and anyway offers few problems in terms of power dissipation owing to the off-detector rack-mounted location. +3V and –3V are required by the ASD board, and +3.3V for all other digital components. Allowing for inefficiencies of 10% for transmission loss, line-noise reduction and other contingency, 110 mW per channel is the total estimated power requirement.

Table 12-9 Power consumption required for on-chamber TGC electronics

ASD	46 mW / channel
LVDS receiver on Patch-Panel for the ASD signals	13 mW / channel
Patch-Panel logic	20 mW / channel
Slave Board	10 mW / channel
High-p_T board	5 mW / channel
TTC, etc.	5 mW / channel
Read-out	0.5 mW / channel
DCS	0.5 mW / channel
Others (regulator loss, line drop, contingency, etc.)	10 mW (10%)
Total	110 mW / channel

Given an estimate of 400,000 channels in total we expect an overall requirement of approximately 45 kW power to be supplied to the on-chamber TGC electronics. To reduce the loss in low-voltage power supplies (LVPS), switcher-type power supplies will be used with a careful design to reduce the line noise and EMI noise. A switching power supply of approximately 800 W output was tested under magnetic field strength of up to 500 Gauss and it was verified that it works without problems. Since the expected magnetic field strength in the environment of these LVPSs is around or below 200 Gauss, we see no problem on this account. The radiation environment will be considered by utilizing components that the ATLAS Collaboration has confirmed are radiation tolerant. We will perform tests of the LVPS in appropriate radiation conditions to confirm that the supplies will survive in the expected radiation environment.

LVPSs of approximately 1 kW output power will be placed along the edge of each MDT wheel where eight electronics racks of height 2.5 m are placed, totalling 16 racks. Hence one rack corresponds to one eighth of the TGC wheels for the LVPSs. Each rack contains three of these LVPSs so that a total of 48 LVPSs are used. This should be sufficient for the TGC electronics system. Power routing will be done through heavy copper lines to the TGC octant block. Lighter lines will supply power to each Patch-Panel from where power will be distributed to ASD and Slave Boards. Assuming an efficiency of 85% for the LVPSs, 180 W per LVPS will be dissipated or 540 W per rack in total. This will be cooled by heat-exchanger and cooling-fan units.

12.12.2 System grounding and shielding

The TGC system adheres to the ATLAS grounding policy outlined in [12-11].

Local chamber ground: Thin gap chambers are constructed as either doublets or triplets. In both cases, ground planes are located outside the sensitive volume, enclosing the wire and strip planes. These ground planes, together with U-shaped copper shields around all four sides of the trapezoidal TGC unit, form a Faraday cage. Paper honeycomb panels of 5mm thickness are glued onto both sides of the ground plane. The prime functionality of the honeycomb is to sustain the gas pressure, but it also serves as an electrical insulator. The Faraday cage enclosure of each unit serves as its local ground.

Global grounding: TGC units are mounted on aluminium support frames forming a ladder structure. The Faraday enclosure, which is the chamber ground, is electrically insulated from the support. Within a ladder of TGC counters, Faraday enclosures are connected to a ground strap which runs along the ladder in the radial direction. The strap is a sizeable copper braid so as to offer a low-impedance path. These straps are connected to a thicker copper bar or braid, itself connected to a detector-wide ground. A overall TGC grounding diagram is shown in Figure 12-32.

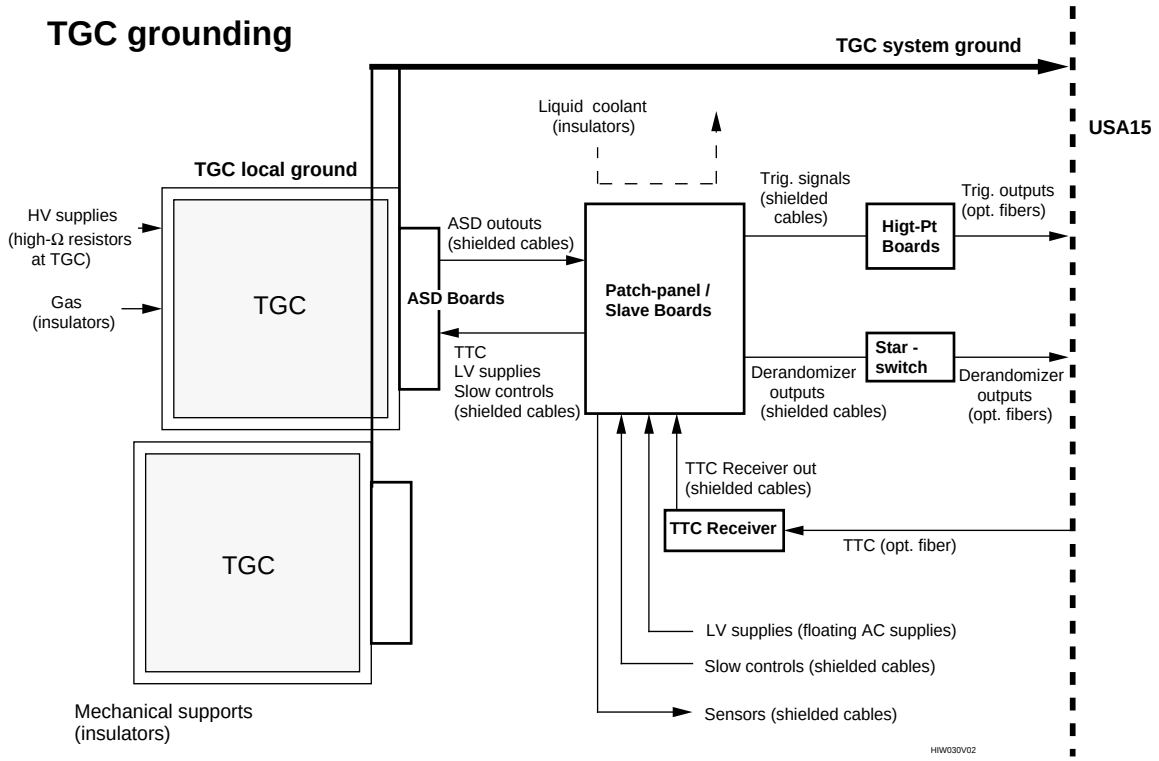

Figure 12-32 A schematic diagram of the grounding scheme for the TGC trigger.

12.12.2.1 Grounding detail

High voltage: High voltage is supplied to each TGC at the opposite side of the chamber with respect to the front-end read-out electronics. The high-voltage ground is connected through a series resistor to the local chamber ground. The series resistor, of the order of several kilo-ohms, serves to break potential ground loops. The high-voltage ground at the supply is not floating for reasons of safety.

On-chamber read-out electronics: ASD Boards are mounted on each chamber on two sides of the trapezoid: one side is used for the wire read-out, the other for the strips. Each ASD Board is

enclosed in a shielding box which is connected to the local chamber ground through copper straps. The ASD LVDS signals are sent through shielded twisted-pair cables to Patch-Panels, which are mounted on the TGC support frame together with the Slave Boards as a single PS-pack unit. The ground of the PS-pack is electrically isolated from the support. The cable shield is connected to the chamber ground at one end, and is connected through a capacitor at the other end to break any ground loops.

On-detector electronics: Electronics modules which are not directly mounted on the TGC counters but on the TGC support frame constitute the on-detector electronics. (These are the Patch-Panels, Slave Boards, High-p_T Boards, Star Switch boards, and TTC receivers.) These modules are enclosed in shielding boxes or mini-crates. Shielded cables are used between the on-detector electronics modules unless they are contained within the same shielding box. Optical links will be used between the on-detector electronics and USA15 for trigger signals, Star Switch outputs and TTC signals.

Low-voltage supplies: Low-voltage supplies mounted at the highest radius of the MDT wheel will be allowed to float. The power return is connected to the Faraday enclosure at the chamber, but not to the safety ground at the power supply. A safety ground device as recommended for use ATLAS-wide will be implemented in the power supply.

Detector control: Slow-control signals including sensor outputs and monitor signals will be sent from or to the outside world through shielded cables. The cable shield will be connected through a capacitor at the chamber to break any ground loops. Doublet Pair signals sent to the DCS sub-system will be isolated following the general ATLAS grounding policy.

12.12.3 Cooling of TGC electronics

It is necessary to consider the power dissipated by the on-detector TGC trigger electronics, and how they will be cooled. Those electronics situated in USA15 (Local DAQ masters, Sector Logic and RODs), outside the ATLAS cavern, are not considered here since they can be cooled, via fluid rack cooling, in the conventional way.

The location of the on-detector electronics is summarized in Table 12-10. ASD boards are attached to TGC counters inside the unit shielding box. The power density for the ASD boards is quite low since the channel-to-channel distance is of the order of 1cm. Power dissipated by the ASD boards totals 46mW/cm for each TGC counter (summing both wire and strip channels), giving approximately 5W per linear meter. Locally, 16 channels of ASD chips sit on each ASD board, totalling 0.8W per board. This is a low enough heat density to be cooled solely by the ambient air. Total heat dissipation from all ASD boards is less than 20kW.

Table 12-10 Location of the TGC trigger electronics and estimated power consumption.

Item	Location	kW
ASD board	inside TGC unit package	18
LVDS receiver	on TGC wheel	5
Patch-Panel	on TGC wheel	8
Slave board	on TGC wheel	4
High-p_T board	on TGC wheel	2
Read-out	on rim of TGC wheel	0.2
DCS	on rim of TGC wheel	0.2
TTC	on rim of TGC wheel	2

Most of the TGC electronics are located on the inner surface of the triplet wheel or outer surface of the outer doublet wheel. The Patch-Panel and Slave Board are packaged together in a PS-pack, and the power consumption per PS-pack is 20W over a total area of approximately $1.3\,\mathrm{m} \times 0.7\,\mathrm{m}$, or $\sim 1\,\mathrm{m}^2$. This again is not a large heat

density, but the total heat produced is beyond what can be dissipated into the cavern air. Either cooling pads or cooling pipes with a flowing liquid coolant will be used for each PS-pack. The design will comply with any ATLAS decision on acceptable cooling liquids for on-detector electronics. Figure 12-33 shows the arrangement for cooling a PS-Pack. Simulation calculations of heat transmission as well as measurements of actual heat dissipation for a prototype of this design will be performed.

High-p_T boards are placed in special crates located at the rim of the triplet wheel where space is available. They will be cooled using a conventional heat-exchanger for the crate as a unit.

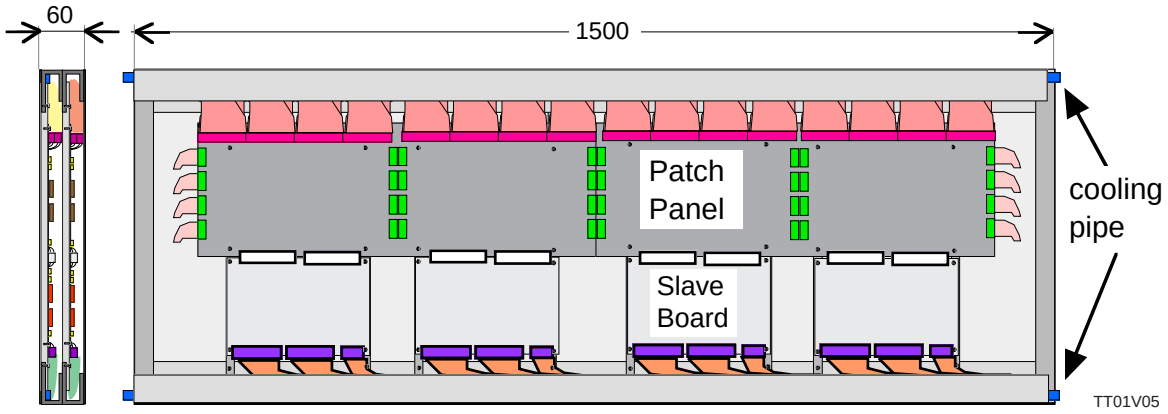

Figure 12-33 PS-Pack showing cooling arrangement and (left) the double circuit board layer structure.

12.13 Prototype experience

In order to test the basic function and performance of the key elements of the level-1 muon end-cap trigger electronics system, prototype test modules were produced using a commercially available FPGA (Field Programmable Gate Array). Although the final logic will be implemented in ASICs, FPGAs allow one to quickly and inexpensively test the concept and set lower limits on the system performance. As a first step, we made a trigger decision prototype module which contained both the lown-p_T trigger circuit (3-out-of-4 coincidence matrix) and the high-p_T trigger circuit. Tests on timing precision for these circuits were performed, as well as a complete check on the function of each element in the module. We also produced additional circuits, such as the 2-out-of-3 coincidence for the triplet layer of TGCs and the bunch-crossing identification, and carried out more comprehensive tests of the trigger electronics system. The experience from the production and tests of those prototype modules are described in the following subsections.

12.13.1 Trigger-logic prototype module

The central part of the level-1 trigger electronics system is the majority coincidence logic, which consists of the low-p_T and high-p_T trigger coincidence matrices. Although these matrices will be located in separate boards in the final system (i.e. the low-p_T matrix in the Slave Board and the high-p_T matrix in the High-p_T Board), we put these matrices onto a single board together with the related circuits for test purposes.

We implemented the majority coincidence logic with Xilinx FPGAs (XC4013). The circuits were made on a 9U VME board, so that the downloading of the FPGA design could be done via VME-bus lines. The Xilinx chip is composed of a matrix of Configurable Logic Blocks (CLBs), each of which acts as a 3-out-of-4 coincidence circuit for the low-p_T trigger unit or a two-fold coincidence for the high-p_T trigger unit. The prototype module was designed to contain the trigger coincidence circuits covering an area corresponding to 32 channels in the TGC pivot plane, as shown in Figure 12-34.

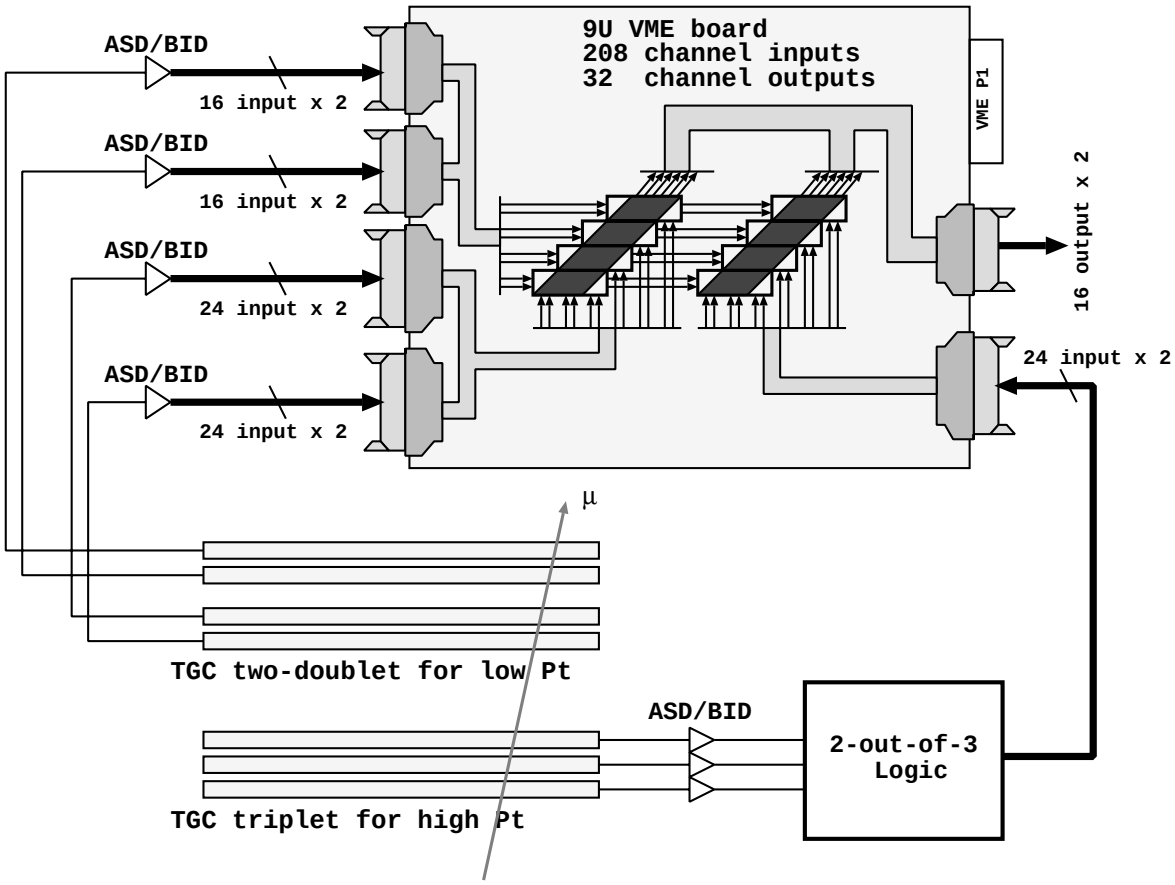

Figure 12-34 Block diagram of the trigger-logic prototype module.

The module accepts inputs of 48 channels from each of the first doublet TGCs and 32 channels from each of the second doublet TGCs (pivot plane) as well as 48 channels from the 2-out-of-3 logic signals of the triplet TGCs.

The outputs of the modules are two sets of 15-bit signals for momentum information: one set for low-p_T and another for high-p_T. Figure 12-35 shows a photograph of the completed trigger-logic prototype board, where the large square chips in the centre of the board are Xilinx chips. Although one chip contains 24×24 CLBs, only a quarter were used due to limitations on the number of input channels and the routing resources between CLBs.

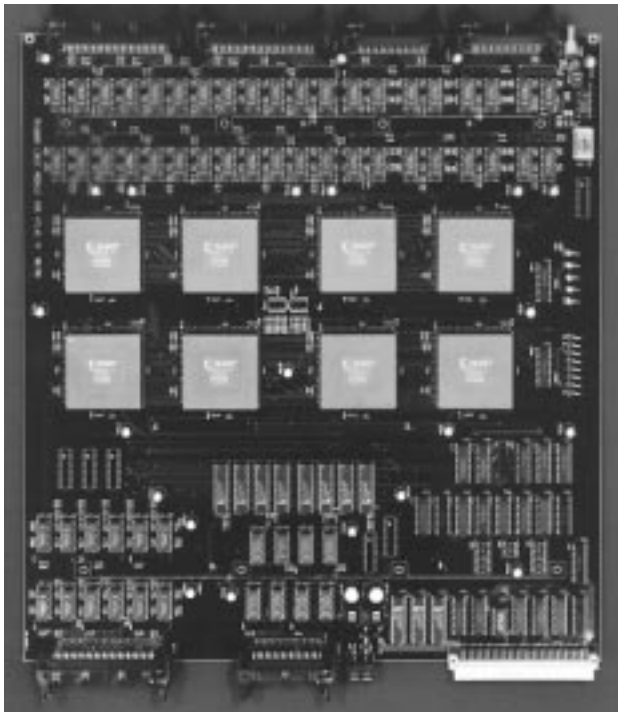

Figure 12-35 The trigger-logic prototype module containing 8 Xilinx FPGA chips(XC4013).

In order to demonstrate how the trigger-logic module works in synchronous mode and to check its operational frequency range, we made a test set-up as shown in Figure 12-36. Only two computer controlled pulse generators, PPGs, were used here in order to make the test as simple as possible. Two independent timing signals (T1 and T2) were synchronized to the clock simulating the bunch-crossing-ID circuits, and the pulse pattern signals from the PPGs were introduced at the corresponding timings. Those signals were then fed into three low-p_T input ports and one high-p_T input port such that the low-p_T and high-p_T trigger conditions could be satisfied. Figure 12-37 shows a test result with a PPG clock frequency of 40 MHz and a trigger gate width of 25 ns, where the hits (T1,T2) giving the trigger coincidence output are plotted in (a) and the hits without the coincidence output in (b). It is clearly seen that the trigger-logic module works 100% correctly.

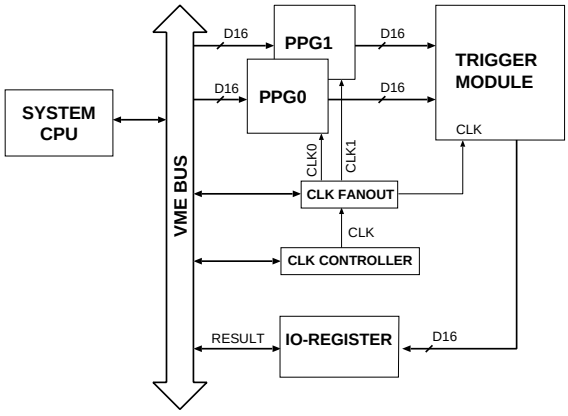

Figure 12-36 Block diagram of the timing test set-up of the trigger-logic module.

The trigger-logic module was designed to work even for the case where the timing distribution of the chamber signals is slightly larger than 25 ns. We made a demonstration test by using the same set-up setting the trigger gate width to 30 ns. The result is shown in Figure 12-38, which agrees exactly with the expected behaviour of the module.

We then examined the maximum frequency at which the trigger-logic module can work, by again using the same set-up and increasing the PPG clock frequency. The prototype module was confirmed to work normally with a clock frequency up to 58 MHz.

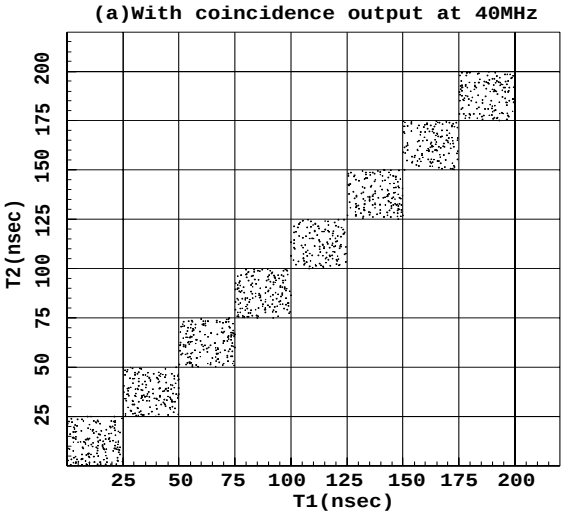

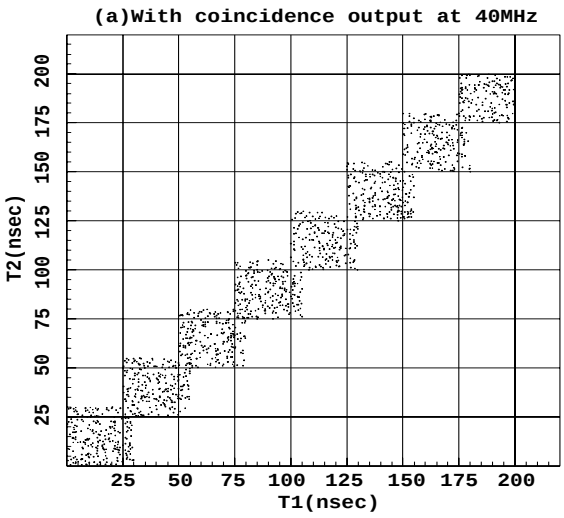

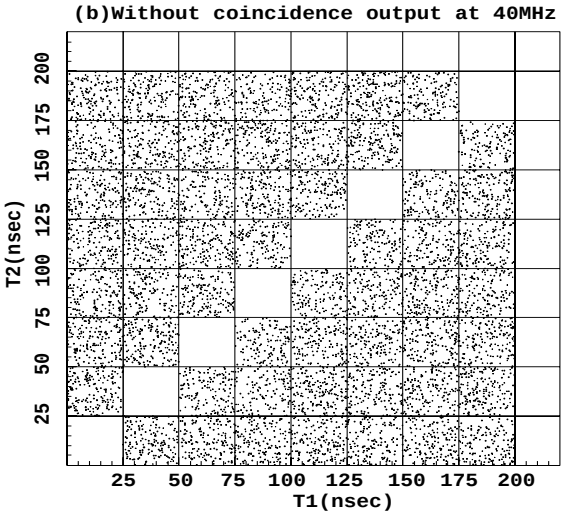

Figure 12-37 Hit map of the timing test of the trigger-logic module, with (a), and without (b), the coincidence output. The clock frequency was 40MHz and the trigger gate width was 25ns.

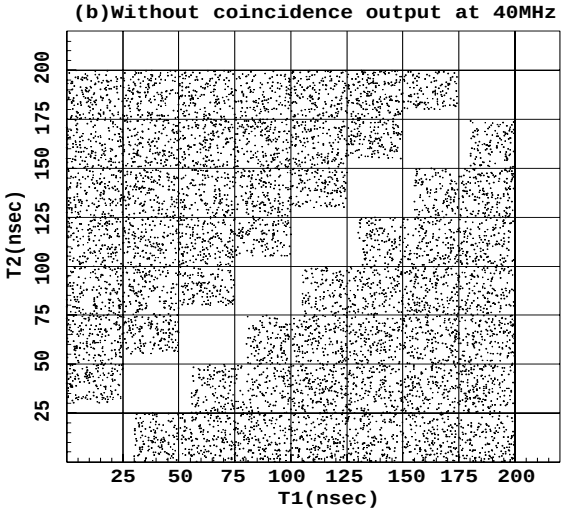

Figure 12-38 Hit map of the timing test of the trigger-logic module, with (a), and without (b), the coincidence output. The clock frequency was 40MHz and the trigger gate width was 30ns.

12.13.2 Bunch-crossing ID Board and overall test of the trigger electronics

The bunch-crossing identification circuit with variable gate width is one of the key elements in the level-1 muon end-cap trigger electronics system, details of which are described in Section 12.1. The circuit was constructed using an FPGA (Xilinx XC4005) on a prototype bunch-crossing ID Board, containing 64 input and 64 output channels. Two such boards were made for the test of the trigger electronics system.

A check was first made of each channel of the bunch-crossing ID circuit by inputting random pulses generated by a PPG module which was operated asynchronously from the 40MHz bunch-crossing clock. The input pulse timing with respect to the bunch-crossing clock was measured by a TDC module. Figure 12-39 shows the result of the check, where the input pulse timings are plotted at the corresponding clock numbers of the output of a channel, for a clock

frequency of 40MHz and trigger gate widths of 25ns and 30ns. For the case when the trigger gate width is greater than 25ns, There is, as expected, a timing range where the bunch-crossing ID circuit gives two consecutive outputs.

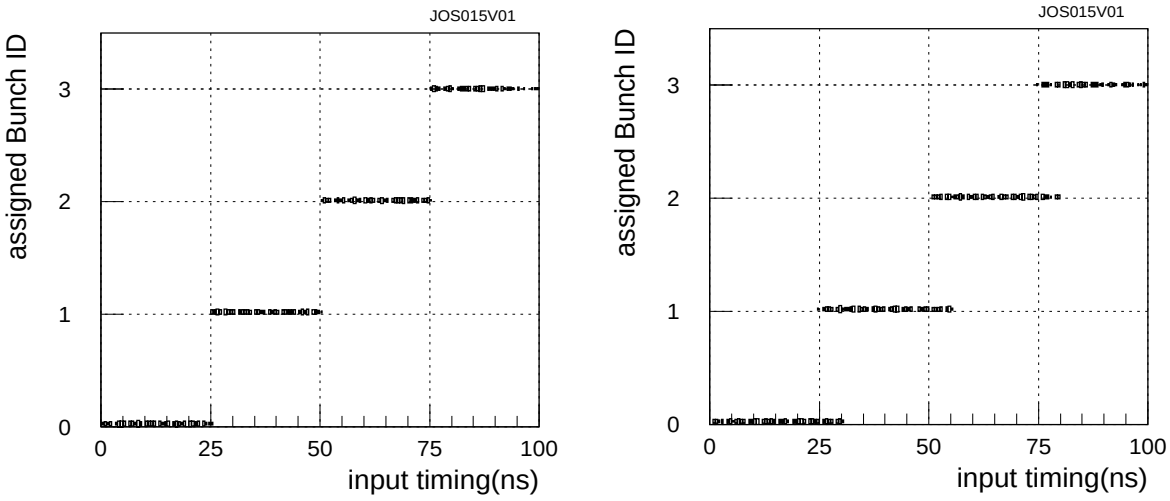

Figure 12-39 Scatter plots of the input pulse timing to the bunch-crossing ID circuit versus the corresponding clock number of the output. The clock frequency was 40MHz and the trigger gate was 25ns (left) and 30ns (right).

We then combined the bunch-crossing ID Boards and the trigger-logic module, and made a test of the accidental trigger rate for the low-p_T trigger. We created 4×16 sets of semi-random pulses using the PPGs operated with 200MHz clocks, and input these pulses to the bunch-crossing ID circuit connecting the outputs to the low-p_T trigger-logic module. The coincidence rates were measured as a function of the frequency of the input random pulses. The result is shown in Figure 12-40 together with the expectation from the calculation showing good agreement.

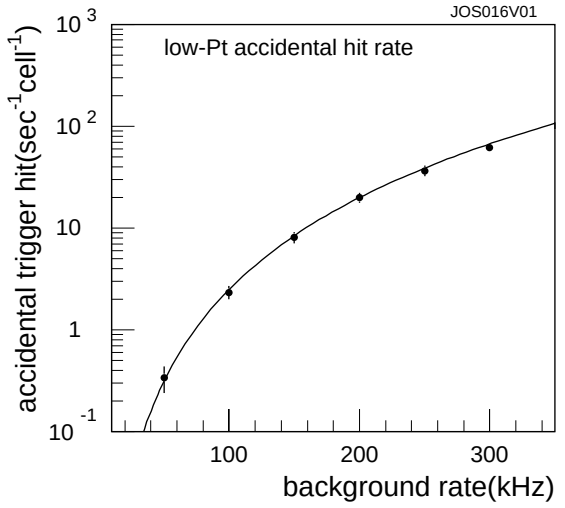

Figure 12-40 Accidental trigger rate for the low-p_T trigger as a function of the random background rate and comparison with the calculation.

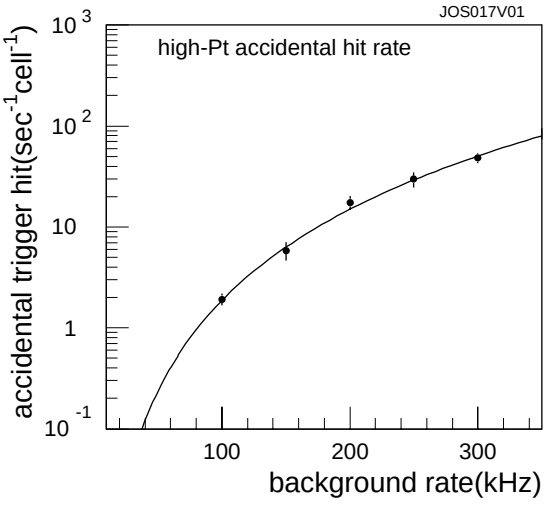

Figure 12-41 Accidental trigger rate for the high-p_T trigger as a function of the random background rate and comparison with the calculation.

We also made the 2-out-of-3 logic circuit for the triplet layer of TGCs by using a programmable logic device (GAL26CV12). After checking the basic function of each element of the circuit, we performed a similar test of the accidental trigger rate for the high-p_T trigger. In this case we put 3×16 sets of semi-random pulses in to the 2-out-of-3 logic circuit, while each of the other 16 sets of semi-random pulses were distributed to the low-p_T circuit giving constant low-p_T coincidence hits. The result is shown in Figure 12-41 and agrees with the calculation.

When the time jitter of the input pulse to the BCID circuit is within 25ns, no problem is observed with bunch-crossing identification. However, for some cases (e.g. TGC strip signals from long strips) the time jitter may slightly exceed 25ns. We thus performed a test to measure the bunch-crossing misidentification rate for the realistic time jitter distribution of TGC signals.

Figure 12-42 shows the simulated time jitter distribution of TGC signals (almost the worst case), which has been confirmed to agree with the measurement of the real signal distribution. The pulses having such a timing distribution were created by the combination of the PPGs and a programmable delay module, and used for the inputs to the same test set-up as that mentioned above. The PPGs were operated synchronously with the bunch-crossing clock and the finer pulse timing ($\Delta t = 1$ns) was created by the programmable delay module. The resulting bunch-crossing misidentification rate and the trigger efficiency are plotted in Figure 12-43 as a function of the gate width. The result is compared with the expectation from the simulation, and shows good agreement. This leads us to the conclusion that the bunch-crossing misidentification rate can be made negligibly small while keeping the end-cap muon trigger efficiency high.

Finally, we carried out some of the above-mentioned tests using clock frequencies higher than 40MHz. The whole trigger electronics system was confirmed to work normally with a clock frequency of up to 70MHz, except for the low-p_T trigger circuit which worked only up to 58MHz.

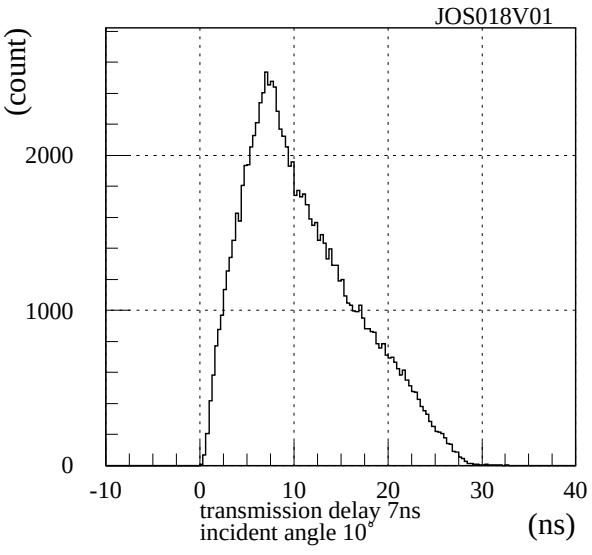

Figure 12-42 Simulated time jitter of TGC signals.

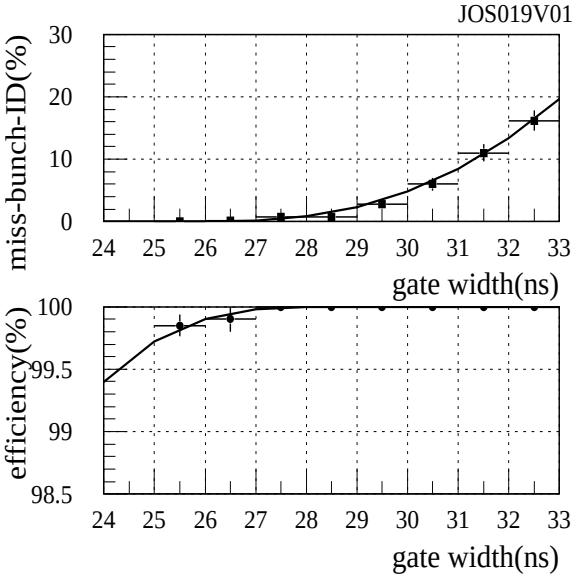

Figure 12-43 Bunch-crossing misidentification rate (top) and trigger efficiency (bottom) as a function of the input signal frequency for the time jitter distribution as shown in Figure 12-42.

12.13.3 Migrating from FPGA to ASIC

As the FPGA chip has much more functionality than is necessary for our use, is not radiation tolerant, and is rather expensive when used in large numbers, it is not our choice for the final system. Radiation tolerant ASIC chips will be produced, which will contain only the necessary functions of the FPGA chip with similar or better performance. For the level-1 muon end-cap trigger electronics system, we think it natural to make two types of ASIC chips: one for the Patch-Panel and the other for the Slave Board and High-p_T Board.

The ASIC chip on the Patch-Panel would possess the following functions (see Section 12.4.2 for the details):

- LVDS receiver,
- Bunch-crossing ID circuit,
- OR circuit for TGC overlaps,
- adjustable delays,
- DAC for threshold voltage,
- test-pulse for ASD.

As it needs both digital and analog circuits to incorporate all these functions, we plan to fabricate the chip using a full-custom 0.6μm CMOS process. The number of gates in the chip would not be very large (~10,000).

There are slightly different coincidence logic circuits in each of the Doublet Slave Board / Triplet Slave Board / High-p_T Board for both the wire and the strip signals. Hence, there are six types of coincidence circuits in total, to be incorporated in a single ASIC chip. We will produce a test chip using the full custom CMOS process, but the final design will use gate array technology. The number of gates in the chip is estimated to be about 200,000.

The reasons such a choice are: only 40MHz logic circuits are required and we have some experience with the radiation tolerance of a similar technology [12-12]. In any case, we will perform radiation tolerance tests for both types of ASIC chip produced by the process finally chosen.

12.14 Construction and assembly

The construction schedule for the level-1 muon end-cap trigger electronics system is summarized in Figure 12-44 and Table 12-11. Milestones for the production are summarized in Table 12-12, and the planning of the Production Readiness Reviews (PRR) in Table 12-13 below.

12.14.1 Construction schedule

The ASD boards are mounted on the TGCs. Since the mass-production phase of the TGC chamber construction will begin in 1999, production of the ASD boards will start in the middle of 1999 in order to coincide with this time-scale. To this end, we have to confirm the radiation tolerance of the ASD chip and boards, and pass the PRR by the end of 1998.

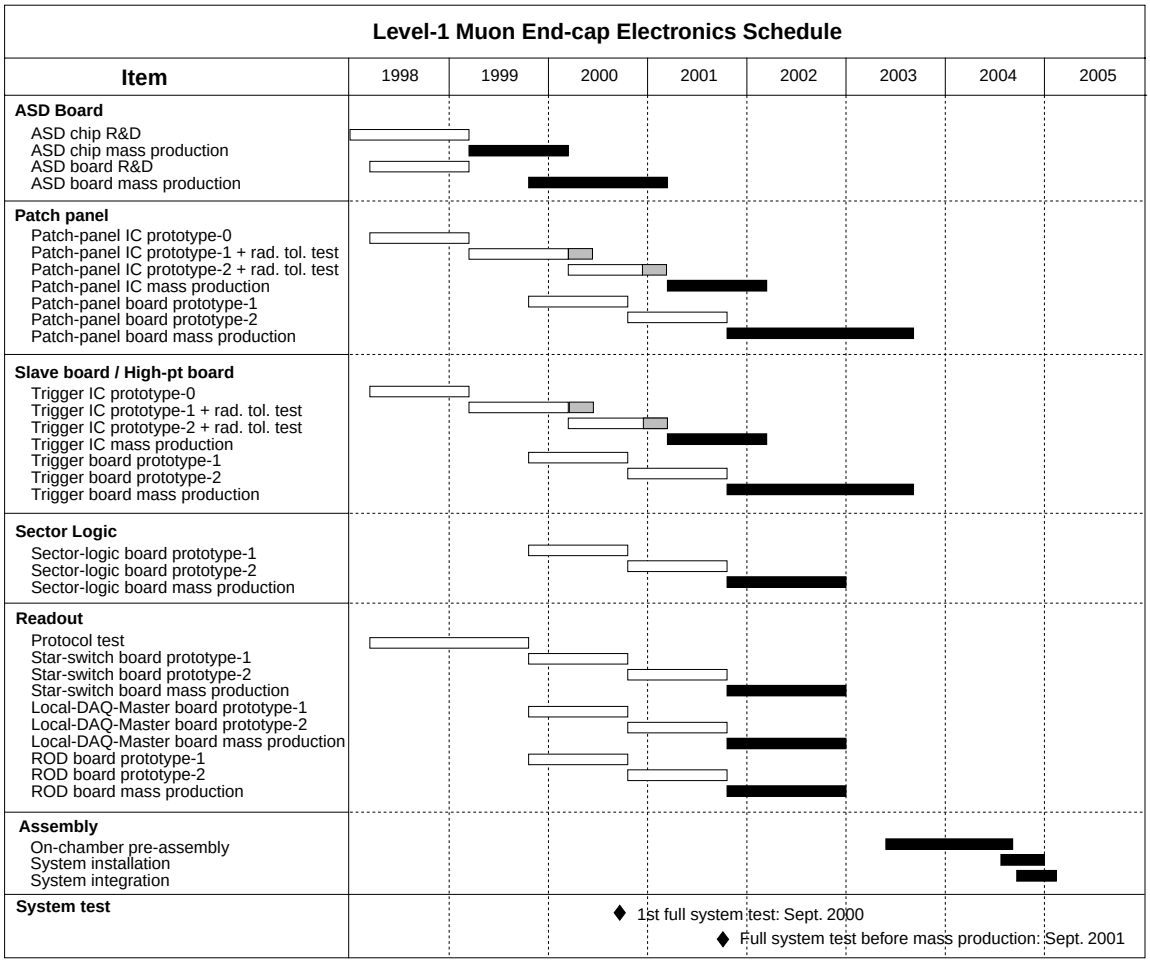

Figure 12-44 Construction schedule and assembly of the level-1 muon end-cap trigger system.

The ASICs to be developed for the level-1 end-cap trigger are the Patch-Panel ASIC and the trigger ASIC. The first prototype for testing the design functionality will be done at the VLSI Design and Education Center, VDEC, in the University of Tokyo. Since VDEC is limited to small quantities, development will continue in collaboration with a major IC manufacturer who can produce large quantities of ICs. The development will continue until the first quarter of 2001. We will confirm the radiation tolerance of these ASICs and then begin mass-production of the ASIC chips in 2001. In parallel with the ICs, we will develop the Patch-Panel board and the trigger boards. Mass production of these boards will start in October 2001 and continue until September 2003.

Production of prototype Sector Logic boards will start at the end of 1999, when the detailed designs are finalized. Since the board is located in USA15, radiation tolerance is not required. We will use commercially available programmable IC chips which maximize flexibility and minimize the time needed for prototyping.

For the read-out system, the protocol test will continue until the third quarter of 1999. The production of prototype boards of the Star Switch, Local DAQ Master, and Read-out Driver will start at the end of 1999. With the exception of the Star Switch board, radiation tolerance is not required. Since the functionality of the IC in the Star Switch is rather simple, we will make use

Table 12-11 Construction schedule for the muon end-cap trigger and DAQ electronics. The dates shown here are the production start and completion dates for each component.

Item	Production start	Production completion
ASD chip	April 1999	March 2000
ASD board	November 1999	March 2001
Patch-Panel chip	March 2001	March 2002
Patch-Panel	October 2001	September 2003
Trigger chip	April 2001	March 2002
Doublet Slave board	October 2001	September 2003
Triplet Slave board	October 2001	September 2003
High-p_T board	October 2001	September 2003
Sector Logic board	October 2001	December 2002
Star Switch board	October 2001	December 2002
DAQ Master board and ROD	October 2001	December 2002

of commercially available one-time-programmable devices (anti-fuse), which are radiation tolerant. Mass production of these boards will finish by the end of 2002.

With the exception of the ASD board, the trigger and read-out systems are closely coupled. We plan to have a single PRR of the whole system before the end of 2000 at the latest. The first full system test will be in September 2000, and the system test before mass production is scheduled for September 2001. These milestones are summarized in Table 12-12.

Table 12-12 Milestones in the construction of the muon trigger electronics system.

Milestone	Date
Read-out protocol fixed	October 1999
First full system test	September 2000
Full system test before mass production	September 2001

12.14.2 Assembly procedure

The ASD boards are mounted on the edge of the TGC doublets and triplets at the last stage of the chamber production line. The TGCs are to be tested for about two weeks with cosmic rays. Unless any failure is found, the boards are never dismounted.

The patch-panel and slave boards are mounted on the TGC wheels. The assembly unit of TGCs is one 24th of each wheel. The assembly of the electronics boards will be done above ground together with the TGCs on a local support frame. Cabling and piping will also be done during the assembly of this sector unit. The assembly is scheduled to start in mid-2003 and to last for about one year.

The High-p_T board and the Star Switch boards are located in the crates around the rim of the triplet TGC wheel. The low-voltage power supplies are installed in racks on the periphery of the MDT wheel. The TGC and MDT wheels are scheduled to be installed in the ATLAS cavern at the end of 2004. Mounting these boards will be done after the installation of the wheel structure.

The sector-logic board, the Local-DAQ-Master board, and the ROD board are installed in the crates in the USA15. This process can be done independently of the detector installation.

12.14.3 Quality Assurance Programme

Given the task of building a detector system that must satisfy rigorous performance and reliability criteria for a project that will run for over 10 years, a quality assurance programme for the end-cap muon trigger is a prerequisite. This is emphasized by the rather limited access to some components of the trigger system once assembly is complete.

The end-cap muon trigger subsystem follows a Quality Assurance (QA) review procedure. Total quality management proceeds from the onset of the design stage and continues through manufacturing, handling, delivery, testing, commissioning and operation / maintenance. The system is divided into many, essentially independent, subsystems. Proper functional definition of each subsystem and specifications for interconnections between them are essential to realise a quality system. These definitions and specifications are described in the Level-1 Muon Trigger User Requirements Document (URD) [12-13] which also specifies procedures of control, testing and maintenance as well as outlining in detail the physics performance goals of the system.

We plan to perform the following tests for the trigger system and its components to guarantee their performance to the standards required:

- Evaluate the failure rate of each subsystem as currently designed in order to make the design more robust and reliable.
- Evaluate the probable frequency of repair to meet the physics goal on each subsystem.
- Evaluate the effect of failure on the trigger system for various types of failure, and design the system so that the more probable failures will cause least problems to the overall system.
- Select components for the resulting level of required reliability.
- Burn-in all essential boards / units to reduce initial failures to a minimum.
- Ensure that we possess sufficient spares for repair purposes for the entire 10-year running cycle. This will include a study of alternative components to replace those that may not be produced throughout the experiment’s lifetime due to a termination of present production processes.
- Establish repair routines for the more routine failures.
- Actual tests of major units / boards under conditions that accurately simulate the expected environment in terms of radiation and magnetic field.

Before proceeding to the mass-production stage, Production Readiness Reviews (PRR) are scheduled for the major components of the system, Table 12-13. We additionally plan a PRR of the entire system before the end of 2000. These reviews will cover all aspects of the system design including integration issues, logistics, quality and safety and will provide an opportunity for Quality Assurance aspects of the design to be thoroughly examined. This together with the Quality Assurance Plan will allow a rigorous review of the status of all aspects of the preparation of the system, including the evaluation of prototype modules.

Table 12-13 Dates of Production Readiness Reviews (PRR) for components of the level-1 end-cap muon trigger system.

Item	Date
ASD chip / ASD board	October 1998
Patch-Panel chip / Patch-Panel	October 2000
Trigger chip / trigger boards (Slave, High-p_T)	October 2000
Read-out boards (Star Switch, Local DAQ Master, ROD)	October 2000

12.15 References

12-1 Ph. Farthouat, *Current Understanding of the Muon Level-1 System*, version 2, ATLAS Note DAQ–NO–089, May 1998.

12-2 *The TGC Excel Parameter Book*, http://www.cern.ch/~lellouch/design/tgc.pdf

12-3 Low Voltage Differential Signalling, LVDS, IEEE 1596.3 standard. See, for example, http://sunshine.cern.ch/LVDS/

12-4 The ATLAS Collaboration, *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

12-5 O. Jinnouchi, O. Sasaki, *Study of muon LVL1 trigger scheme*, ATLAS note DAQ–NO–082, Feb 1998.

12-6 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998. http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

12-7 The so-called 'TP32' version of the background study reported in a private communication from A. Ferrari, November 1996.

12-8 CERN RD12 group, *Timing, Trigger and Control (TTC) Systems for LHC Detectors*, http://www.cern.ch/TTC/intro.html

12-9 *ATLAS policy on radiation tolerant electronics*, ATLAS working document, http://www.cern.ch/Atlas/GROUPS/FRONTEND/WWW/radtol2.ps

12-10 RD49 Collaboration, *Study of the Radiation Tolerance of ICs for LHC*, CERN / LHCC / 97–63, LEB Status Report / RD49, December 1997.

12-11 B. Williams and Ph. Farthouat, *ATLAS Policy on Grounding and Power Distribution*, ATLAS working document, http://www.cern.ch/Atlas/GROUPS/FRONTEND/WWW/gnd_note.ps

12-12 Y. Arai, private communication. The 0.5 micron CMOS process used for the TMC test chip for MDT was confirmed to be radiation tolerant.

12-13 *LVL1 Muon Trigger User Requirements Document (Draft Version 1.4)*, ATLAS working document, March 1998,
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/muons/L1MT980309.ps

13 The muon trigger interface to the CTP

13.1 Introduction

As already discussed in Chapter 9, the muon trigger system is composed of three subsystems (see Figure 13-1) – detector-specific logic (DSL) associated with the RPC and TGC detector systems (see [13-1][13-2] for a detailed description of the TGC and RPC) and the muon trigger / CTP interface (MUCTPI). The DSL subsystems associated with the RPC and TGC

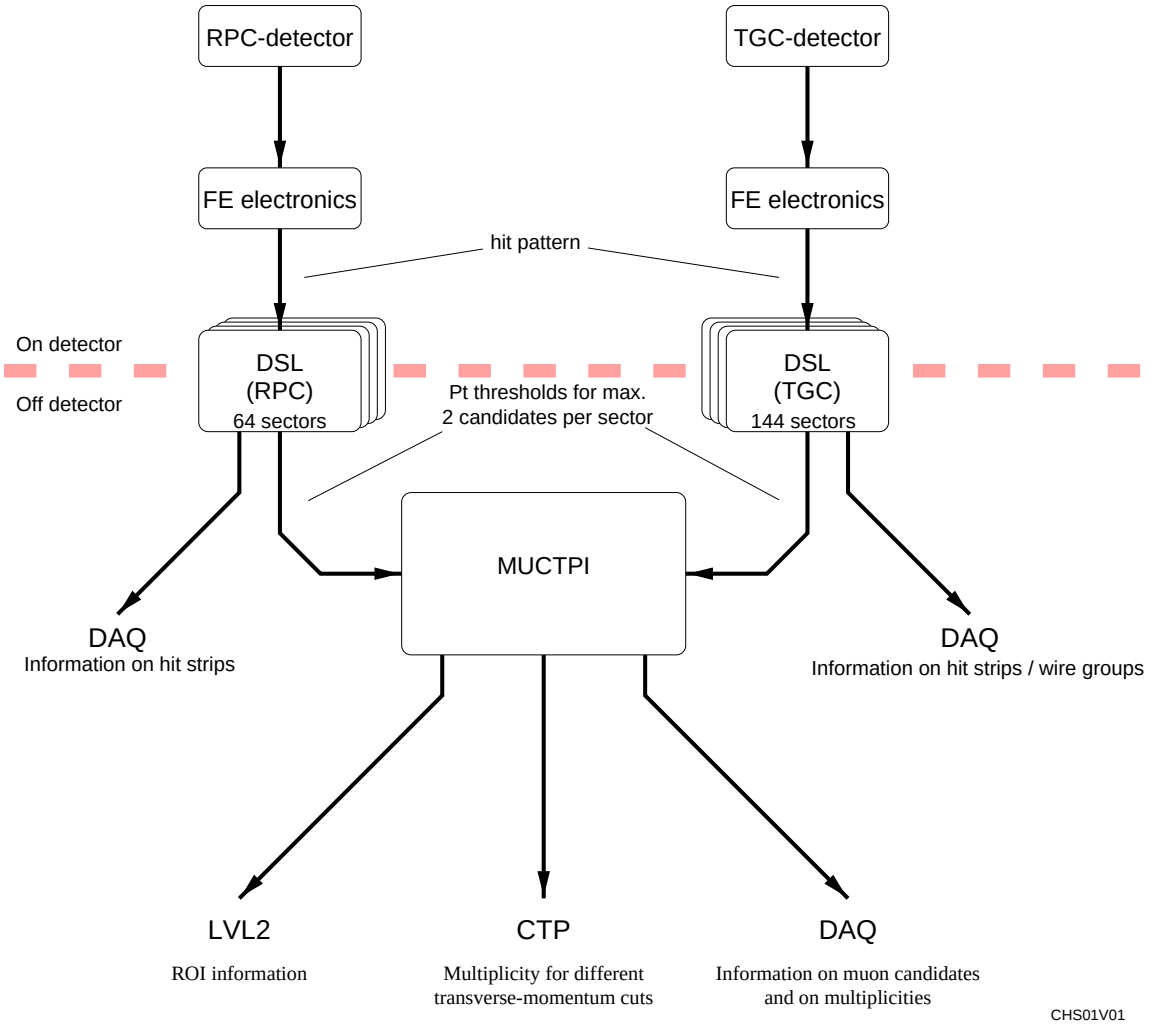

Figure 13-1 Data flow in the muon-trigger system.

systems were described in detail in Chapters 11 and 12 respectively. References [13-3] and [13-4] summarize the trigger schemes for barrel and end-cap regions. This chapter describes the MUCTPI subsystem.

The results from the DSL that form the input to the MUCTPI provide information on up to two muon-track candidates per sector. The information includes the position and p_T range of the track candidates. The MUCTPI combines the information from all of the sectors and calculates total multiplicity values for each of the six p_T thresholds. These multiplicity values are sent to the CTP. Note that in forming the multiplicities care has to be taken to avoid double-counting single muons in regions where trigger chambers overlap. Otherwise doubly-counted single muons could dominate the low-p_T dimuon trigger, giving an unacceptably high rate.

Additional functions of the MUCTPI are to provide data to the LVL2 trigger and to the DAQ system for events selected at LVL1. The LVL2 trigger is sent a formatted copy of the information on candidate muon tracks. This information is used to define regions of interest (RoIs) that drive the LVL2 muon-trigger processing. The DAQ system receives a more complete set of information, including in addition the computed multiplicity values. The information sent to the LVL2 trigger is ordered, according to decreasing in p_T.

13.2 Summary of requirements

The MUCTPI collects data from the DSL; after data processing and formatting, results are sent to the central trigger processor (CTP), the LVL2 trigger, and the DAQ. A detailed list of the technical requirements for the MUCTPI can be found in Ref.s [13-5] and [13-6]. In the following, a short summary of some of the main requirements is given, as needed for the understanding of the systems implementation.

- Some muons cross overlapping trigger chambers in such a way that hits are produced in both chambers. If the chambers belong to different sectors, the DSL may send two muon candidates caused by the same particle to the MUCTPI. The MUCTPI must make sure that such candidates are counted only once. To do this it uses information indicating if a muon candidate has been found in a region where trigger chambers overlap.
- The DSL can send up to two muon candidates per sector to the MUCTPI. In order to allow for some flexibility in forming the LVL1 trigger decision, two methods of forming the multiplicities have to be provided: for each sector either all muon candidates are taken into account, or only the candidate with the highest p_T contributes to the multiplicity calculation. The two options have to be independently programmable for each of the six p_T-thresholds.
- The maximum overall multiplicity that needs to be handled by the system is seven candidates. Larger multiplicity values are rounded down to seven.
- For all data which are read out via the DAQ system, pipelines have to be implemented which hold data for a time corresponding to the latency of the LVL1 system. In addition it must be possible to read out data from a programmable window of up to ±2 bunch-crossings (BCs) width around the LVL1 trigger. This is needed for setting up the timing of the system and also facilitates monitoring of activity in the trigger chambers shortly before and after the event which caused the trigger.
- Since the data received by the MUCTPI come from different parts of the detector the MUCTPI must align the incoming data in time, compensating for different times of flight and signal propagation delays, so that only data corresponding to the same bunch crossing are used to form multiplicities and for defining RoIs.

- The latency of the MUCTPI, measured from the arrival of the last piece of data from the DSL until the last piece of information has arrived at the CTP, must not exceed eight BCs.
- The interface to the LVL2 trigger must operate without data loss up to a LVL1 trigger rate of 100 kHz. The latency for providing the trigger information to LVL2 should not exceed 100 μs.
- The DAQ interface must cope with the maximal expected occupancy at full luminosity without losing data.
- Sufficient online monitoring information must be provided in order to check the correct functioning of the MUCTPI during beam running and to facilitate fast localization and diagnosis of possible problems.

13.3 Interfaces

The MUCTPI has interfaces to the DSL, the CTP, the LVL2 trigger and to a readout buffer (ROB) for the DAQ (see Figure 13-1). Other interfaces not shown in the figure are to the TTC system (provides clock and trigger signals), to the run-control system and the monitoring system, and to the DCS (for monitoring of power supplies and crate cooling).

13.3.1 Interface to DSL

As discussed above, the DSL is divided into sectors readout each of which provides information to the MUCTPI. In total there are 208 sectors, 2×32 for the barrel, 2×48 for the end-caps and 2×24 for the forward regions. For each sector, one 32-bit word is received at the BC rate (40 MHz). Figure 13-2 illustrates the detailed format of these words [13-7]. Each word contains information on up to two muon candidates. The p_T thresholds of the candidates are encoded with numbers from one to six. The location of each candidate in the sector is given by the RoI address. For the barrel and end-cap, flags indicate if the candidate was found in a zone of overlapping barrel sectors or of overlapping barrel and end-cap sectors. These flags are used to avoid double-counting muons in these zones when forming the total multiplicities. An overflow flag signals the presence of more than two candidates in a sector. Since the barrel DSL can handle at most one candidate in a group of four RoI regions (so-called pad), and the end-cap and forward DSL can handle only one candidate per subsector, additional overflow flags indicate if more than one candidate occurs in these regions. In all cases of an overflow, the highest p_T muon candidates are kept. Candidates within a sector are ordered according to p_T. Candidate 1 refers to the highest-p_T candidate and candidate 2 to the second-highest-p_T one.

The data words are received in the MUCTPI by 16 so-called octant boards which will be described below. Each board receives data from 13 sectors covering a full octant for one end of the detector.

13.3.2 Interface to CTP

The connection from the MUCTPI to the CTP has to be as fast as possible to minimize the latency. Since the MUCTPI and the CTP will be located very close to each other in the same rack or in adjacent ones, the cable length is only of order a metre. The six 3-bit multiplicity values will be sent as differential electrical signals, for example using twist-and-flat cables.

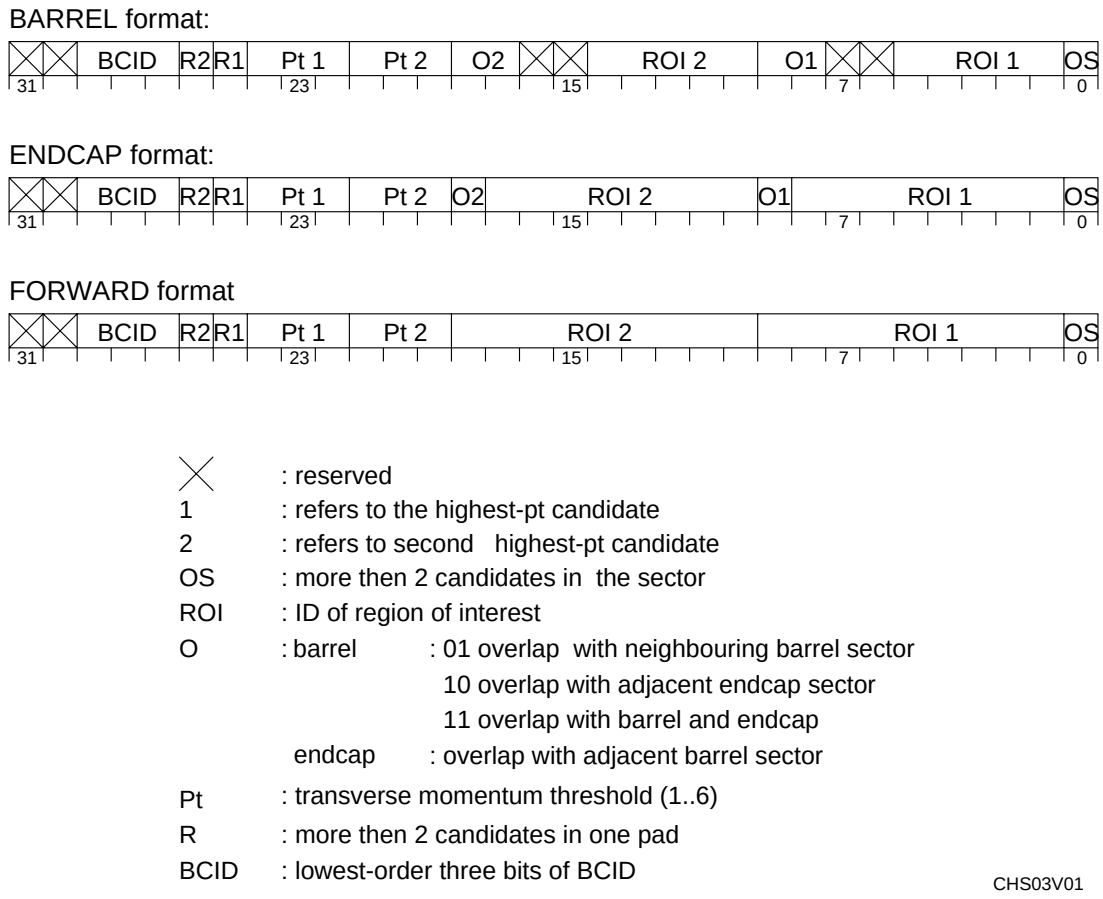

Figure 13-2 Format of the input data coming from the DSL.

13.3.3 Interface to LVL2

Information on muon candidates must be sent to the LVL2 interface so that the LVL2 trigger processor can build up regions of interest which it then investigates closer using information from other subdetectors. For every event triggered by LVL1 the following information must be sent:

- The event number and the BCID number
- The number of muon candidates to be sent
- A flag indicating if not all found muon candidates could be sent to LVL2. [The total number of candidates sent to the LVL2 trigger will be limited to a fixed number.]
- For each muon candidate which is sent to LVL2:
    a. The p_T range of the candidate.
    b. The sector identifier and subsector address of the candidate.
    c. A flag indicating if it was the highest-p_T or the second-highest-p_T candidate in its sector.
    d. Flags indicating if the candidate was found in a region of overlapping sectors.

e. Flags indicating if more than two candidates were found in the sector or more than one candidate in the subsector.

The physical implementation of the interface has not been finally decided. A possible solution which exists today is the S-LINK interface [13-8].

13.3.4 Interface to DAQ

For each event triggered by LVL1, the information that is sent to LVL2 is also sent to the DAQ system via a ROB. In contrast to LVL2, which receives the information only for the selected bunch crossing, the DAQ is sent information in a time frame of up to ±2 BC around the triggered one. The DAQ receives as additional information, also in a time frame of up to ±2 BC around the triggered one, the multiplicity values calculated by the MUCTPI.

As discussed in Section 13.4.4 below, the physical link to the DAQ will be the ATLAS standard readout link.

13.3.5 Interface to the TTC

The MUCTPI receives various standard signals from the TTC system: the LHC clock (BC) from which the local system clock is derived, the bunch-counter reset signal (BCR) used to check the synchronization of the readout, and the LVL1 trigger decision (L1A) used to initiate the readout for selected events.

13.4 Functional partitioning and specification

The MUCTPI is divided into a number of building blocks which are housed in one 9U VME crate shown in Figure 13-3. In addition to the modules drawn, there will be a VME master in the crate (e.g. commercial CPU module).

The different functionalities of the MUCTPI are implemented in three types of VME modules which are connected to each other via an active backplane. The functionality of the MUCTPI is shared by these system components as follows (see Figure 13-4):

- 16 so-called octant boards (MIOCTs) receive data corresponding to an octant in the azimuthal direction and half the detector in the η direction. They form muon-candidate multiplicities for this region, correctly taking into account the overlap zones between barrel and end-cap sectors. There is no overlap between muon trigger chambers associated to different octant boards.
- The Muon Interface to the CTP (MICTP) contains the driver unit to the CTP.
- The Muon Interface to Read Out Driver (MIROD) drives, after some data formatting, data to the ROI builder of the second-level trigger and serves as the interface to the Read Out Buffer (ROB).
- All modules are connected via the Muon Interface Backplane (MIBAK). It contains two components: The active part forms the total candidate multiplicities by adding the multiplicities of the MIOCT boards. The passive part contains a bus system to transfer data from the MIOCTs and the MICTP to the MIROD.

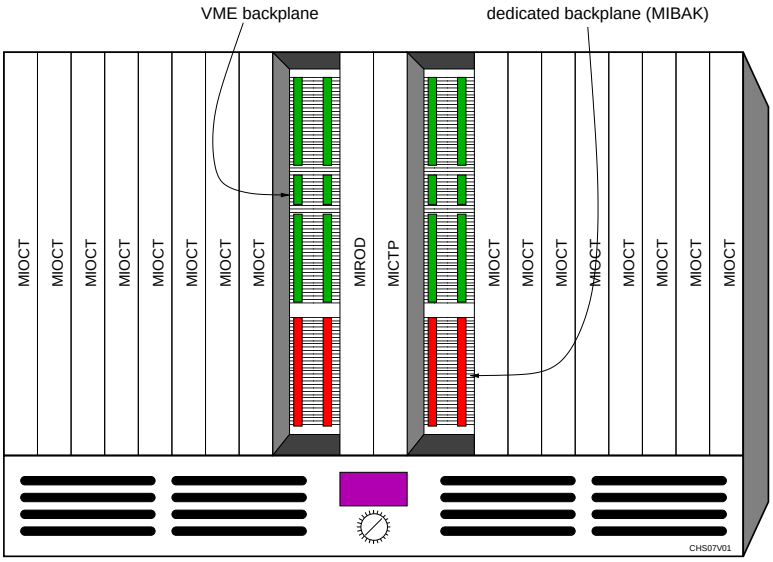

Figure 13-3 The layout of the MUCTPI crate.

13.4.1 The Octant boards

Sixteen Muon Interface Octant boards (MIOCTs) receive data from the DSL. The octant boards are divided into two groups which cover the positive and the negative rapidity regions of the detector. Each octant board receives data from four barrel, six end-cap, and three forward sectors. The region covered corresponds to one octant in azimuthal for half the detector.

Figure 13-5 shows the block diagram of one octant board. After signals are received from the DSL they are synchronized to the phase of the internal BC clock at the input of the MIOCT boards. Sampling is done on the positive or negative edge of the internal clock depending on the position of the input signal transition with respect to the internal clock. A TDC [13-9], incorporated in the octant board, is used to determine the appropriate sampling transition and, during running, to monitor in the timing of incoming signal transitions. Note that the variation in this timing is expected to be negligible because the timing is determined by properties of the interconnection cables and the fixed relative phase of the BC clock in the sector logic and the octant board.

After synchronization, the different input signals have to be aligned in time. This is necessary because of different processing latency between the RPC and TGC-based sector logic, and also because of different cable lengths and times of flight for different parts of the system. The alignment stage in the MIOCT boards is performed using configurable-length shift registers (range 0–10 BCs)[1]. The right configuration for the length of the registers can be determined using test events sent by the DSL to the MUCTPI. The data word of each sector contains a bit field with the three least significant bits of the BC (see Fig. 13-2). The length of the shift registers

1. Note the distinction between programmable and configurable parameters. The former can be changed simply and rapidly through memory accesses. The latter are used for parameters that do not change once the system has been set up — they can be changed through relatively slow and complicated procedures, for example by loading new FPGA configuration files.

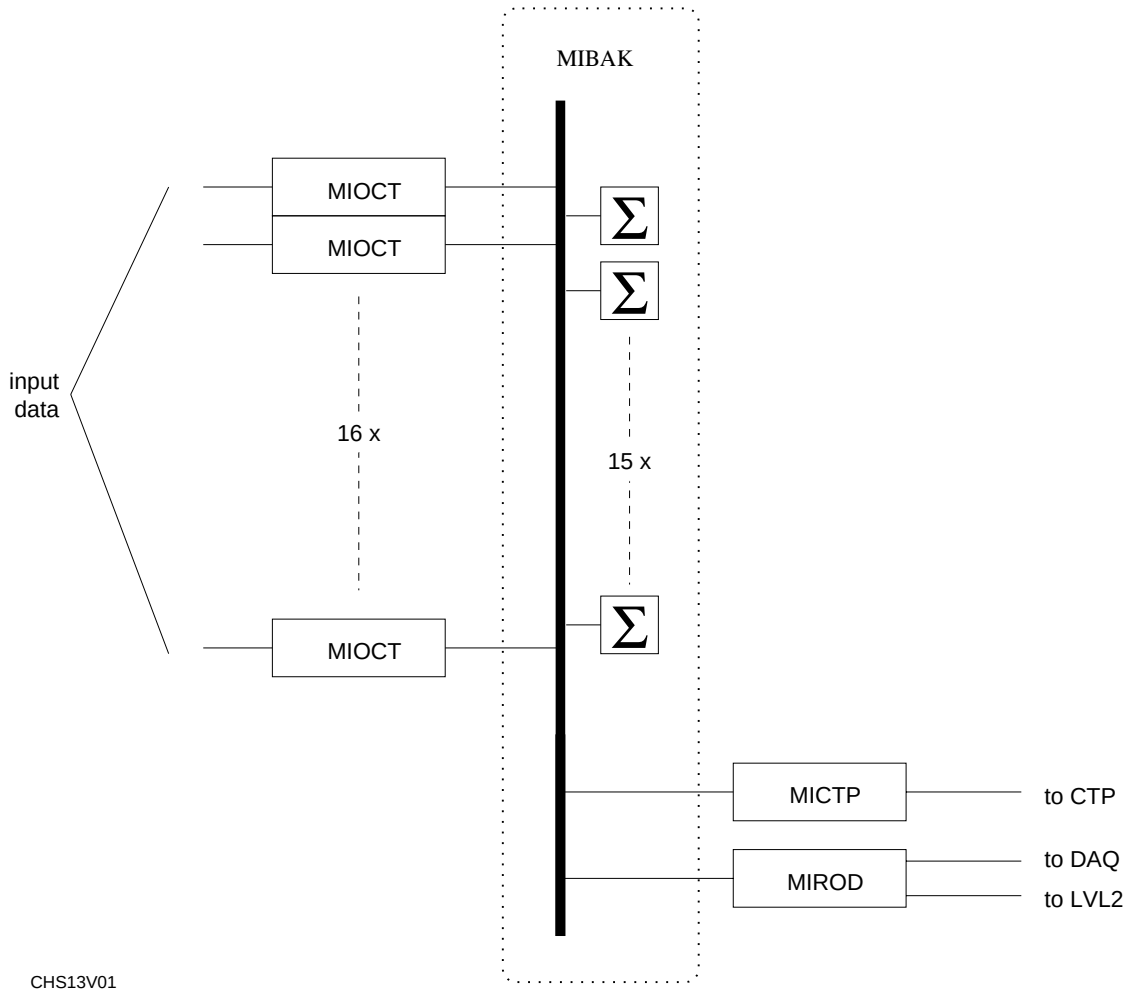

Figure 13-4 Functional partitioning of the MUCTPI.

has to be configured so that the BC numbers of all data words coming from the same L1A are equal.

All the data received from the DSL are stored in pipeline memories for the duration of the latency of the LVL1 trigger. In case of a LVL1 Accept (L1A), data of all input channels have to be read out of the pipelines by the derandomizers and written into the local readout buffer. A programmable window around the triggered BC can be defined. Up to ±2 bunch crossings around L1A can be read out. At a L1A rate of 75 kHz, the total data rate in an octant board is 176 Mbit / sec (13 input channels; 32-bit input word extended by 4-bit channel ID; maximum 5 BC long time frame; 75 kHz rate).

Before data of the 16 octant boards can be sent via the internal backplane to the MIROD, a zero-suppression stage is needed in order to reduce the data rate to a tolerable level. Subsequently, data belonging to one L1A are formatted into a data package (with header and trailer) and buffered until they are read out via the backplane. The data format of these packages is shown in Figure 13-6. The first word of every package forms a header containing a four-bit head-identification pattern (bits 31 to 35), the identification number of the MIOCT module (MNBR), the event number (EVID), and the BCID. What follows are one word for each sector containing

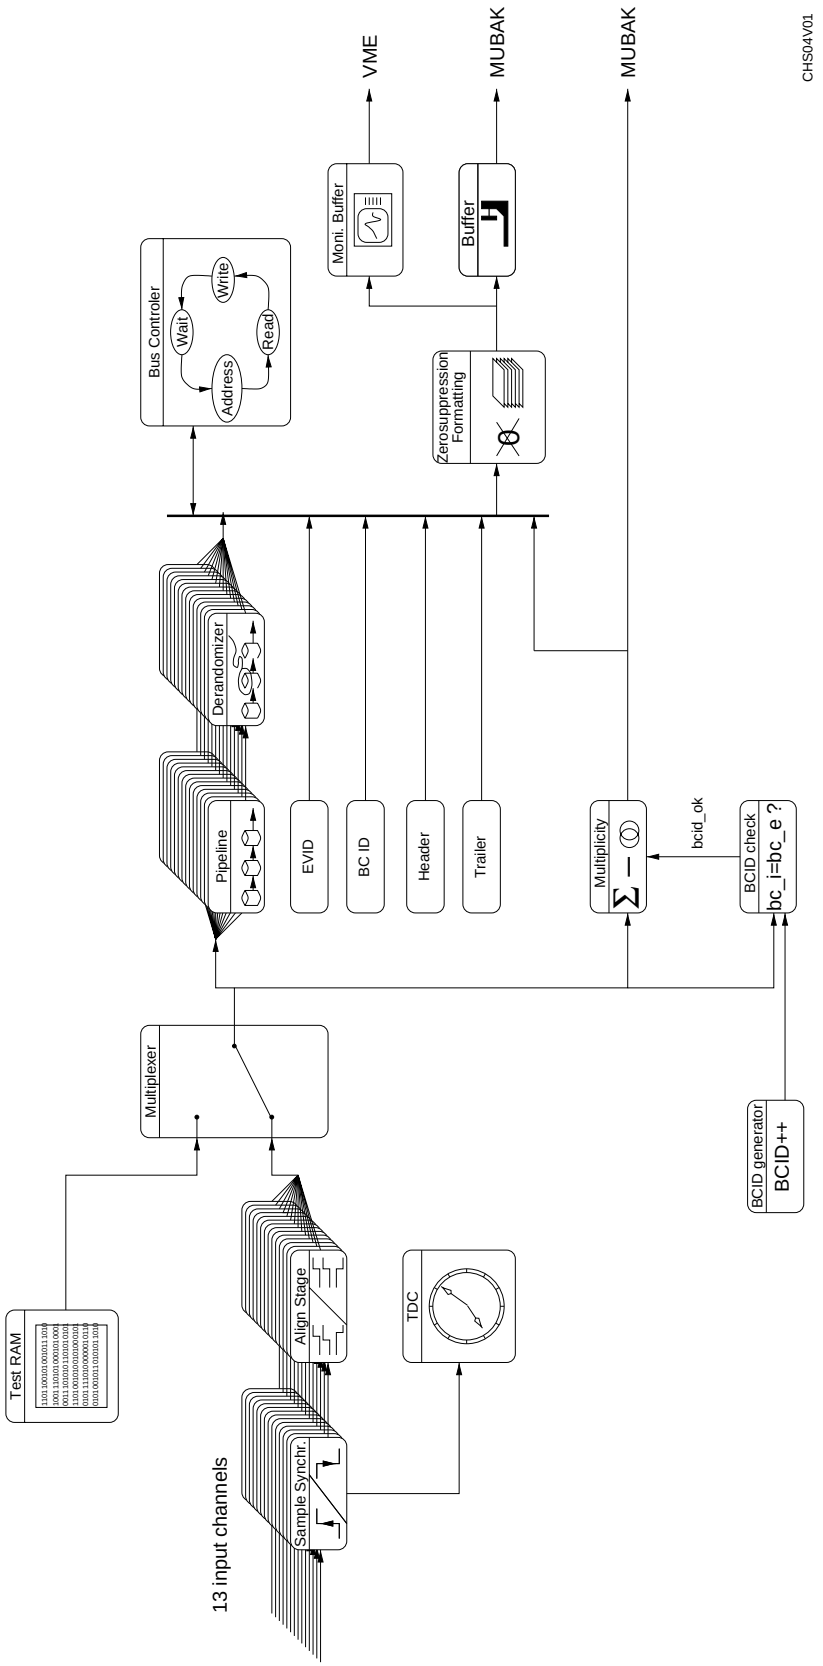

Figure 13-5 Block diagram of the octant board.

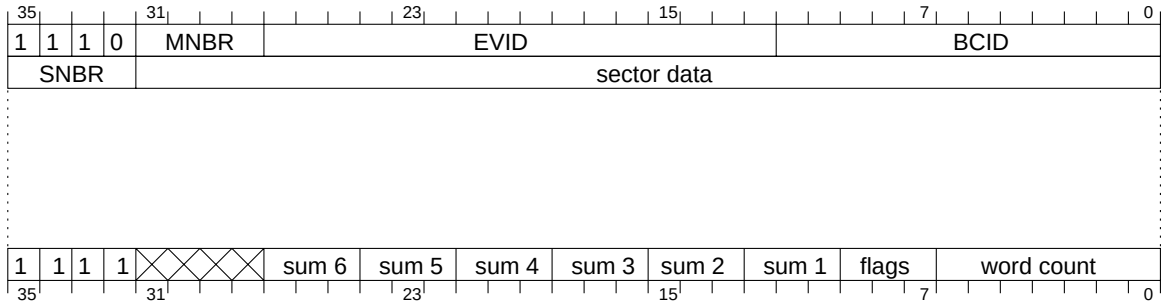

Figure 13-6 Format of the data packages sent from the MIOCT modules to the MIROD via the backplane.

at least one muon candidate. The format of these words corresponds to the format of the input data to the MIOCT module (see Section 13.3.1). Additionally, the sector number (SNBR) is added in the four most significant bits. A trailer, identified by the pattern 1,1,1,1 in the four most significant bits, closes the data package. It contains the six multiplicity sums, three bits reserved for flags, and a word count indicating the total length of the data package in words.

The multiplicity logic counts the number of muon candidates for each BC clock in the six different transverse-momentum classes taking account of possible overlaps between adjacent barrel sectors and barrel and end-cap sectors. Figure 13-7 shows the position of these regions in

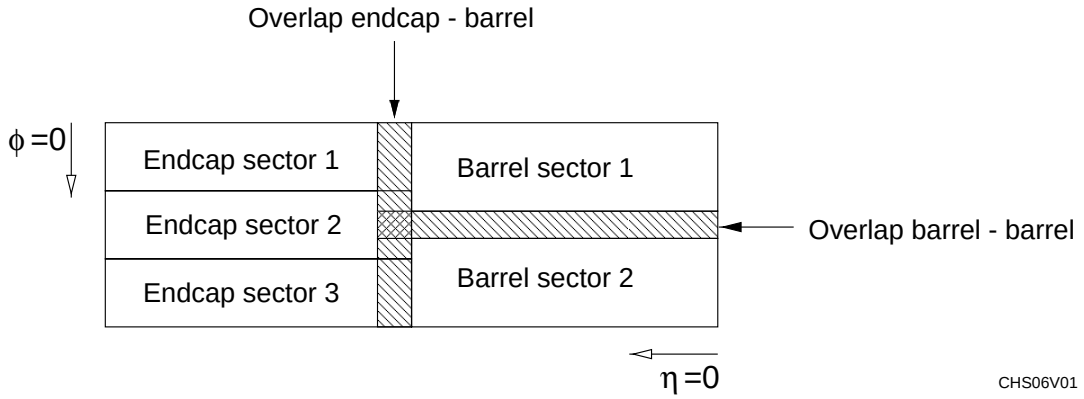

Figure 13-7 Position of the overlap regions of the muon trigger chambers.

the muon trigger chambers. Each octant board covers two such structures. There is no overlap

between sectors associated to different octant boards. For each group of sectors corresponding to Figure 13-7, a fast truth table is used to detect the presence of muon candidates in different sectors which could be caused by the same particle. Figure 13-8 illustrates how the truth tables, together with a set of lookup tables, avoid the double counting of muon candidates in overlap regions. If the truth table detects two muon candidates in the same overlap region belonging to different sectors, it asserts a veto for one of the candidates. The subsequent adder stage will ignore candidates which have been vetoed. The outputs of all first-stage adders are summed to form six 3-bit sums which are clipped at a maximal count of seven.

In order to test the system without relying on data from the DSL, the input stage can be switched to a test RAM (256×32 bits) that can be loaded with an arbitrary test pattern via VME. Monitoring information is acquired from the system via its VME interface. In addition, VME is used to configure the FPGAs of the whole system. This allows for high flexibility even in the case when minor modifications are desired.

13.4.2 The interface board to the CTP

The interface board to the CTP, MICTP, collects the multiplicity sums for the six p_T thresholds over the custom backplane, MUBAK, described in Section 13.4.3 below, and transmits them to the CTP. The MICTP is responsible also for distributing time-critical control signals to the rest of the MUCTPI system.

13.4.2.1 Multiplicity formation for the CTP

Figure 13-9 shows a block diagram of the MICTP. The upper part shows the building blocks receiving the total muon multiplicities from the MIBAK backplane; these are then driven to the CTP. In order to be able to test the CTP without relying on data from the MIOCT boards, the MICTP contains a programmable test-pattern RAM which can be multiplexed to the output drivers. Additionally a synchronization pattern can be driven to the CTP. This is used by the CTP to synchronize the input signals to the internal system clock.

The multiplicity sums that are sent to the CTP are also stored in a pipeline memory. In case of a LVL1 trigger, data from the pipeline are stored in a derandomizer. These data are subsequently read out by the MIROD module (see Section 13.4.4) over a custom bus on the backplane.

For monitoring purposes, eight counters for each p_T threshold count the occurrences of 0,1,2,...7 muon candidates per bunch crossing. At the end of each LHC turn the contents of these counters are copied into a set of registers which can then be read via VME. At the same time the counters are reset.

13.4.2.2 Distribution of LVL1 control signals

The MICTP distributes the system clock to the rest of the MUCTPI. Under normal operation this clock is the LHC BC clock. For test purposes, an internally-generated 40 MHz clock can be used instead. Five other time-critical control signals are sent to all modules of the system via the backplane: The signals ECR (event counter reset) and BCR (bunch counter reset) reset all local event and bunch counters. The L1A signal is distributed to all modules to indicate a valid LVL1 trigger; this signal initiates the transfer of data from the pipelines to the derandomizers. The signal labelled ‘synch moni’ is used to synchronize the writing of monitoring information into

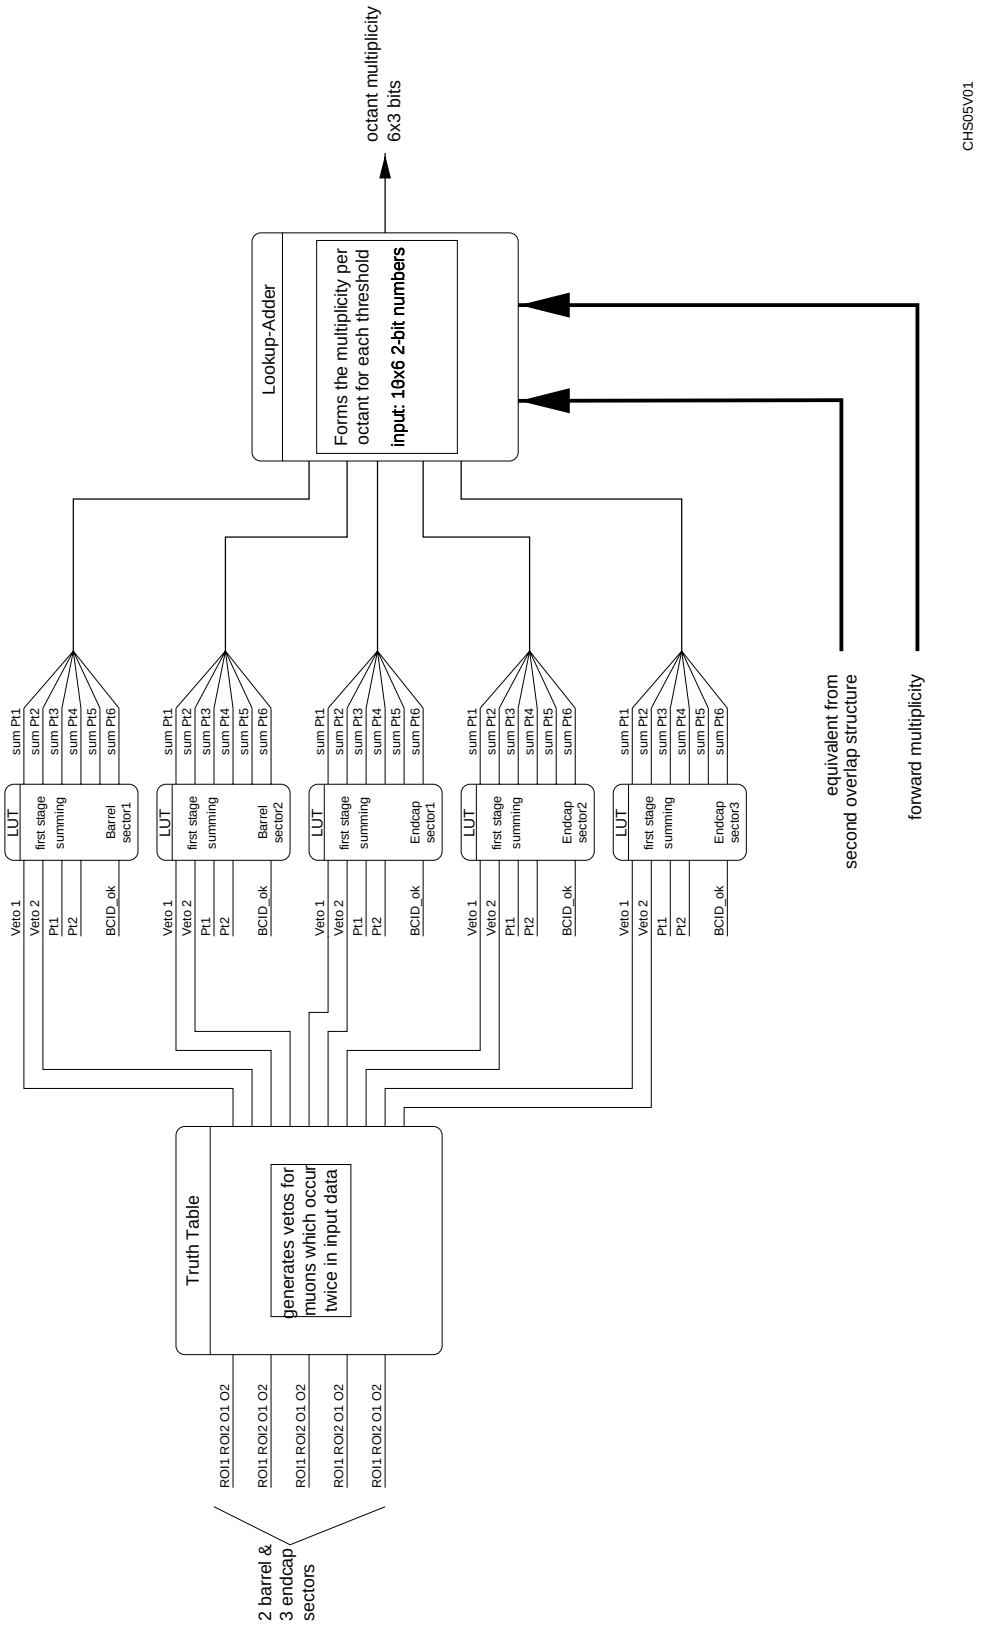

Figure 13-8 Logic performing the sum of muon candidates within an octant board.

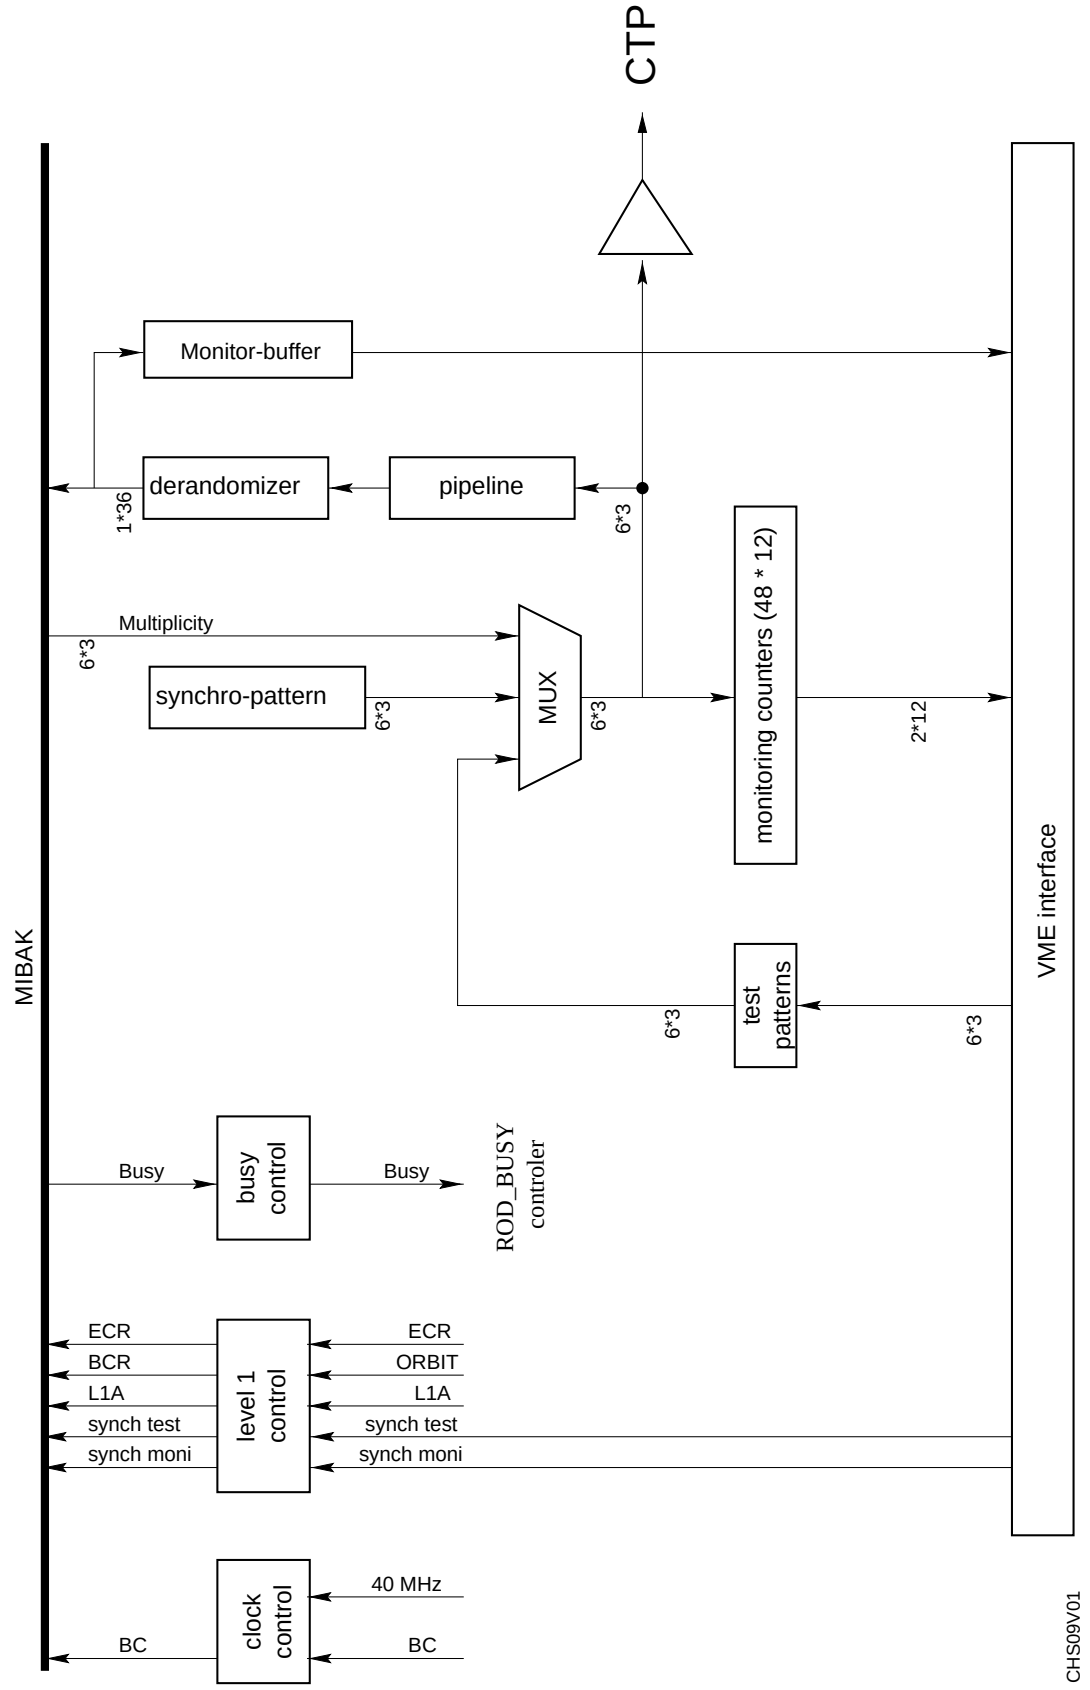

Figure 13-9 Block diagram of the MICTP.

dedicated buffers among all modules of the system, making it possible to investigate data belonging to the same L1A in all modules of the MUCTPI by reading out via VME the dedicated buffers. The signal ‘synch test’ is used to synchronously initiate the test cycle in which the octant boards are fed by the local test-pattern RAM.

In addition the MICTP monitors the BUSY line on the MIBAK backplane. Modules of the MUCTPI can activate this line in case they cannot accept more L1A triggers because their internal buffers are nearly full. In case of an activated BUSY in the system, the MICTP drives a BUSY signal to the BUSY_ROD module, which handles all BUSY signals generated in ATLAS (see Chapter 20).

13.4.3 The backplane

The backplane of the system consists of two parts. The upper part (J1, J2, J0) is a standard VME backplane. The lower part (J3), called MIBAK, is a custom backplane designed to form multiplicity sums over the MIOCTs, to move data from the MIOCTs and the MICTP to the MIROD, and to distribute time-critical signals in the system.

The octant boards are arranged in the MUCTPI crate as shown in Figure 13-3. Each octant board drives six 3-bit sums corresponding to the six different transverse-momentum thresholds onto the backplane. Additional summers are installed on the backplane itself, and connected as shown in Figure 13-10. It is possible to form the sums over the full detector in four layers of adders. The final results are received by the MICTP. The advantage of such an arrangement with respect to a solution where all active components are placed into the octant boards, is that signals have to traverse a minimum number of connectors. This leads to small signal distortions and minimizes propagation delays. Additionally the number of pins of the backplane connectors are minimized. In principle also an analog variant, in order to form the sums on the backplane, could be considered. After evaluating the times which would be necessary for digital-to-analog conversions before summing and analog-to-digital conversion of the results, no reduction of the system latency with respect to the digital solution can be expected. The digital version is preferred since it is less prone to noise on the backplane and therefore easier to implement.

In order to transfer data from the MIOCTs and the MICTP to the MIROD, a simple one-directional bus with high data throughput based on a token architecture has been developed. It is implemented in ECL technology in order to reach high readout rates at high data security. The bus is 36 bits wide in order to transfer the data packages described in Section 13.4.1. A simple handshake facilitates fast readout of all MIOCT buffers, as described in more detail in Section 13.4.4.

The MIBAK is used to distribute time-critical control signals to the MUCTPI system. The backplane layout must make sure that these signals arrive at all relevant modules of the MUCTPI crate with a fixed phase with respect to the BC signal. This ensures that in all modules these signals are synchronized to the same bunch-crossing.

13.4.4 The interface board to the LVL2 trigger and to the DAQ

Figure 13-11 shows a block diagram of the interface board to the LVL2 trigger and to the DAQ (MIROD). The MIROD module has to collect data from the octant boards and the MICTP before

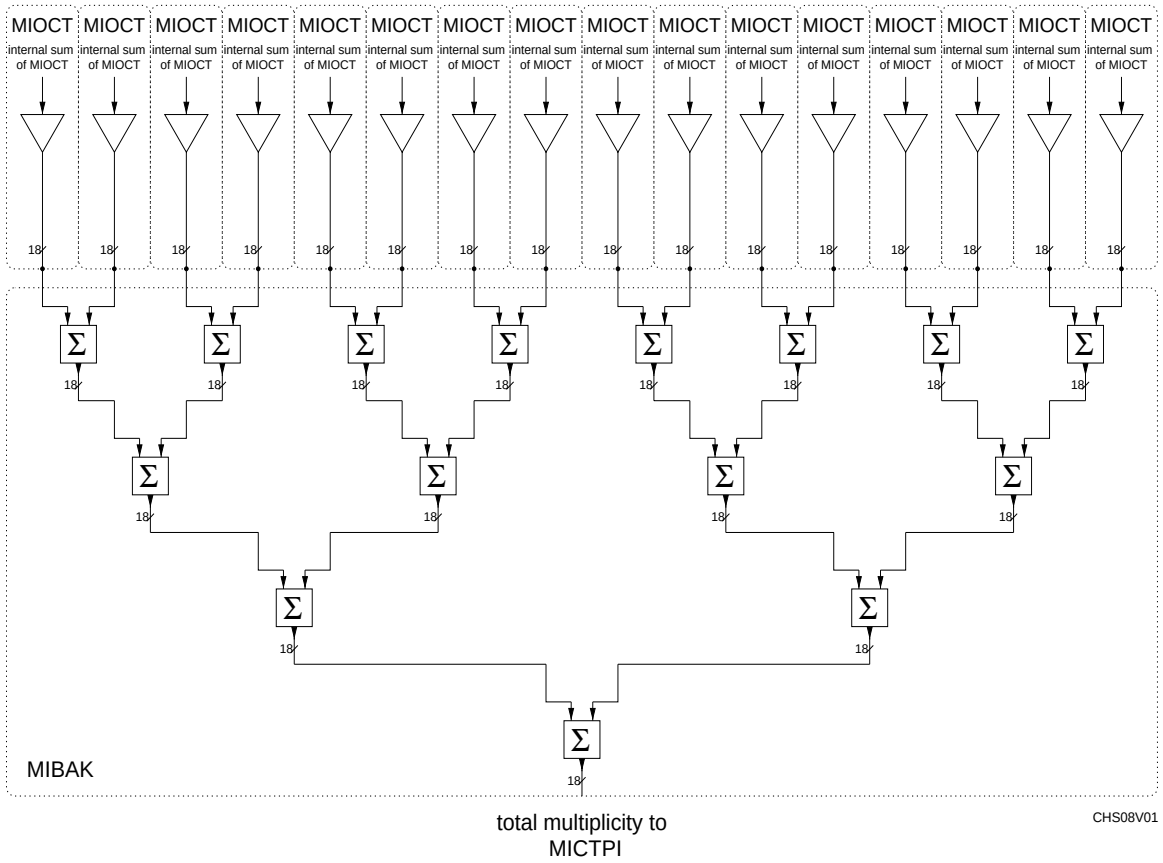

Figure 13-10 Schematic view of the part of the MUCTPI backplane which serves to form digital multiplicity sums among the octant boards.

sending them to the region of interest builder of the LVL2 trigger and the ROB of the DAQ system. The data are transferred over the token-ring bus implemented on the MIBAK. The bus is unidirectional and the MIROD module is the master for all transfers. After a L1A signal all slave modules have to indicate with a dedicated line on the bus that their data are ready for readout. The lines of all slave modules are combined in an OR. The MIROD module initiates a readout cycle as soon as all slave modules have their data available. It puts the token on the token ring which initiates the readout of the first module in the token-ring chain. The token subsequently propagates through all slave modules and returns to the MIROD module after the last slave has been read out. The synchronous readout works at a frequency of 40 MHz.

In order to provide the LVL2 trigger with the information necessary to form the RoIs, data received from the MIOCT modules on all muon candidates are classified according to their transverse momentum and then written into an intermediate buffer. The data used by the LVL2 trigger for the RoI generation belong to a single BCID. Within the window of data which is sent to the DAQ, the time slice for the data sent to LVL2 can be programmed in the MIROD.

In the MIROD, a p_T threshold can be programmed so that only muons exceeding this threshold will be sent to LVL2. In addition, it is possible to optionally take into account only the highest-p_T muon candidate of each sector. The total number of muon candidates transferred is limited to a fixed number (10 in the present design). A flag indicates if, due to this limit, not all candidates

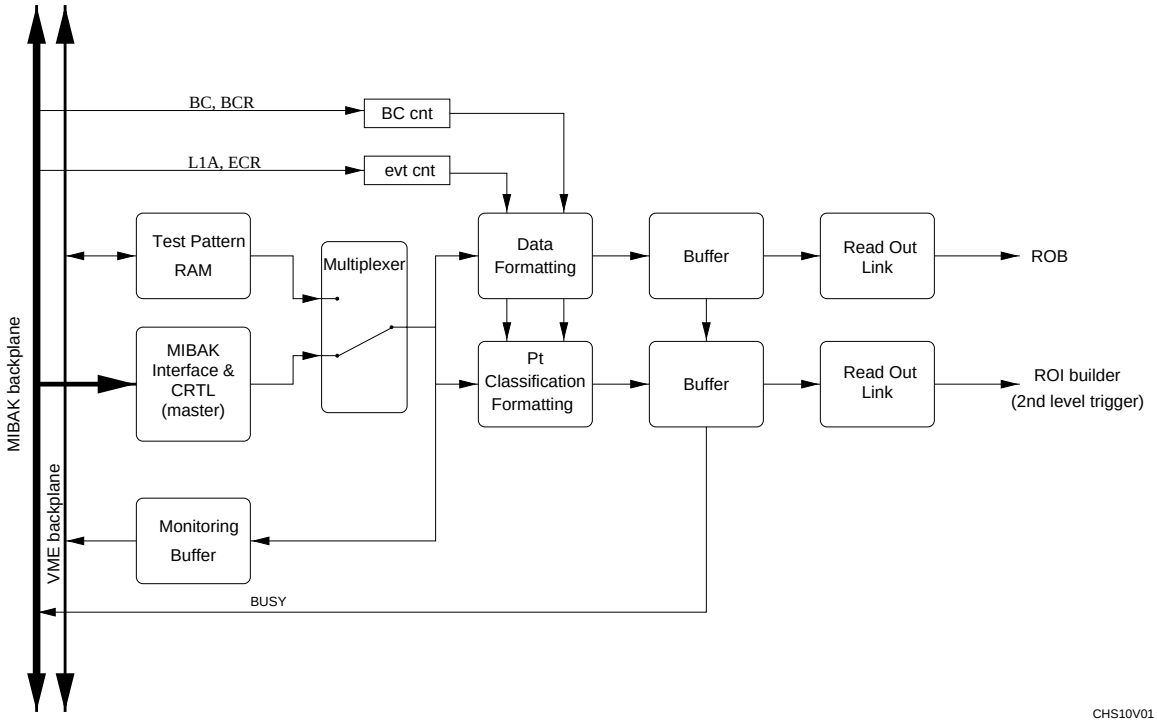

Figure 13-11 Block diagram of the MIROD.

are transferred. Candidates are ordered according to p_T before they are sent, the highest-p_T ones being sent first.

The MIROD module performs checks regarding the consistency of the data packages received from the backplane. For example, a check is made that all data sent for a given L1A come from the same BC. In case of inconsistency, the data sent to the ROB are flagged.

The MIROD module constantly monitors the status of its internal buffers. It asserts a busy signal in order to inhibit the generation of more triggers, which could lead to data loss, in case a buffer fills up close to its capacity.

In order to facilitate the monitoring of data arriving in the module, a buffer is filled with part of the incoming data. With help of a dedicated signal on the backplane, the filling of this buffer can be synchronized to the filling of equivalent buffers in the MIOCT boards. To be able to test the LVL2 trigger system and the DAQ system from the MIROD onwards, it is possible to feed the MIROD with test-data from a programmable RAM.

A prototype implementation for the links to LVL2 and to the ROB is the S-LINK [13-8]. It allows the transfer of data at a rate of 1.28 GBit / sec (32 bits at 40 MHz).

13.5 Overall latency calculation

The MUCTPI is designed such that it will not contribute more than eight BCs to the overall latency of the LVL1 trigger system. Figure 13-12 shows how the MUCTPI system components

contribute to that latency. The latency estimates have been done on the basis of electronics components which are on the market today.

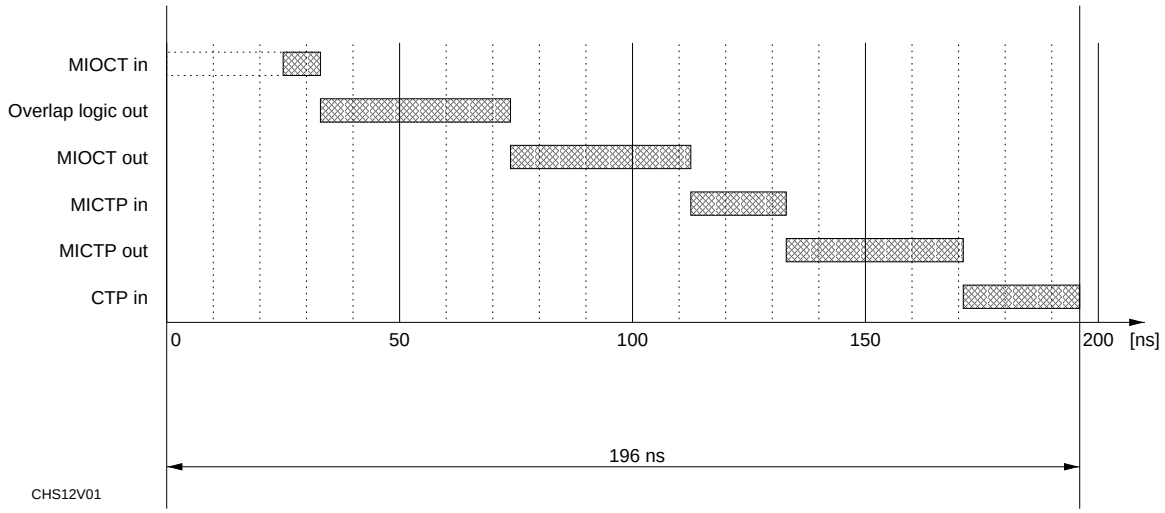

Figure 13-12 Breakdown of the latency contributions for the MUCTPI system components.

13.6 Summary of monitoring and test facilities

Important aspects of monitoring and testing have been pursued in the design of the MUCTPI system:

- By providing all modules with test-pattern memory it is possible to test every module independently of other hardware in the trigger system.
- By providing all modules with monitoring buffers, which are filled with data from the internal data flow and can be read via VME, it is possible to check the data coherence by comparing data read from the buffer with results computed from the input data. Additionally this information will be read out continuously during running conditions by the monitoring software. Histograms like, for example, hit maps and rate histories will be provided in order to monitor the correct functioning of the system.
- By providing a method to write synchronously in all modules to the monitoring buffers data of the same event, it is possible to trace data from the input to the output of the MUCTPI and compare the activity of the MUCTPI with expectations.
- Statistics on the formed multiplicities are provided by a set of counters in the MICTP for each LHC turn.

In order to make use of these online monitoring facilities, the MUCTPI contains a real-time VME controller. Efficient software is needed to read out the monitoring information while evaluating it and presenting it to the operators, so that possible problems in the system can be tracked down quickly.

More online monitoring is provided by the LVL2 and event-filter systems where the information of the MUCTPI can be combined with data from other detector components like for

example the muon precision chambers. These systems will be able to calculate online trigger efficiencies or noise levels in the muon trigger system.

13.7 Protoyping plans

A demonstrator programme is in process for the MUCTPI. A small protoype system will be built containing a MUCTPI crate with at least two MIOCT boards, a MIBAK backplane, a MIROD, and a MICTP board. At least two MIOCT boards are necessary in order to test the multiplicity formation via the backplane. The demonstrator prototype modules deviate from the final versions of the modules in the following respects:

- The MIOCT board has only one physical input for one sector. This is not a severe restriction for the demonstrator since arbitrary test patterns can be loaded in the test RAMs for the other sectors, and the hardware can be tested without restrictions using these test patterns.
- The signal receiver of the MIOCT board is implemented on a small mezzanine card. This allows different technologies to be tried out for the interface to the DSL.
- The overlap logic in the MIOCT module is implemented such that it ignores both muons of an end-cap sector if both of them overlap with the same barrel sector. This means that in these very rare cases the calculated multiplicity will be one too low. A modification of the truth tables handling the overlap will overcome this problem in the final version of the module.
- In the demonstrator MIOCT board, there is no memory foreseen in order to capture incoming data and read them out for monitoring purposes ([13-5]).

13.8 Construction and schedule

The MUCTPI components will be implemented using standard programmable logic components like FPGAs and CPLDs. The whole system can be built using technology available today.

Figure 13-13 shows the time schedule for the development of the MUCTPI. It is foreseen to have a minimal system consisting of at least two octant boards, a backplane, a MICTP, and a MIROD available end of 1999 so that the whole system can be tested in a test beam scheduled then. Afterwards possible changes can be implemented and the final system will be in production. A second test beam in 2002 is used to test the complete system. If necessary, after the test beam there is one more year available to implement modifications. The complete system will be available in the beginning of 2004.

13.9 Quality assurance and review procedures

A set of procedures and rules must be followed to assure the required quality of the system.

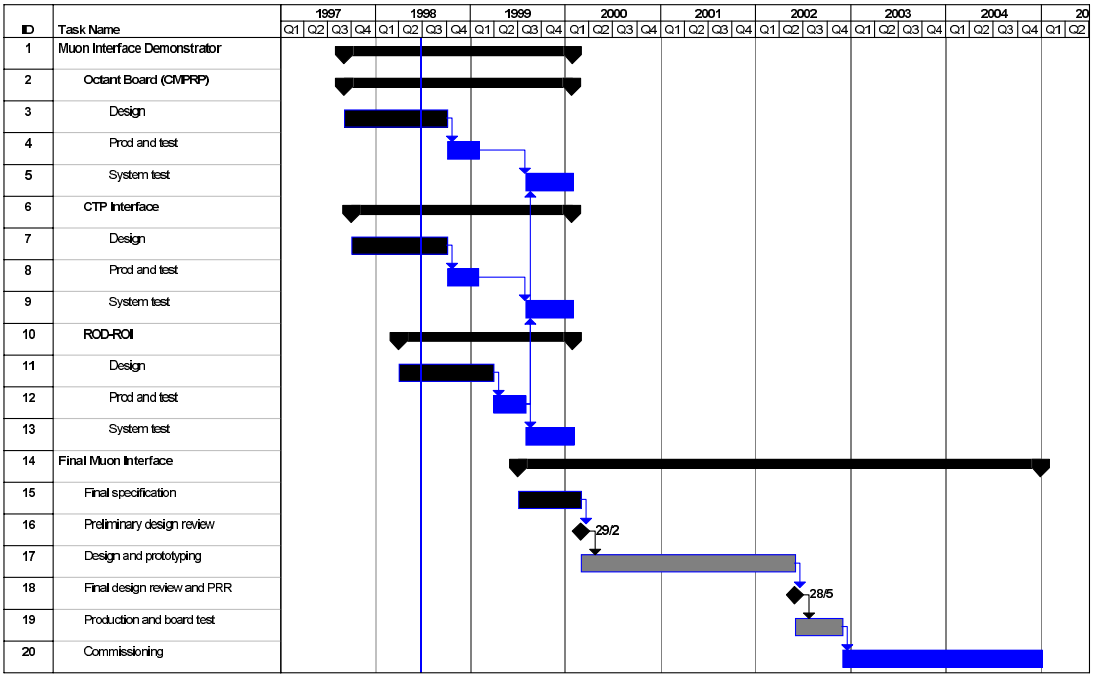

Figure 13-13 Time schedule for the MUCTPI.

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

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

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

The FDR will be held before the MUCTPI design is sent for manufacture, and is intended to catch any design or engineering errors before budgetary commitment. This review will necessarily be of a more technical nature than the initial review, but there should be few problems to detect by this stage. It will be merged with the ATLAS production readiness review (PRR).

This review work has already been done for the prototyping. A user requirement document (URD) for the LVL1 muon trigger system has been written [13-5] and reviewed a first time by representatives of each subsystem.

Prototypes of each element will exist and be tested in the coming year. During that time, the final specifications will be written The PRR will review them as well as the performance and functionality of the prototype system.

The construction of the MUCTPI components will follow usual rules, namely:

- Electrical tests of the PCB before assembly.
- Burn in of the boards after assembly by leaving them under power without cooling for two days.
- Test of the full boards.

Each element will receive a serial number and a database will be set up to store the information relative to each board: origin of the components which populate the board, results of tests, history of failure and location.

13.10 References

13-1 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, 1994.

13-2 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC 97–22, May 1997.

13-3 E. Petrolo et al., *The First Level Muon Trigger of ATLAS in the Barrel Region*, Third Workshop on Electronics for LHC, London, UK, September 1997.

13-4 O. Jinnouchi, O. Sasaki, *Study on muon level-1 trigger scheme for TGC*, ATLAS note DAQ–No–82, February 1998.

13-5 *LVL1 Muon Trigger User Requirements Document (Draft version 1.4)*, ATLAS working document, ATL–DA–ES–0002, March 1998.

13-6 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

13-7 P. Farthouat, *Current Understanding of the Muon LVL1 System*, version 2, ATLAS note DAQ–No–89, May 1998

13-8 O. Boyle et al., *The S-LINK Interface Specification*, CERN working document, March 1997.
http://www.cern.ch/HSI/s-link/spec/

13-9 J. Christiansen, *32 Channel general purpose Time to Digital Converter*, internal CERN note, January 1997.

14 Muon-trigger algorithms and performance

14.1 Introduction

In this chapter we present the main results on the expected physics performance of the muon-trigger system. A detailed simulation of the relevant detector components as well as the logic of the trigger has been made and is used to both optimize the system and confirm that the appropriate performance requirements, as outlined in Ref. [14-1], have been met.

The muon-trigger simulation is discussed (Section 14.2), and the analysis of the geometrical acceptance using this simulation reported (Section 14.3). The algorithms used by the trigger system to discriminate muon transverse momentum are introduced (Section 14.4) and the performance of the trigger analysed (Section 14.5). The trigger rates arising from known processes producing muons have been calculated (Section 14.6), as well as those from other particle fluxes present in the experimental hall (Section 14.7).

The discussion in this chapter closely follows that of [14-2], where a somewhat more detailed treatment of some topics can be found.

14.2 Simulation of the muon-trigger

In order to understand the performance of the ATLAS level-1 muon-trigger detailed simulations of both the ATLAS detector and trigger have been made. These include modelling of the passive material of the detector, of the active volumes and detector response of the muon-trigger counters, and detailed simulation of the hardware and logic of the muon-trigger itself. This has been done in the framework of the standard ATLAS simulation packages, using GEANT 321 [14-3] to describe the detector geometry in the DICE [14-4] program, and the ATRIG [14-5] package to simulate the trigger response.

14.2.1 Material of the detector and magnetic field

The material through which muons pass prior to traversing the trigger counters plays an important part in determining the acceptance of the trigger coincidence windows through the contribution of multiple scattering to the measured residual in the trigger chambers. The material in front of the trigger chambers is that of the end-cap and barrel toroids and the inner detector and calorimeters, both hadronic and electromagnetic. In both the barrel and the end-cap this material is dominated by the tile and liquid-argon calorimeters, as shown in Figure 14-1. Total material between the interaction point and the muon system constitutes between 10 and 15 absorption lengths, with somewhat more material in the region of the barrel-end-cap interface at $|\eta| \sim 1.15$. For lower momentum particles this material serves as a barrier between the interaction point and the trigger chamber planes; in the barrel, muons with p_T below 3 GeV are absorbed. All this material is included in the DICE simulation package. Material arising from cables, cooling components and other elements of the detectors that have yet to be finalized are not included in the simulation. These do however constitute an amount of material that can be neglected relative to that present in the current detector model.

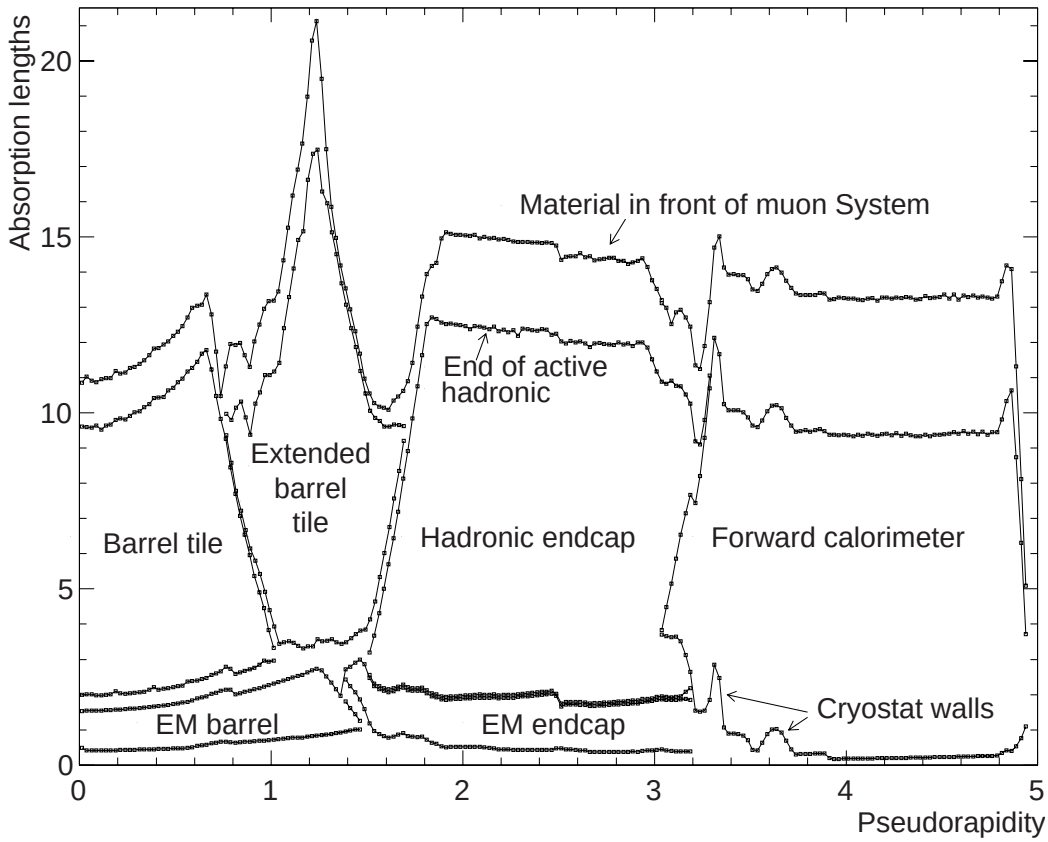

Figure 14-1 Contributions of the various ATLAS detectors to the absorptive thickness shielding the muon system from primary collision products.

A precise map is made of the three-dimensional magnetic field arising from the barrel and end-cap toroids as well as the field of the central solenoid and hadron calorimeter which acts as the return yoke. This enables the path of charged particles in the simulation to be accurately tracked. The magnetic configuration used to calculate the field map is described in [14-6].

14.2.2 The trigger chambers

In the ATLAS detector simulation particular care has been taken in describing the different technologies of the muon spectrometer: monitored drift tubes (MDT), constituting the precision chamber system[1], and resistive plate chambers (RPC) and thin gap chambers (TGC), forming the trigger system, are all simulated in detail with regard to the geometry and material composition. All detector units are divided into sensitive and non-sensitive regions. Layers of sensitive regions are interleaved with layers of passive material. The basic parameters and complete layout of all muon detectors are described in the ATLAS muon database [14-7]. A three-dimensional view of the simulated trigger system is shown in Figure 14-2 and a longitudinal view in Figure 14-3.

1. Cathode strip chambers (CSCs) provide the precision measurement of muon position and momentum for $\eta > 2$, in the first muon station. The CSCs are not modelled in the present simulation; MDT chambers have been simulated in their place.

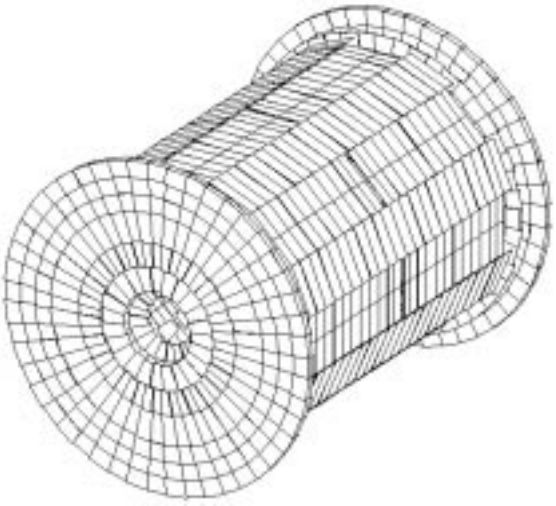

Figure 14-2 Three-dimensional view of the simulated trigger system: resistive plate chambers (RPCs) in the barrel region, thin gap chambers (TGCs) in the end-cap region.

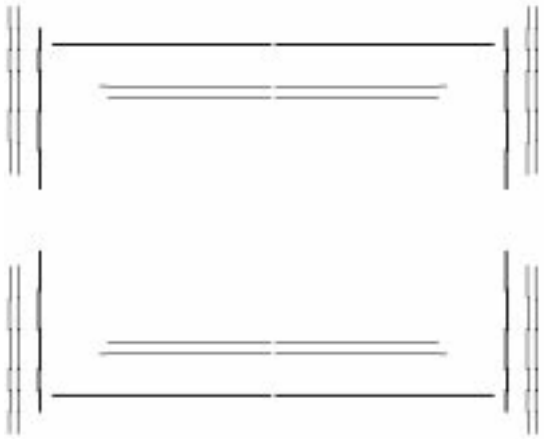

Figure 14-3 Transverse view of the simulated trigger system with a longitudinal cut at ϕ=0°.

Trigger signals for the level-1 muon-trigger are provided by the trigger chambers: RPCs in the barrel and TGCs in the end-cap. In order to understand the many triggering issues raised by the performance of the trigger chambers these chambers have been modelled in detail within the DICE package. For the TGC a detector response is assumed such that hits are produced on only a single wire. In the RPC cluster size is modelled with an upper limit of two strips hit per track. Full hit efficiency for charged particles traversing active volumes is assumed in both subsystems. Within the active volumes spacers within the gas gaps are modelled. Effects that cause the chambers to deviate from their nominal geometry, such as temperature related effects and chamber displacement or deformation are not simulated.

This simulation has been used to optimize the design of the chamber geometry and the read out segmentation. It exactly reproduces the currently envisaged configuration of the chamber layout. This precise modelling allows the investigation of geometrical issues such as overlaps and acceptance holes which are crucial to the understanding of the trigger performance. For each subdetector in the simulation one layer of strips or one layer of wires is defined as a sensitive plane and digitization is modelled using the GEANT tracking of particles through the simulated detector volumes to determine the hit position.

14.2.3 Simulation of the trigger logic

The digits from the simulation of the trigger chambers are passed to a detailed simulation of the hardware containing the trigger logic. The full logic of the trigger, as well as the coincidence matrices in which coincidences are formed, is modelled. The details of the trigger electronics are discussed in Chapters 11 and 12, and the modelling of the logic in the simulation in Section 14.4.2.

The only significant aspects of the trigger that are not included in the simulation used for these performance studies are the timing information and some simulation of expected noise.

Dedicated GEANT simulations and numerical calculations have been performed to evaluate the effect on trigger rates and performance from these phenomena [14-8], [14-9].

14.3 Muon-trigger system acceptance

14.3.1 Introduction

In an apparatus as large and complex as the ATLAS muon spectrometer it is important to evaluate the effect of the geometrical acceptance on the trigger performance [14-10]. The analysis of the geometrical acceptance of the trigger system is based on the layout description available in the ATLAS muon database [14-7]. The layout of the trigger system, RPCs in the barrel and TGCs in the end-cap are shown in Figures 14-2 and 14-3.

14.3.2 Trigger chambers

The trigger system in ATLAS is composed of the RPC ($|\eta| < 1.05$) and the TGC ($1.05 < |\eta| < 2.4$) subsystems (Chapter 10). The acceptance of the low-p_T trigger in the barrel system is defined by those tracks producing hits in at least three of the four inner trigger planes, and that of the high-p_T trigger by those tracks within the low-p_T acceptance and giving hits in at least one of the two planes of the outer station. For the TGCs the low-p_T acceptance is defined by tracks with three from four possible hits in the outer two trigger stations (two doublet units); the high-p_T trigger acceptance by tracks within the low-p_T acceptance and having at least two hits in the three planes of the inner trigger station (triplet unit).

In Figure 14-4 the geometrical acceptance for very high-p_T muons as a function of η, integrated over the azimuthal angle ϕ is shown for both the low- and the high-p_T configurations. This calculation considers only the geometry of the trigger system, and all generated tracks traversing the appropriate sensitive regions are considered to lie within the acceptance. No treatment of trigger coincidence windows is included in the analysis. The acceptance is limited by losses due to the presence of magnet and support structures (support feet of the inner detector and calorimeter and coil support ribs), and space allowed for services.

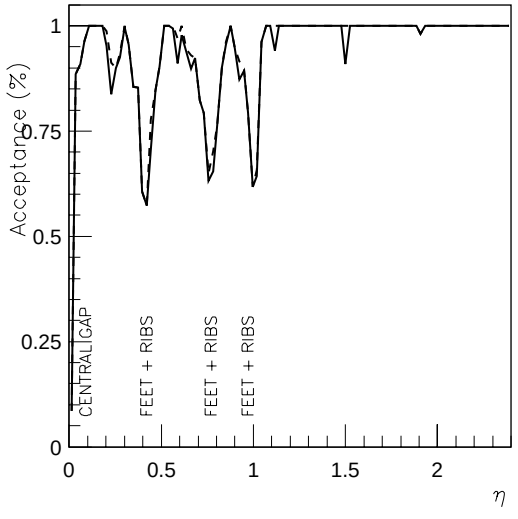

Figure 14-4 Trigger acceptance integrated over ϕ as a function of pseudorapidity for low-p_T (dashed line) and high-p_T (solid line) trigger settings. The acceptance loss at $\eta = 0$ can be seen as well as that due to support structures (feet and ribs).

The acceptance as a function of ϕ, averaged over η, in the range $|\eta| < 2.4$, is shown in Figure 14-5. The loss in acceptance due to the feet is seen in the two regions around 270 degrees, as well as the eight-fold modularity due to detector and magnet structure. Tracking very high-p_T muons through the

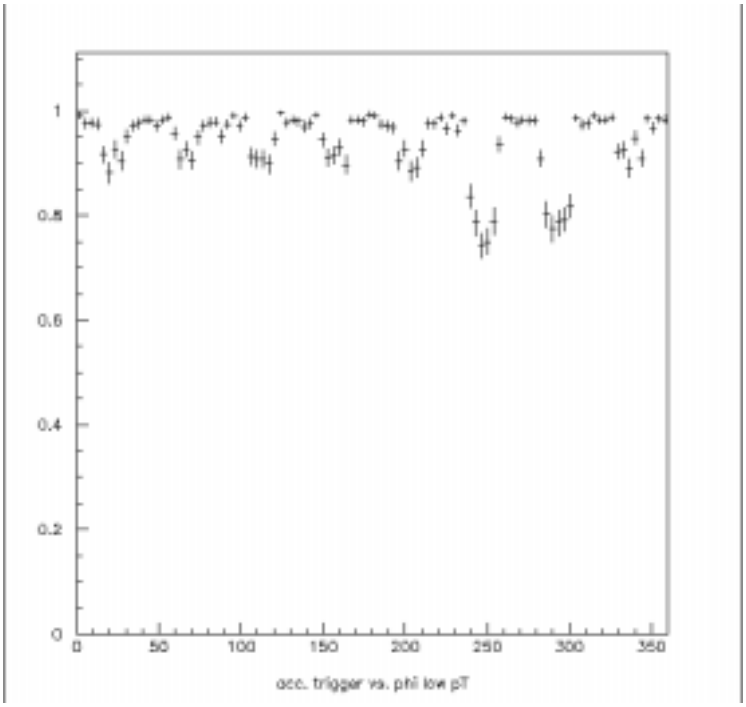

Figure 14-5 Trigger acceptance as a function of ϕ, averaged over η, in the range $|\eta|<2.4$. The structures around 270° arise from the support feet and whilst the periodic acceptance loss is from the coil support ribs.

trigger system, the total acceptance loss in the region $|\eta|<2.4$ and integrated in azimuth is ~6% for the low-p_T coincidence, whilst for the high-p_T coincidence the acceptance loss is ~6.5%. The average acceptance of the muon-trigger system in the region ($|\eta|<2.4$) is shown in Table 14-1, normalized to the system coverage ($|\eta|<2.4$). In these calculations the contributions to the acceptance loss due to alignment corridors (which necessitate holes in the trigger counters) and holes in the TGCs for the displacement cables of the inner detector are neglected. The fraction of single muons that are double counted, considering only the geometrical overlap between the barrel and end-cap systems in the transition region, is ~0.77% and ~0.15% for the low- and high-p_T coincidences respectively, within the trigger acceptance ($|\eta|<2.4$). It should be emphasized that logic exists in the MUCTPI (Chapter 13) to ensure that only a single level-1 trigger is given for such triggers; this calculation excludes the effect of such logic.

14.3.3 Transition region

At the transition between the barrel and the end-cap subsystems ($|\eta|$ ~1.05) a non-pointing gap, along the z-axis has been left, as can be seen in Figure 14-3. This allows permanent access to the barrel volume and for the passage of services. This region has been investigated to determine the level of acceptance loss and double-counting.

The trigger acceptance and the probability that single muons are identified in both subsystems in this region have been

Table 14-1 Average acceptance of the Level-1 muon-trigger system, normalized to the region $|\eta|<2.4$, (neglecting the alignment corridors and rail holes).

	barrel	**end-cap**	**global system**
low-p_T trigger	0.909	0.976	0.940
high-p_T trigger	0.885	0.974	0.935

investigated in the region $0.8 < \eta < 1.2$ accounting for magnetic field and the material of the inner detectors. This shown in Table 14-2 for various p_T values.

Table 14-2 Double counting as a fraction of acceptance in the region $0.8 < \eta < 1.2$, in %, in the trigger system for muons of 500, 20, 6 GeV p_T. The acceptance shown is for μ^+, with the results for μ^- in brackets. (The acceptance for μ^+ and μ^- shown here is reversed for negative η.)

(0.8< η <1.2)	Double-counting rate		Acceptance	
p_T (GeV)	low-p_T (%)	high-p_T (%)	low-p_T (%)	high-p_T (%)
500	4.55 (4.2)	1. (1.)	93.1 (93.0)	92.3 (92.6)
20	8.2 (1.2)	3.2 (.3)	95.9 (90.2)	95.5 (89.0)
6	23.5 (0.)	6.9 (0.)	100.0 (77.0)	99.8 (78.2)

In the present layout an additional layer of RPCs has been added at the end of the barrel in the sectors covered by large chambers in such a way as to make the system as projective as possible in the bending plane. The additional chamber substantially increases the trigger acceptance in the region $1.0 < |\eta| < 1.1$, relative to the layout described in [14-6].

The present layout, used for all acceptance calculations here, does not include chambers around the coil support ribs. A modified layout has been investigated where very small chambers are added either side of each rib at $\eta = 0.3$, 0.7 and 1.0 in all six small standard sectors. The dimension of these new chambers is about $760 \times 300\ mm^2$ (active area) and their addition would result in an increase in acceptance of about 4%, normalized to the barrel η region $|\eta| < 1.05$.

14.3.4 Higgs to four muons: studies of the detector acceptance

Higgs boson decays to four muons are a benchmark process for muon detectors at the LHC. In order to evaluate in detail the effects of the muon-trigger system acceptance on the Higgs detection capabilities of ATLAS, we have tracked a sample of Standard Model Higgs decays to four muons ($H \rightarrow ZZ^* \rightarrow 4\mu$) through the trigger detector simulation, for several values of the Higgs boson mass in the range 120–180 GeV. We have estimated that even for a dimuon-trigger more than 99% of the Higgs events generated within the trigger acceptance ($|\eta| < 2.4$), for all generated Higgs masses, have two muons in regions covered by the trigger counters. Clearly the single muon-trigger will have an efficiency for triggering such events extremely close to 1. Details of the acceptance subject to demanding three or four detected muons are shown for a 140 GeV mass Higgs boson in Table 14-3.

Table 14-3 Geometrical acceptance of the Level-1 muon-trigger for the process $H \rightarrow ZZ^* \rightarrow 4\mu$, for a 140 GeV mass Higgs boson. Acceptance is shown as a function of the number of muons lying within the trigger acceptance.

140 GeV Higgs mass	Low-p_T acceptance	High-p_T acceptance
Four muon detection	0.709	0.681
Three muon detection	0.963	0.955
Two muon detection	0.999	0.998

14.4 Track finding algorithms

14.4.1 Overview

The level-1 muon-trigger is based on the measurement of muon trajectories in two or three different planes (called stations). Muons are bent by the magnetic field generated by the toroids and their angle of deflection depends on their momentum, the field integral along their trajectory and the Coulomb scattering in the material lying in front of the trigger planes. The energy loss fluctuation is also important for low-p_T transverse-momentum triggers.

The differences from the infinite momentum track, 'residuals', are measured using three trigger stations, Figure 14-6. The trigger plane farthest from the interaction point, in the end-cap, and

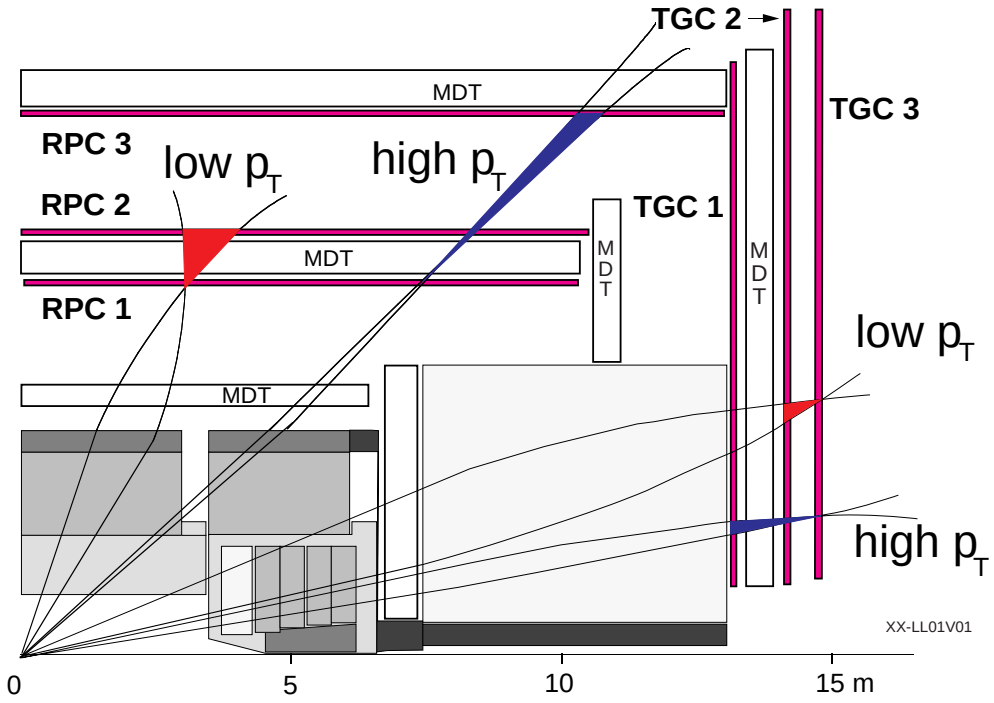

Figure 14-6 The level-1 muon-trigger scheme.

nearest to the interaction point in the barrel, is called the pivot plane and two different lever arms from the pivot to the other two trigger planes provide two different measurements of the residuals. The two different lever arms allow trigger thresholds to cover a wide range of momenta: the shorter lever arm (the pivot and station 2) covers a lower-momentum region and the longer one (the pivot and station 1 for the end-cap, the pivot and station 3 for the barrel), a higher-momentum region.

The residual distribution in a region (r, ϕ) on the pivot plane reflects the momentum spectrum of muons passing through that region. The level-1 muon-trigger uses the residual distribution in order to discriminate muons with transverse momentum above some threshold from those below the threshold.

Each hit found in station RPC1 (TGC3) is extrapolated to station RPC2 (TGC2) along a straight line through the interaction point. A coincidence window is then defined around this point, where the window's size depends upon the required p_T threshold. The low-p_T trigger condition is then satisfied if, for both projections

- there is at least one hit within the coincidence window
- at least one of the two low-p_T stations has hits in both trigger planes, to satisfy the three-out-of-four majority logic.

The size of the coincidence window defines the p_T threshold applied in the trigger: the broader the window the lower the threshold.

A similar procedure is performed for the high-p_T trigger, where the planes of RPC3 (TGC1) together with the pivot plane are used. The high-p_T trigger is satisfied if

- the track passes the low-p_T criteria, and
- in the barrel at least one hit in the two trigger planes of RPC3 are in coincidence, and in the end-cap if at least two of the three planes of TGC1 in the η view, and one of the two planes of TGC1 in the r–ϕ view, are within the appropriate coincidence window.

In the end-cap, read out segments (wire groups and strips) in the three trigger stations are aligned projectively so as to point to the nominal primary-vertex position. Wire groups in the three successive active layers of a triplet unit in station 1 are staggered by one third of their width relative to one another. Those in the two active layers of the doublet units of TGC2 and TGC3, and the strips in all stations, are staggered by half their width. This staggering makes effective position resolution one third or one half of the width of the actual read out segment.

In the barrel strip segments are similarly projective. Staggering is by approximately half a strip width in the central region of the detector ($\eta < 0.6$) and by somewhat less in the higher η regions.

The muon-trigger is divided into regions in η–ϕ where independent thresholds can be set, i.e. where independent trigger windows can be used. This segmentation is discussed in Sections 11.1.2 and 12.3.

A time coincidence among hits is also required, to

- identify the bunch crossing,
- reduce the trigger rate from accidental coincidences induced by the cavern background.

The timing procedure of each of the two subsystems is described in Sections 11.2.4.1 and 12.4.6; the proposed procedure is equivalent to the setting of a time gate of 22 ns in the barrel and 31 ns in the end-cap, where logic exists to ensure that the appropriate bunch crossing is triggered.

14.4.1.1 Declustering algorithm

Due to charge spread, cross talk or background, it is possible that a few neighbouring read out segments (wire groups or strips) may fire. Such hits can potentially produce multiple triggers. To reduce the combinatorial effect on the number of triggers, a declustering procedure is applied in both the barrel and end-cap trigger logic. In the barrel only the central hit (or two hits, in the case of an even number) is allowed to participate in the formation of trigger coincidences. The maximum multiplicity that can be handled is a system parameter; above that,

a certain number of hits will survive, depending on the reduction rules. In the end-cap, centre finding logic is implemented which selects only the second channel of sequential hits.

14.4.2 Simulation of the trigger logic

The purpose of the trigger simulation is not only to evaluate the performance of the system, but also to optimize the trigger logic design. The trigger simulation should be flexible and duplicate in detail the hardware of the trigger logic in order to optimize parameters such as read out segmentation and the size of the coincidence matrices, as well as to debug the logic itself. This optimization is fed back to the design of the logic.

All the logic of the hardware of the trigger system is duplicated in the simulation of the trigger logic used for the performance studies of this document. All coincidences are performed using the appropriate granularity of the relevant trigger matrix and the rules and constraints on the allowed multiplicity in the logic at any single point are followed within the simulation. In the barrel, the staggering of the RPC strips and the declustering algorithm are not yet included in the simulation used here. No timing information is currently included in the simulation of either subsystem; this is not likely to impact significantly on the trigger performance since the timing has been optimized to include all hits from true prompt muon. Dedicated studies have been made to investigate this in both barrel and end-cap systems [14-8], [14-9].

14.5 Trigger coincidence windows and efficiency

14.5.1 Transverse-momentum resolution

The coincidence windows of the level-1 muon-trigger are designed to offer momentum discrimination by including muons above threshold and excluding those below threshold which will be bent beyond the coincidence window. The p_T resolution at threshold in turn is a function of the trigger detector geometry (in particular the lever arm between stations 1 and 2, and station 3 – see Figures 11-1 and 12-1), the magnetic field intensity and its inhomogeneities, the Coulomb scattering in the central calorimeter and the width of the interaction region. Whilst this is mostly relevant to the primary bending direction (r or η) there is also some bending in the orthogonal ϕ-plane.

The trigger scheme is shown in Figure 14-6. In the barrel, where a muon is seen to leave a hit in both of the low-p_T stations the projective extrapolation of the hit in station 1 (assuming an origin at the nominal p–p interaction point) to station 2 can be made assuming the muon path to be a straight line (i.e. the infinite momentum approximation). The distance between the extrapolated point in station 2 and the actual hit detected is defined as the separation parameter, d. In the end-cap system the low-p_T trigger coincidence is made using station 3 (the pivot plane) and station 2.

The histograms in Figure 14-7 show as an example the value of the signed separation parameter d, for $p_T = 6$ GeV muons, in a sector of the end-cap trigger chambers at $\eta = 1.5$. The two peaks shown in the histogram correspond to positive and negative muons. In the ideal case two narrow peaks are expected, approximately equidistant from and either side of $d = 0$, one for each muon charge. In reality, the physics effects mentioned above introduce a significant smearing of these distributions and as a consequence the two peaks are broadened. This coincidence plot is a useful estimator of the p_T resolution at any trigger threshold relative to the bending seen in the trigger counters. The method followed consists of fitting each peak with a Gaussian function with parameters σ and μ; the parameter μ is the average value of d (for a given muon charge) and is proportional to $1/(p_T - p_T0)$ (where p_T0 is the average p_T lost in traversing the calorimeter), while σ, the width of the Gaussian, represents the physics limitations on the measurement of d. Therefore the ratio σ/μ is a measure of the transverse-momentum resolution, whilst μ quantifies the amount of bending. In order to evaluate the contribution of each source of smearing, the parameter d has been plotted with the relevant physics effects introduced into the simulation one at a time.

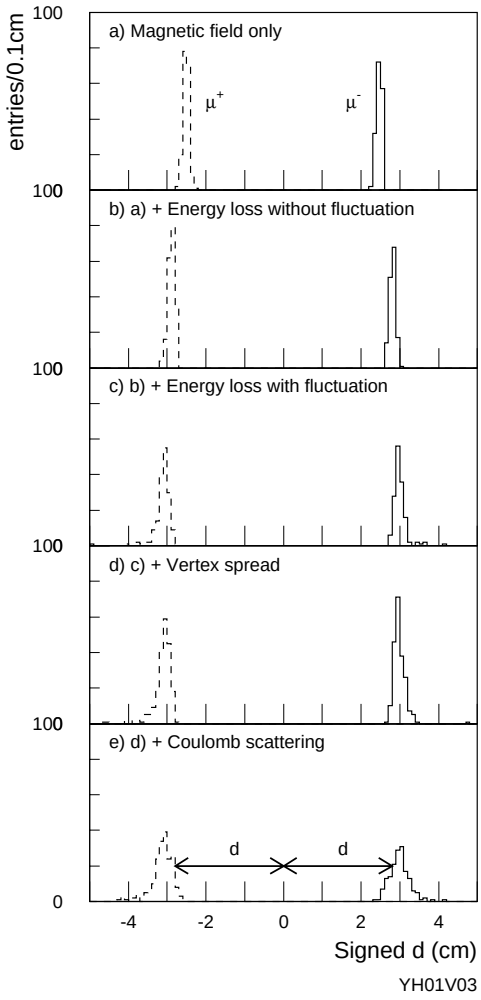

Figure 14-7 Coincidence plot for 6 GeV p_T muons, both charges, generated in the end-cap close to $\eta = 1.5$, $\phi = 0.0$–$2.0°$. Several histograms are shown in the same picture, representing: a) magnetic field only, (no vertex spread, Coulomb scattering, energy losses or detector response); b) as a) + energy loss without fluctuations; c) as a)+ energy loss with fluctuations; d) as c) + vertex spread; e) as d) + Coulomb scattering. The detector granularity is not considered in this plot.

Throughout the barrel detector Coulomb scattering is the most important physics effect limiting the precision of the p_T measurement for the low-p_T, although its effects are somewhat η– and ϕ-dependent. The typical resolution figure is 30%, rather uniform throughout the η–ϕ space covered.

Figures 14-7 and 14-8 show the coincidence plot for 6 GeV p_T muons in the end-cap system; the typical momentum-resolution figure is about 10% where this is mainly limited by the Coulomb scattering in the calorimeter. However, there are ϕ regions in the η interval $1.0 < |\eta| < 1.8$ where this resolution is significantly worse. This effect is due to the combination of the magnetic fields produced by the barrel and end-cap toroid coils. Since in these regions the field integral is low, the associated d parameter is low, and the resulting muon momentum resolution $\delta p_T/p_T = \sigma/\mu$ is worse than elsewhere.

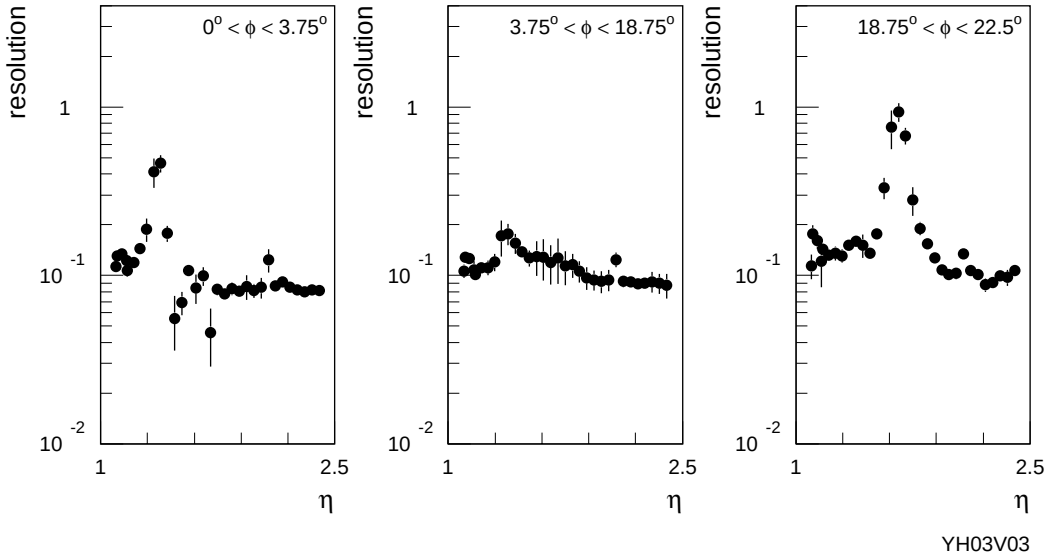

Figure 14-8 Transverse-momentum resolution, $\delta p_T/p_T$, of the end-cap system for 6 GeV p_T muons as a function of η and ϕ.

This is shown in Figure 14-9 (showing data in the region $\eta = 1.3$, $\phi = 0.0°$–$2.0°$) where the two opposite-sign muon peaks are almost overlapped. The result of this is that almost all p_T discrimination in such regions is lost.

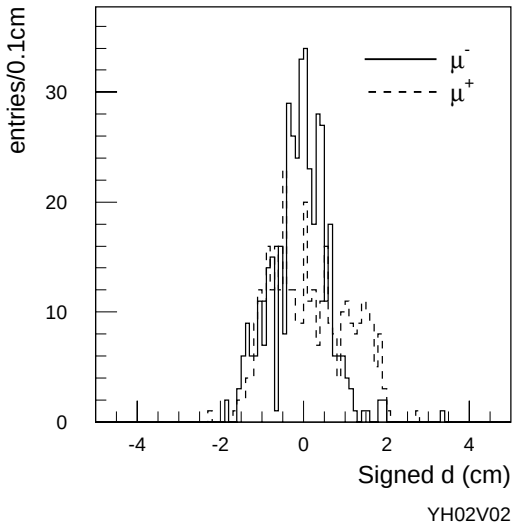

Figure 14-9 Coincidence plot for 6 GeV p_T muons, for both charges, generated at $\eta = 1.3$, $\phi = 0.0$–$2.0°$. The two charged-muon peaks are overlapping each other almost completely, demonstrating the loss of almost all p_T discrimination.

As shown in Figure 14-10 these poor-resolution regions are located along the barrel and end-cap coils, with η values of ~1.3 (ϕ coordinates of the end-cap coils) and ~1.6 (ϕ coordinates of the barrel coils). Figure 14-8 shows the p_T resolution in the end-cap as a function of η and ϕ, for three different ϕ regions across the whole η range of the end-cap acceptance. In the barrel there are no such overlapping fields, and as such the resolution is rather constant across both η and ϕ.

For the high-p_T trigger the p_T estimation is made using hits in station 1 (station 3) and station 3 (station 1) in the barrel (end-cap) system. The trigger scheme adopted also uses the coincidence in station 2 for track confirmation; this is to reduce the trigger rate from fake muons, rather than to offer any improvement in the momentum discrimination of the trigger. Here the resolution is limited mainly by the Coulomb scattering and the width of the interaction region ($\sigma_{vtx} \sim 5.5$cm). At p_T = 20 GeV the transverse-momentum resolution is $dp_T/p_T \sim 0.40$ in the barrel and $dp_T/p_T \sim 0.15$ in the end-cap. The higher performance of the end-cap system is largely due to the higher field integral provided by the forward toroids [14-6] and the different trigger counter layout that allows the end-cap trigger system to make full use of the available field integral; indeed, the barrel trigger counters are placed inside the magnet, while the end-cap trigger system is located

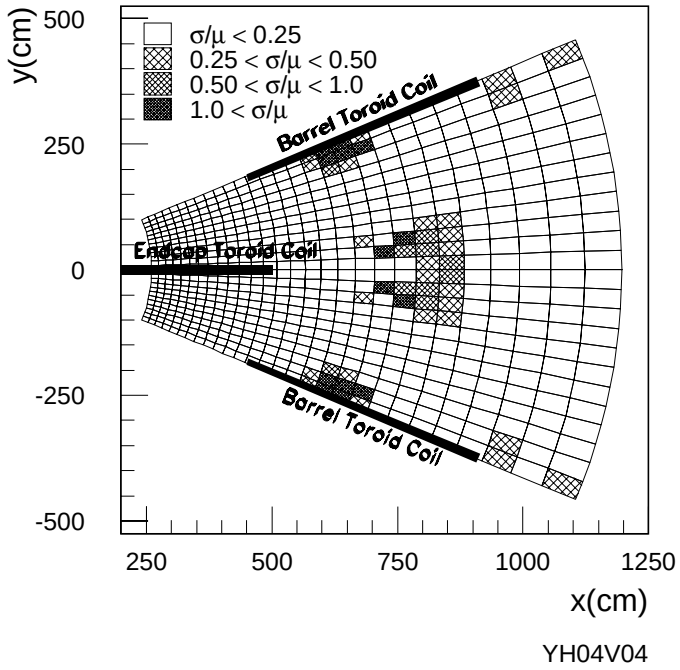

Figure 14-10 Position of the regions in the end-cap system where momentum resolution is poor. A single octant is shown where the subdivisions represent trigger sub-sectors, and the dark lines show the position of the barrel and end-cap toroid coils.

beyond the toroids. The high-p_T trigger in the end-cap system shows the presence of the same low field-integral regions as seen for low-p_T muons.

Plots of the *d* parameter are used to determine the coincidence windows for high- and low-p_T trigger thresholds. The resolution plot is used to identify regions where the p_T discrimination of the trigger system is poor.

14.5.2 Coincidence windows

The sizes of the trigger windows that determine the p_T threshold have been computed for the proposed trigger system by tracking single muons through the ATLAS detector using Monte Carlo simulation techniques (see Section 14.2). The size of a coincidence window is defined such that 90% of the muons of each charge within the detector acceptance, generated with transverse momentum equal to the threshold, are accepted. The absolute normalization is determined by the trigger efficiency for muons with momentum far above the threshold, and corresponds to the geometrical acceptance of the trigger detector.

The coincidence windows of the trigger will ultimately be defined as areas in two dimensions on the surface defined by inputs to the coincidence matrix in the two projections η and ϕ. The boundaries of the window for any one trigger threshold will be determined using simulated muons, and once data have been collected using true muons in ATLAS. The studies made here have approximated the two-dimensional trigger window to a simple rectangle. This model reproduces exactly the barrel coincidence logic and represents a good approximation to the end-cap trigger coincidence logic.

As discussed in Section 14.5, the window size depends on the trigger-counter geometry, the magnetic-field intensity, the Coulomb scattering and the energy-loss fluctuation in the central calorimeter, and the width of the interaction region (σ_{vtx}~5.5 cm). The detector granularity also contributes to the momentum resolution; it has been designed by optimizing the opposing requirements of trigger performance and number of trigger channels. Since the geometry and the large variations in both granularity and field strength in the end-cap make this region more complex these studies have concentrated on that subsystem.

The sizes of coincidence windows for the 6 GeV low-p_T threshold in the barrel are shown in Figure 14-11 for standard sectors. Since the position of trigger stations of the special small

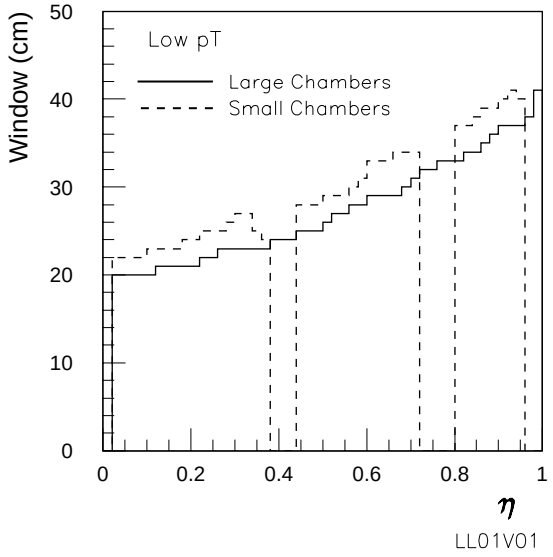

Figure 14-11 Window size in the barrel as a function of η in both the small and large chambers, for 6 GeV muons in the bending plane of the low-p_T system.

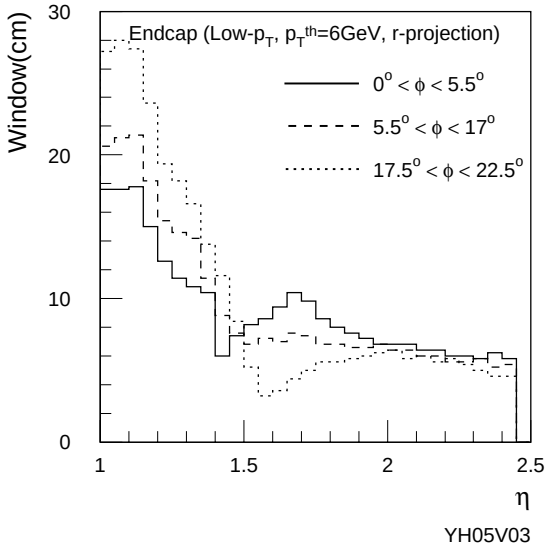

Figure 14-12 Window size in the end-cap as a function of η and φ for 6 GeV muons in the bending plane of the low-p_T system. Three φ intervals are considered: the end-cap toroid region (0° to 5.5°), the barrel toroid region (17° to 22.5°), and the region between the two (5.5° to 17°).

chamber sectors is the same as the standard small sectors (they differ only in that the coverage along the beam-line direction is reduced), the coincidence-window size is the same so long as muons cross both trigger stations of the low-p_T system.

In the end-cap low-p_T system the coincidence windows show a pronounced φ-dependence. This is due to the complex magnetic fields produced by the combination of the barrel and end-cap toroid coils. Neighbouring η–φ regions are seen with significantly different field integrals, and this has a direct impact on the size of the coincidence windows.

Figure 14-12 shows the size of the coincidence window, in the η projection, in the end-cap as a function of η and φ, for a threshold of 6 GeV p_T in the low-p_T trigger system. These are optimal window sizes calculated independently of the actual detector granularity. The η and φ values are those of the track position on the pivot plane.

The coincidence window is not symmetric around the extrapolated point of the infinite momentum path: Muons of a given charge, bending towards high $|z|$ regions (in the barrel) or high r regions (end-cap), cross station 2 at a distance from the extrapolated point larger than that for muons of opposite charge. This effect increases with $|\eta|$ in the barrel and with r in the end-cap. This leads to asymmetric coincidence windows, formed by independent left and right half-windows. The difference between the right and left half-window is 7 / 10 cm in the large / small chambers at $|\eta|$~0.9 in the barrel and 3.5 cm at $|\eta|$ close to 1 in the end-cap. The size of each half-window is evaluated accounting for the relative trigger efficiency of positive and negative muons independently, to minimize possible effects on charge-asymmetry measurements.

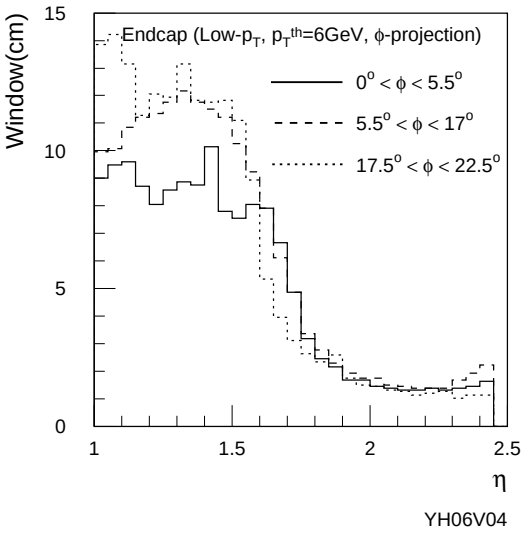

Figure 14-13 Window size in the end-cap for 6 GeV muons in the low-p_T system, r–ϕ projection as a function of η and ϕ. Three ϕ intervals are considered: the end-cap coil region (0° to 5.5°), the barrel coil region (17° to 22.5°) and the region between the two (5.5° to 17°).

The coincidence window for 6 GeV transverse-momentum thresholds in the r–ϕ projection is shown in Figure 14-13 for the end-cap; in this case the window is defined such that

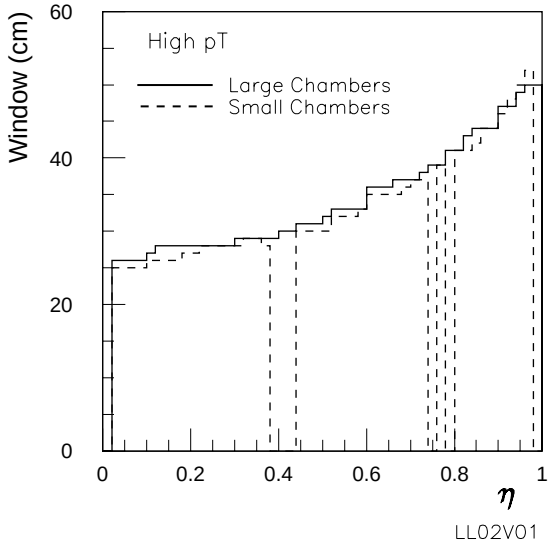

Figure 14-14 Window size in the barrel system as a function η in both the small and large chambers, for 20 GeV muons in the η projection.

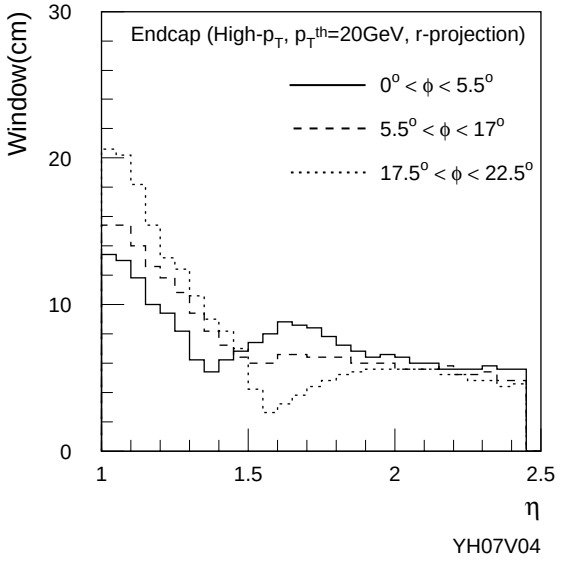

Figure 14-15 Window size as a function of η and ϕ for 20 GeV p_T muons in the η projection end-cap high-p_T system. Three ϕ intervals are considered: the end-cap coil region (0° to 5.5°), the barrel coil region (17° to 22.5°) and the region between the two (5.5° to 17°).

more than 99% of the muons within the geometrical acceptance are accepted. The coincidence in the r–ϕ plane serves to reduce the fake muon-trigger rate induced by the cavern background. In

the final windowing scheme both η and *r*–ϕ windows will be applied simultaneously in a true two-dimensional window. In Figure 14-13 it is possible to see the effect of the non-toroidal field in the transition region: In this region, muons also bend in the *r*–ϕ view, whereas, in an ideal toroidal system, only the Coulomb scattering in the calorimeter and the field in the central solenoid would be expected to produce deviation from the original muon trajectory.

A similar analysis has been performed for the nominal high-p_T threshold, p_T^{th} = 20 GeV. Figure 14-14 shows the size of this coincidence window, in the η projection, for the barrel, and Figure 14-15 shows the same projection for the end-cap. The same coincidence in the other projection is shown in Figure 14-16 for the end-cap. Here also the ϕ-dependence of the trigger window size in the end-cap system is visible. For the r–ϕ view, the coincidence window is again defined such that more than 99% of the muons with $p_T^{th} = p_T^{th}$ are accepted.

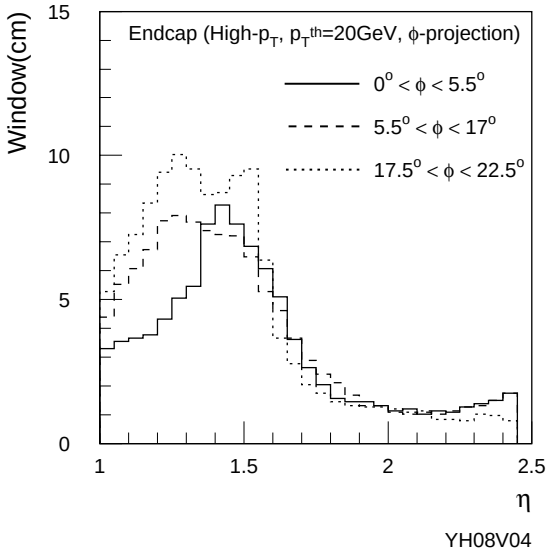

Figure 14-16 Window size in the end-cap as a function of η and ϕ for 20 GeV muons in the high-p_T system, r–ϕ projection. Three ϕ intervals are considered: t the end-cap coil region (0° to 5.5°), the barrel coil region (17° to 22.5°) and the region between the two (5.5° to 17°).

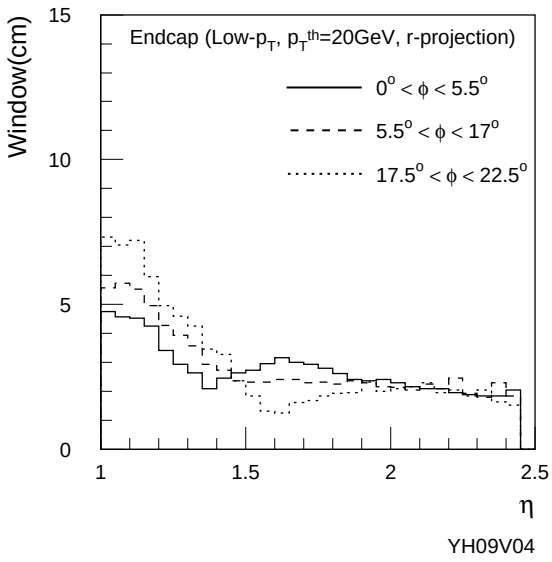

Figure 14-17 Window size in the end-cap as a function of η and ϕ for 20 GeV muons in the low-p_T system, bending projection. Three ϕ intervals are considered: the end-cap coil region (0° to 5.5°), the barrel coil region (17° to 22.5°) and the region between the two (5.5° to 17°).

Figure 14-17 shows the size of the window in the η projection for 20 GeV p_T muons in the low-p_T end-cap system. The window is approximately a factor of five smaller than that for 6 GeV muons, reducing significantly therefore the random trigger rate. The evaluation of the coincidence window size and the intrinsic momentum resolution has a direct impact on the optimization of the trigger detector granularity, largely with regard to the primary bending projection. The detector pitch should be fine enough that no significant deterioration of the intrinsic resolution is introduced by the system, and yet cost constraints favour a coarser pitch and fewer channels. In the barrel, the typical window size is around 35 cm, both for low and high-p_T threshold, with an intrinsic momentum resolution of about 40% at low-p_T threshold and about 15% for the high-p_T threshold. A read out pitch of the order of 3–4 cm, as proposed in this document, is adequate to maintain the physics performance of the system. Conversely, for the end-cap system the coincidence window can be as small as 6cm, with an intrinsic resolution of 10% at low-p_T threshold and about 15% for the high-p_T threshold. A wire-group size of about 14 (21) mm for the doublet (triplet) TGC stations, corresponding to an effective read out pitch of

7mm (from the effective resolution of half a wire-group, see Section 14.4), introduces an additional small smearing effect, reducing the resolution slightly, but still quite adequate for the trigger.

The detector granularity in the other projection is less crucial for the trigger performance and it has been optimized to provide an accurate second coordinate track measurement for muon reconstruction.

14.5.3 Trigger efficiency and masking

In order to evaluate the level of rejection of muons below any one trigger threshold by the trigger system, single muons over a wide range of momenta have been generated and passed through the ATLAS detector and trigger simulation. The barrel and end-cap systems have been studied independently, and the trigger efficiency evaluated as a function of p_T. In the end-cap where the window size and the trigger efficiency have a strong η dependence, the efficiency has been evaluated as a function of η (integrating over ϕ). These calculations are performed in a fiducial rapidity range within the geometrical acceptance of each system.

The trigger efficiency has been evaluated by simulating the trigger logic using coincidence windows defined as explained in the previous section. The efficiency is given simply by the ratio of the number of triggered muons to the number generated within the η fiducial interval, including also geometrical acceptance effects. The trigger efficiency is plotted in Figures 14-18 and 14-19 for the 6GeV p_T threshold in the low-p_T system, for both muon charges, as a function of p_T, for the barrel and end-cap.

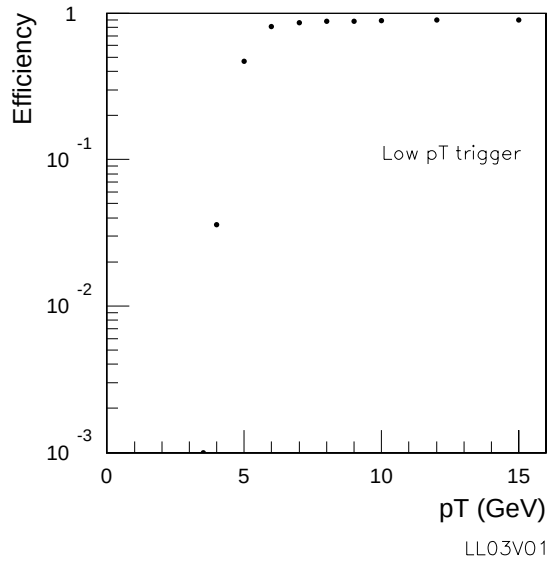

Figure 14-18 Trigger efficiency for the barrel for the 6 GeV threshold in the low-p_T system.

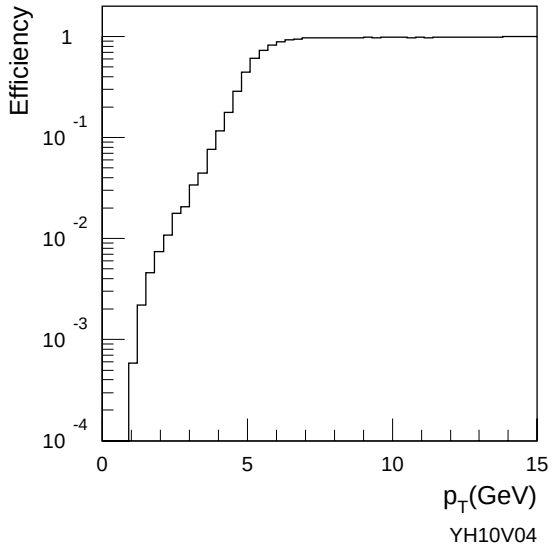

Figure 14-19 Trigger efficiency for the end-cap for the 6 GeV threshold in the low-p_T system.

Figure 14-20 shows the trigger-efficiency curve for η around 1.3 in the low-p_T end-cap system as a function of ϕ. It can be seen that the discrimination of the trigger in the ϕ interval with poor resolution is weaker than for other ϕ values. This is the consequence of the regions with low field integral, as discussed in Section 14.5.1.

The trigger efficiency for low-p_T (<6 GeV) muons in the barrel is lower than in the end-cap since in the barrel the central calorimeter acts as a barrier to muons with $p_T < 3.0$ GeV; this threshold being the average energy loss for a muon that traverses the calorimeter at normal incidence. In the end-cap the calorimeter geometry differs and the energy absorption sets a limit on p rather than p_T: $p_T > E_{loss} \times sin\theta$, where E_{loss} is about 4 GeV and $sin\theta$ can be as small as 0.17, thus bringing the calorimeter threshold below 1 GeV.

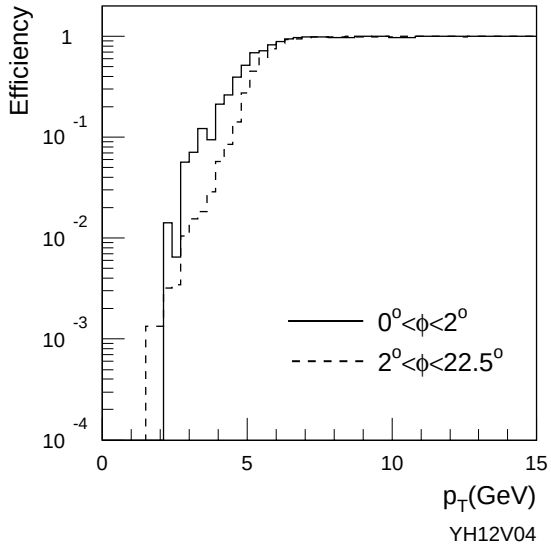

Figure 14-20 Trigger-efficiency curve for 6 GeV threshold in the end-cap low-p_T system, for muons generated with η around 1.3 for a) ϕ = 0–2°; b) ϕ = 2–22.5°

The level of precision of the calculation of trigger efficiency at very low p_T ($p_T < 2$ GeV) is of particular importance for the estimation of the trigger rate induced by muons from in-flight decays of light mesons. The cross-section for this process is large and very steeply rising with decreasing p_T around $p_T = 2$ GeV (see Section 14.6.1.1). This aspect is of particular importance for the end-cap system, where low-p_T muons are not naturally removed by the material of the central calorimeter. The present statistical error on the efficiency at low-p_T is about 15%. Efforts have been made to cross-check both the simulation of the trigger and the mechanics of the efficiency calculation. An entirely independent implementation of the simulation of the end-cap trigger chamber geometry and trigger logic has been made and the efficiency calculations presented here duplicated. It is seen that the independent calculation reproduces overall efficiencies presented here, including those regions where most rate is seen ($p_T < 2$ GeV), at the 10% level. This suggests that systematic errors on the efficiencies arising from coding and simulation uncertainties are at or below this level.

Efficiency curves, as a function of p_T are shown in Figures 14-21 and 14-22 for $p_T^{th} = 20$ GeV in the high-p_T system for the barrel and end-cap.

For both low- and high-p_T trigger thresholds, triggered muons with generated momentum below threshold have been studied in detail in order to verify the performance of the simulated trigger logic. These studies have focused on the end-cap system since here the trigger performance is more sensitive to the complex magnetic field in the forward region than to the more uniform field in the barrel.

Figure 14-23 shows hit positions in the pivot plane of the low-p_T end-cap trigger system (station 3) for triggered muons with p_T in the range 1 to 5 GeV where the threshold value was 6 GeV. A significant fraction of triggered muons below threshold are concentrated in those η–ϕ regions where the p_T resolution is poor, see Figure 14-10. However these regions get more diffuse for lower-p_T muons.

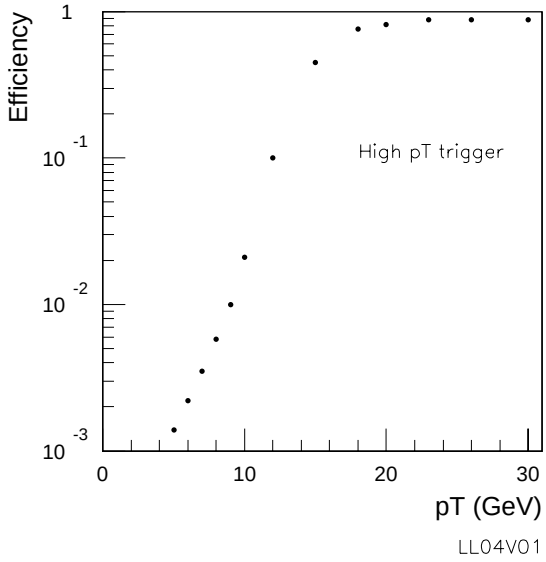

Figure 14-21 Trigger efficiency for the barrel system as a function of p_T, for the 20 GeV high-p_T threshold.

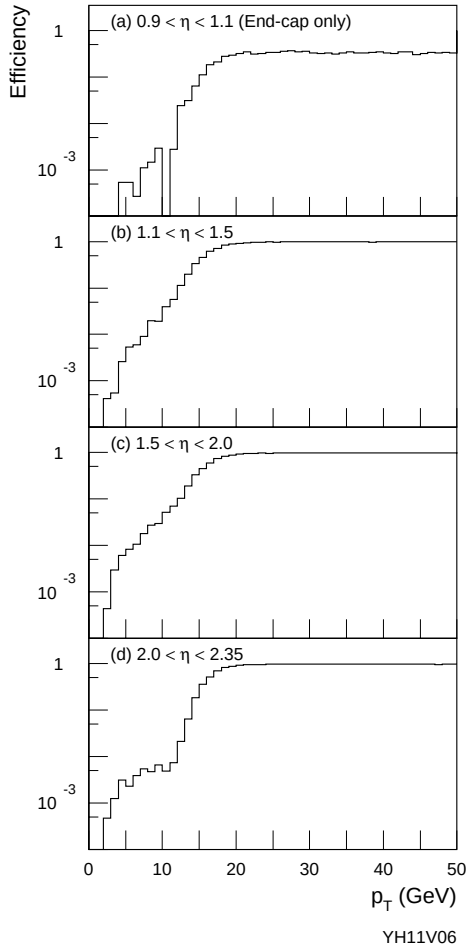

Figure 14-22 Trigger efficiency for the end-cap system as a function of p_T for the 20 GeV p_T threshold, for various η ranges.

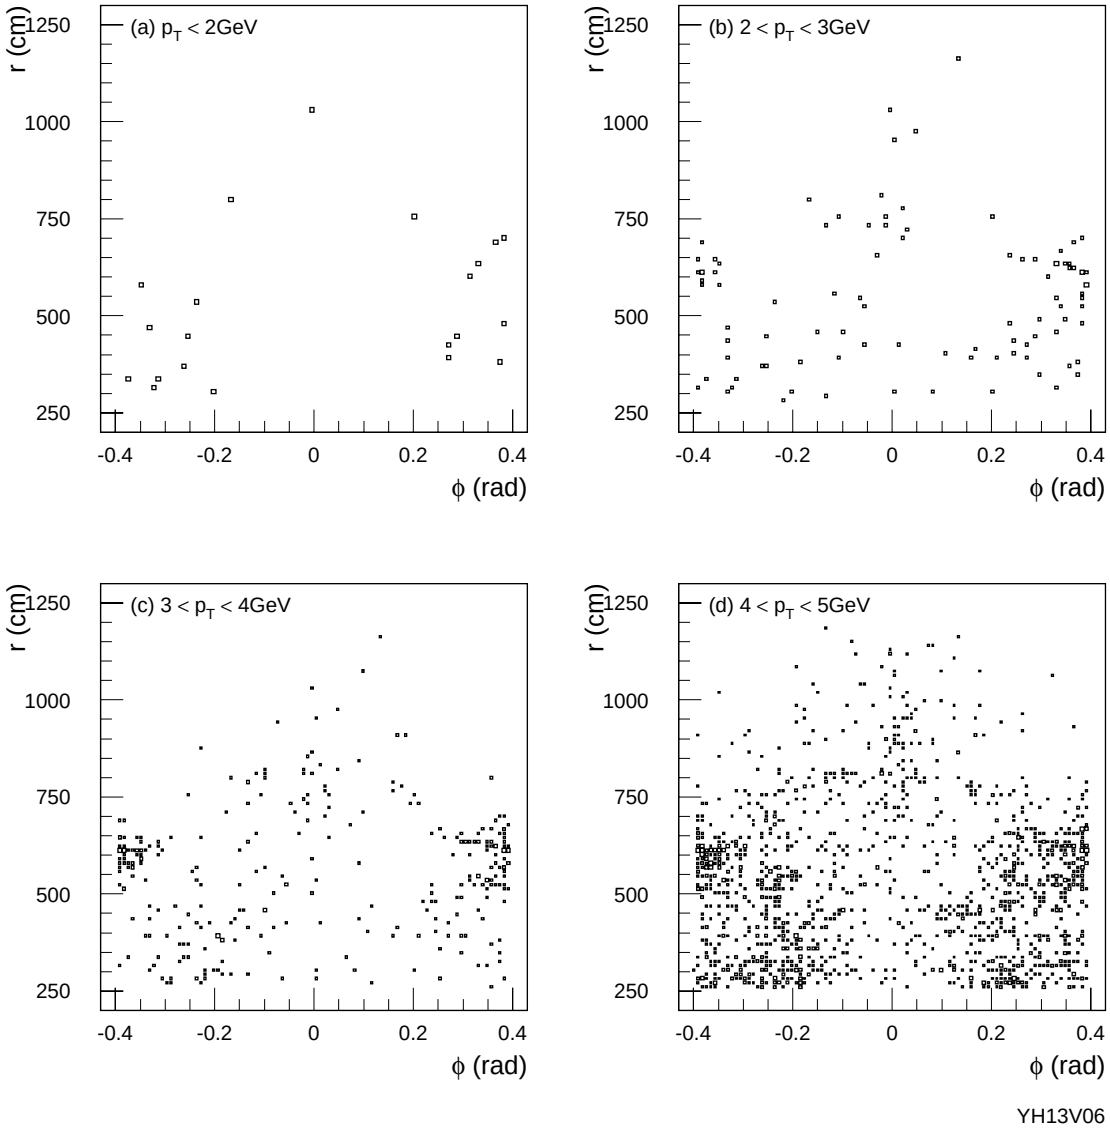

Figure 14-23 Source of efficiency from muons below threshold for the low-p_T end-cap system. Hit positions of muons below threshold triggered by the low-p_T system with p_T^{th} = 6 GeV are shown. The generated p_T was: a) < 2 GeV; b) 2–3 GeV; c) 3–4 GeV; d) 4–5 GeV.

One approach to reducing the rate from such low-p_T muons is to 'mask', or remove from input to the coincidence logic regions known to have poor momentum resolution. However it should be noted that a significant fraction of these triggers lie outside the poor resolution regions. Figure 14-24 shows the trigger efficiency for 2 GeV muons in the low-p_T system plotted versus the trigger acceptance for 10 GeV muons, when η–ϕ regions with poor momentum resolution are masked. It can be seen that the rejection of low-p_T muons is weakly increased when the masking procedure is used, at the price of a visible efficiency loss for higher-p_T muons.

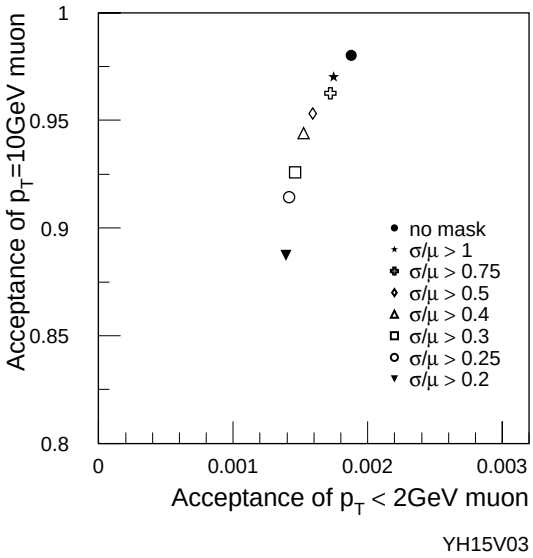

Figure 14-24 The effect of masking on the low-p_T trigger in the end-cap system. Trigger efficiency for muons with $p_T < 2$ GeV is shown versus that for 10 GeV p_T muons, for various masking schemes. The trigger threshold is 6 GeV.

The situation is similar for the high-p_T trigger threshold. The lower-p_T muons triggered with high threshold are more focused in the regions with poor momentum resolution than in the low-p_T case, as shown in Figure 14-25. However, randomly distributed triggered muons in the pivot plane are seen, particularly in the forward regions. An analysis of these triggers indicates that these are made by an accidental coincidence of the muon track and accompanying soft electrons.

Figure 14-26 shows the 5 GeV muon-trigger efficiency plotted against the trigger acceptance for high-p_T muons in the high-p_T end-cap system with 20 GeV threshold, when η–ϕ regions with poor momentum resolution are masked in the high-p_T system. In this case the masking of bad regions is moderately successful in reducing further the trigger rate from lower-p_T muons poorly measured by the system.

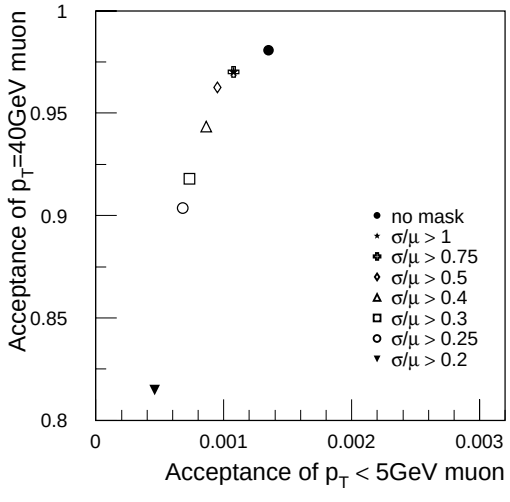

Figure 14-26 The effect of masking on the high-p_T trigger in the end-cap system. Trigger efficiency for muons with $p_T < 5$ GeV is shown versus the trigger efficiency for 40 GeV p_T muons for various masking schemes. The trigger threshold is 20 GeV.

14.5.4 Threshold flexibility

An important requirement of the muon-trigger system is that it possesses the flexibility to vary in small steps the trigger threshold across the range of p_T from 6 to 30 GeV; this allows the trigger rate to be reduced according to the background conditions and to produce RoIs for the Level-2 trigger with the lowest threshold possible. Whilst the 6 and 20 GeV thresholds are those for which the trigger system has been optimized it is important to demonstrate the flexibility of the trigger to set thresholds other than these. This performance is illustrated here by the threshold flexibility of the end-cap system.

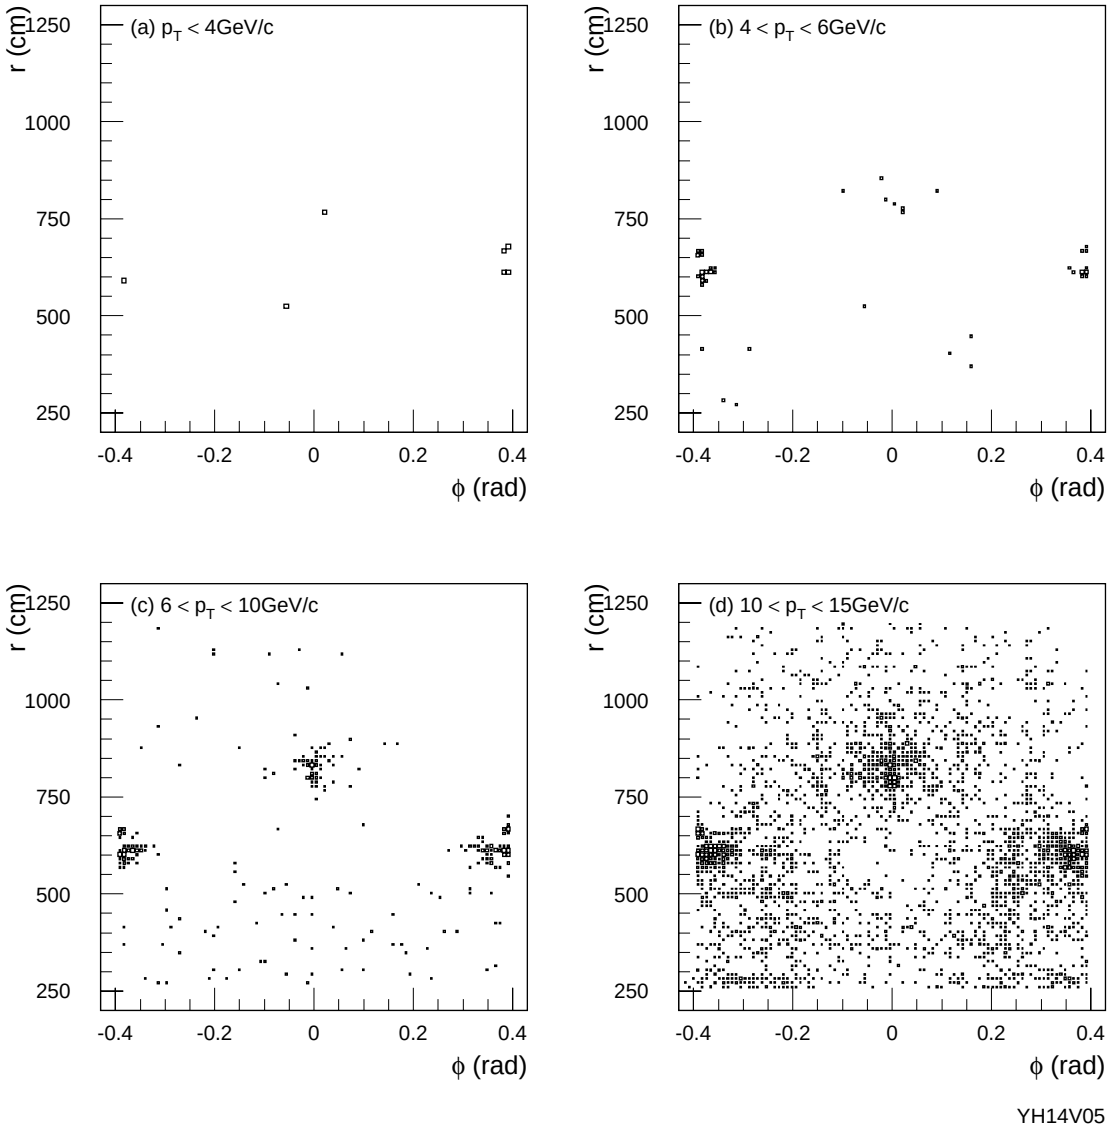

Figure 14-25 Source of efficiency below threshold for the high-p_T end-cap system. Hit position of muons triggered by the high-p_T system with p_T^{th} = 20 GeV is shown. The generated p_T was: a) < 4 GeV; b) 4–6 GeV; c) 6–10 GeV; d) 10–15 GeV.

There are two points to be investigated:

- the consequence, for the full system, of setting a threshold as high as 30 GeV;
- the ‘cross-over’ transverse-momentum threshold from the low-p_T system to the high-p_T system.

The physical quantities to be analysed are the coincidence window size and the discrimination power to reject muons with p_T below threshold.

Figures 14-27 and 14-28 show the range of coincidence-window sizes for the end-cap low and high-p_T systems, respectively, as a function of p_T threshold. The trigger-efficiency curves for each of these thresholds are shown in Figure 14-29.

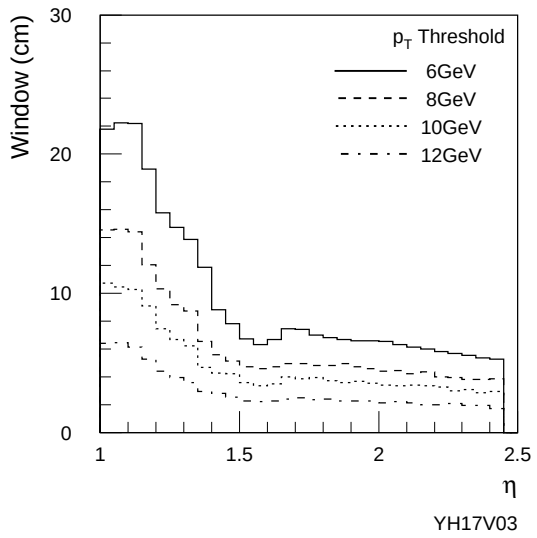

Figure 14-27 Coincidence window size in the main bending direction averaged over ϕ, as a function of p_T threshold, for the end-cap low-p_T system.

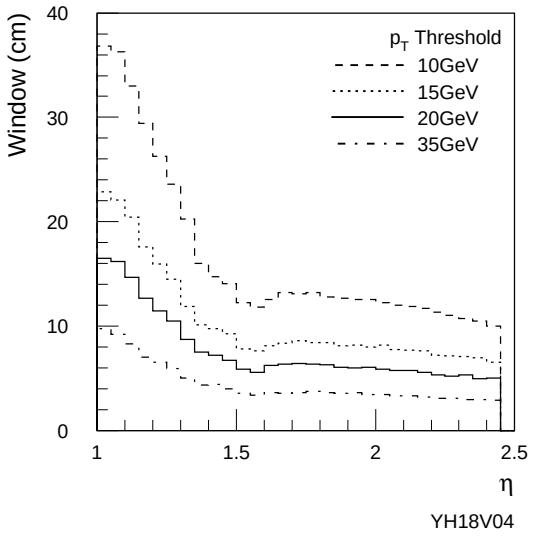

Figure 14-28 Coincidence window size in the main bending direction averaged over ϕ, as a function of p_T threshold, for the end-cap high-p_T system.

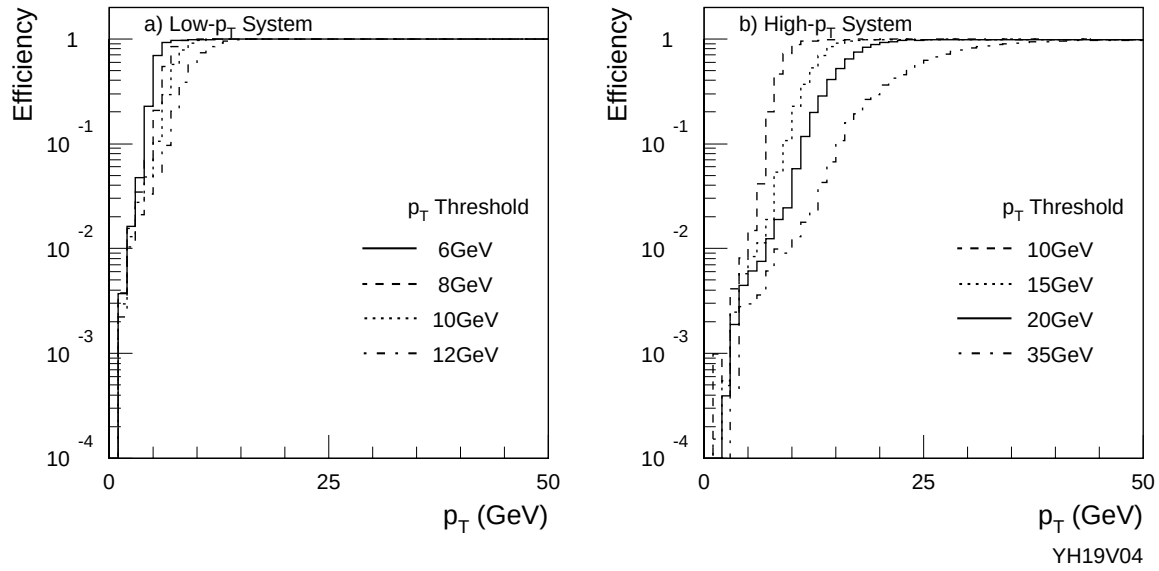

Figure 14-29 Trigger efficiency curves for a range of p_T thresholds for (a) low-p_T and (b) high-p_T end-cap trigger systems.

The consequences of this study are the following.

- The low-p_T system can be used to set a threshold as high as 10 GeV in the end-cap system. The main limitation arises from the trigger-detector granularity.
- The high-p_T system can be used to set a threshold at 10 GeV with efficiency close to 1. Thus we have a significant overlap in the threshold potential between the high and low-p_T systems in the end-cap. Being able to use the high-p_T logic at such low thresholds allows the possibility of an extremely robust trigger for b physics in the event that background conditions prevent the low-p_T logic from being usable.

- Thresholds significantly above 10 GeV/c can be set only using all planes of the trigger system, i.e. by using the high-p_T plane additionally. The consequence of this is a coincidence window size almost twice as large as that at 20 GeV.

The detector and trigger simulation has been used to optimize the granularity of the end-cap trigger counters and to determine the dimension of coincidence matrices. This has been done using the single muon data generated for the above studies and producing a read out segmentation scheme (in terms of a wire grouping in the r/η direction) that optimizes the performance requirements of the trigger (as set out in [14-1]) with the constraint on the total number of channels arising from cost limits. The final segmentation contains 220,000 wire channels and 95,000 strip channels. The inputs to the coincidence matrices from the wire groups consist of 31 δr inputs (in the range –15 to +15) and 15 from $\delta\phi$ inputs (–7 to +7). Details of the coincidence matrices and their inputs can be found in Chapter 12.

14.6 Trigger rates: prompt and secondary muons

The rates that will be seen in the ATLAS level-1 muon-trigger have been calculated by performing a full Monte Carlo simulation of the response of the ATLAS detector and level-1 muon-trigger to single muons to extract trigger efficiencies as a function of muon p_T and η. These efficiencies have then been convoluted with cross-sections extracted from Monte Carlo simulation of the physics processes that give rise to muons in the detector. This procedure has been performed for prompt single-muon production, from b and c hadrons and from W and Z decays, and for decays in flight to muons of π and K mesons.

The response of the detector and trigger ensemble to muons produced by the interaction of the beam and the LHC machine components (‘beam halo’), as well as to cosmic muons, has been investigated through dedicated Monte Carlo studies.

14.6.1 Acceptance to muons from interaction products

14.6.1.1 Cross-sections of prompt processes and meson decays in flight

Inclusive muon cross-sections at LHC arising from decays of b and c hadrons, of top quarks and of W and Z decays have been calculated using the Monte Carlo program Pythia, version 5.7 [14-11]. These cross-sections, integrated in the kinematic region $|\eta| < 2.7$ and $p_T > 3$ GeV, are shown as a function of p_T in Figure 14-30. The cross-sections are dominated by semileptonic decays of b and c hadrons for $p_T > 8$ GeV.

Light hadrons emerging from p–p collisions contribute charged-particle background via two mechanisms:

- decays in flight (hadron $\rightarrow \mu X$) in the inner tracker cavity,
- hadronic debris produced in calorimeter showers and in other interaction between hadrons and detector components.

Here we consider the consequence on the trigger rate of the secondary muons originating from the π/K decays. The cross-sections of such muons have been calculated using the DPMJET Monte Carlo program [14-12], and the results are shown in Figure 14-30. In the calculation the

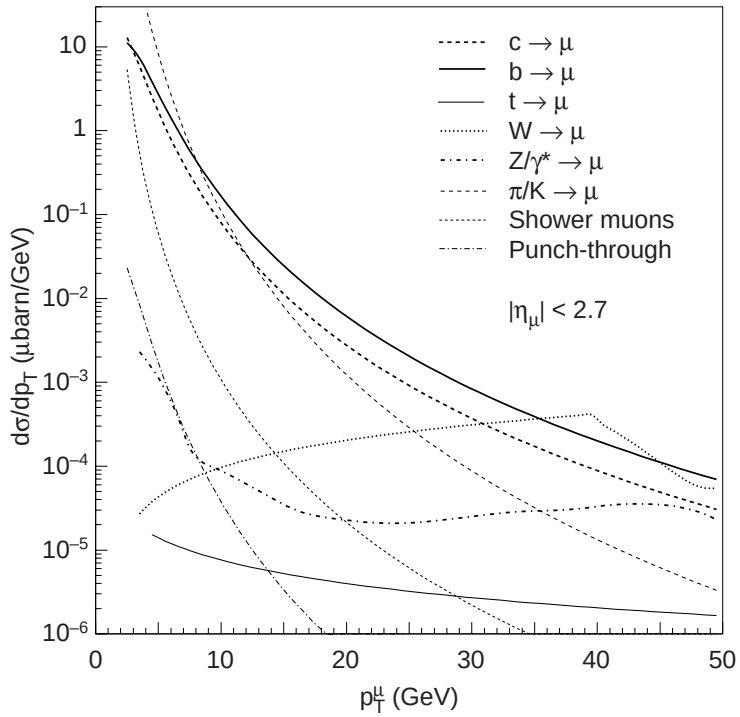

Figure 14-30 Muon cross-sections as a function of p_T at production, integrated in the region $|\eta| < 2.7$.

total inelastic pp cross-section is normalized to 80 mb. The Pythia 5.7 program has been used as a check of this cross-section. At lower transverse momenta ($p_T < 8$ GeV) in-flight decays of π and K mesons are the dominant source of muons and thus of trigger rate in the level-1 muon system.

14.6.1.2 Rates from prompt muons and meson decays in flight

The calculated trigger efficiencies for single muons have been convolved with the cross-sections plotted in Figure 14-30 to yield the total expected trigger rate from each production process in the barrel, end-cap and the sum of the two systems. In the end-cap this convolution has been performed in three bins of η to account for the significant dependence on η of both the cross-section and the efficiency of the level-1 trigger in the end-cap. In the barrel such a treatment is not necessary. The estimated rates are shown in Table 14-4.

The source of this rate varies in each of the two subsystems and for the process involved, although in almost all regions of the detector the rate arises largely from muons from in-flight decays of π and K mesons. This rate derives almost exclusively from lower-p_T (less than a few GeV) muons seen in the trigger detectors in the region $1.5 < |\eta| < 2$. Such muons are not seen in the barrel due to the shielding effect of the calorimeter material.

Uncertainties in trigger rates

Uncertainties in the trigger rate arise from several sources:

- The prompt rate is a small fraction of the total rate seen by the level-1 muon-trigger at low-p_T and less than half at high-p_T, and so its uncertainty does not dominate the errors on the total-rate estimate. There are however significant uncertainties in the prompt muon cross-sections estimated by Pythia at LHC energies. The most important of these

Table 14-4 Trigger rates, in kHz, expected in the barrel, end-cap and combined muon system arising from various physics processes. These rates are calculated by convolving the single muon cross-section from each process with the efficiency of the level-1 trigger to single muons. The low-p_T rates assume a luminosity of 1×10^{33} cm^{-2}s^{-1} and the high-p_T rates a luminosity of 1×10^{34} cm^{-2}s^{-1}.

	Process	Barrel	End-cap	Combined system
Low-p_T (6 GeV)	π / K decays	7.0	9.8	16.8
	b	1.9	2.1	4.0
	c	1.1	1.3	2.4
	W	0.004	0.005	0.009
	Total	**10.0**	**13.2**	**23.2**
High-p_T (20 GeV)	π / K decays	0.3	1.8	2.1
	b	0.4	0.7	1.1
	c	0.2	0.3	0.5
	W	0.035	0.041	0.076
	Total	**0.9**	**2.8**	**3.8**

originate from the unknown factorization and renormalization scales in the structure functions and the running coupling constant; these are likely to introduce cross-section uncertainties at the level of a factor two. Additional uncertainties arise from the unknown quark mass, yielding a factor of 1.5 on the total uncertainty. The accuracy with which cross-sections measured at the Tevatron are reproduced by the simulation used here gives an indication of the total size of such uncertainties. The model used here is consistent with inclusive b-production rates measured by the CDF experiment [14-13]. The extrapolation of cross-sections at Tevatron energies to those at the LHC introduces uncertainties less than a factor of two [14-14]. However all of these quantifications of uncertainties in the cross-section apply to the region $p_T > 6$GeV. The uncertainties in cross-sections at p_T values far below this (such as $p_T \sim 2$ GeV, where the majority of the level-1 muon-trigger rate originates) are likely to be greater and no authoritative estimates are available.

- Uncertainties in the modelled muon rate from in-flight decays of light mesons are dominated by model dependence and the parameters governing the parton showering and hadronization. An indication of the uncertainties in the cross-section estimation has been obtained by comparing two independent Monte Carlo programs. All rate calculations have been performed using DPMJET, giving the rates discussed above, and with Pythia as a cross-check. The two programs have been optimized independently and so can be expected to reflect the divergence of the different approaches to modelling of the cross-section. Whilst Pythia has the possibility to describe soft processes, DPMJET is a dedicated simulation of inelastic hadron collisions tuned to existing data. (Both programs here are normalized to a total inelastic cross-section of 80 mb.) Whilst the differential cross-sections $d\sigma/dp_T$ and $d\sigma/d\eta$ modelled by the two programs are of rather similar shape, Pythia consistently produces a somewhat lower cross-section than DPMJET. For this reason the rate estimate obtained using the Pythia cross-section is typically 35% lower than that from DPMJET. This discrepancy is considered to be an indication of the systematic uncertainties on this cross-section, particularly with regard to the modelling of

the p_T distribution. Uncertainties on the total inelastic cross-section are likely to be smaller than 30%.

- The limited statistics of the efficiency calculation result in significant uncertainties on the total accepted rate. This is particularly true of the π/K decays where the accepted rate is a result of the convolution of a very low efficiency with a very high cross-section. The statistical uncertainties on the accepted rate are of the order 10–15%, and are thus generally small in comparison with the expected model uncertainties on the input physics processes.
- In principle the rate can depend upon the accuracy of the modelling of the detector response and the ongoing optimization of the detectors themselves. (In the end-cap for example the rate studies were performed using a read out segmentation that is slightly different from that which the current optimization yields.) In general however the effect the detector response will have on the final rate is small not least because of the declustering algorithms which ensure that at various stages in the trigger chain the multiplicity of hits is highly controlled. Whilst the details of the detector response and the interaction with the declustering algorithms do introduce uncertainties on the above rates they are expected to be small. During studies of the efficiency of the level-1 trigger the effect of substantial changes in the amount of passive material in the ATLAS detector has been investigated; these changes have been found to significantly affect rates, since these are dominated by rather soft muons. The material modelled by the ATLAS detector simulation is a rather accurate representation of the currently envisaged design and future developments are likely to involve the addition of material (such as cables, cooling structures, etc) rather than its removal, and these in turn are likely to suppress rather than enhance the total trigger rate from prompt and secondary muons.

  Similarly the optimization of the read out segmentation of the trigger detectors, particularly in the case of the wire grouping of the TGC chambers in the end-cap, will affect the sharpness of the threshold that can be set in the trigger. However the optimal segmentation, given the cost constraints on the number of channels in the total system, is rather close to that in the trigger simulation used here, and it is expected that this uncertainty will change total rates rather marginally.

  Masking of low-resolution regions in the detector, again most relevant in the end-cap (see Figure 14-10), is foreseen. This will serve to reduce rates at the expense of acceptance. Such possibilities are currently under study, and it is expected that the rate in the end-cap can be reduced by a few tens of per cent. The analysis of real ATLAS data will allow an optimal masking scheme to be implemented.

14.6.2 Trigger rate from non-proton-proton muons

14.6.2.1 Cosmic muons (barrel system)

Despite the significant depth at which the ATLAS experiment is situated, cosmic rays contribute to the trigger rate in the muon system. This rate arises largely from the access shafts located above the experimental hall. The ATLAS cavern is located about 75 m underground and access made possible by two parallel shafts (9 m and 12.6 m in diameter) about 60 m deep. The axes of both shafts are centred on the beam line and they are separated by a distance of 25 m.

Cosmic muons [14-15] have been simulated at ground level above the experimental area and propagated through the shafts and the rock using Monte Carlo techniques. Muons reaching the experimental hall have then been tracked through the ATLAS detector and trigger simulation. The first 14 m of material below ground level have a density of $1.6\,g/cm^3$; the remaining material down to the experimental hall has been simulated as rock with $2.5\,g/cm^3$ density. To reach the ATLAS detector, a perpendicular muon traversing the full depth of rock must have an energy above 20 GeV.

The surface used for the generation of cosmic muons was a square at ground level centred on the detector position with a side of length that has been varied from 80 m to 150 m. Hits recorded by the RPC system have been passed through the trigger simulation; the ratio of the number of triggered muons to the number generated at the surface gives the trigger probability per incident muon for the given area of cosmic production. The trigger rate from cosmic muons is given by the rate of muons incident on the generation area multiplied by the calculated trigger probability. The trigger rate has been studied as a function of the area dimension in order to evaluate the edge effects due to the limited size of the production surface.

By normalizing the incident cosmic rate to $100\,Hz/m^2$ (the approximate rate of the muon component at sea level), we found a trigger rate in the low-p_T system (for a 6 GeV p_T threshold) below 200 Hz. The corresponding rate for the high-p_T system (for a 20 GeV threshold) is much lower (below 10 Hz). This rate is more than one order of magnitude lower than the rate from pp interaction muons, but still sufficient to offer the potential for their use in the calibration of the muon-trigger in the barrel region.

Figure 14-31 shows the η–ϕ scatter of cosmic muons that give a trigger in the barrel low-p_T system. Most of the triggered muons are in the central region of the detector, in the top and the bottom of the muon-trigger system, as might be expected. It is expected that cosmic muons, given the estimated rate calculated here, can be used to calibrate the muon-trigger in the barrel region, potentially for both geometrical alignment and timing studies. This valuable data can be exploited far in advance of the first beams in LHC.

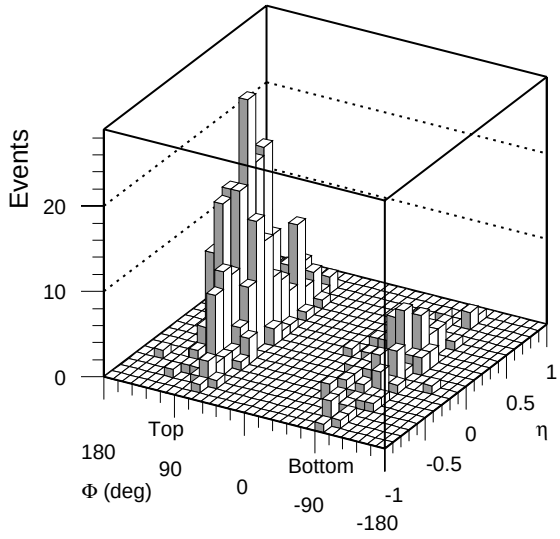

Figure 14-31 Cosmic muons in the barrel system. η–ϕ coordinates of cosmic muons triggering the low-p_T trigger system (p_T threshold of 6 GeV).

14.6.2.2 Beam halo muons (end-cap system)

It has long been known at hadron colliders that very substantial muon fluxes can be seen, particularly in the more forward regions, arising from muons produced through the interaction of the proton beam with residual beam gas in the beam pipe, and with the limiting apertures of the machine such as collimators. The muons are produced largely at a small angle relative to the beam direction and at low radius, and can have potentially catastrophic implications for the rate of muon-triggers seen in these regions. In this respect the experience of previous detectors has been valuable, and the projective geometry of the end-cap muon-trigger chambers serves to control this problem.

A comprehensive study of muons produced in interactions between the LHC beam and the machine components has been performed for the CMS experiment using a detailed simulation of such processes [14-16], considering the latest information on the design of the LHC machine. The differences between the beam conditions in ATLAS and CMS are slight enough that this simulation is also entirely relevant for ATLAS. The interactions modelled are those of a beam of 530 mA at 7 TeV with luminosity 1×10^{34} cm^{-2}s^{-1}, with all machine components within 1000 m of the interaction point simulated. The muons produced in such interactions have an energy distribution very sharply peaked at energies of a few GeV and with a tail extending to 500 GeV. The greatest flux of muons at the cavern entrance is below a radius of ~1m relative to the beam line, well below the geometrical acceptance of the end-cap muon-trigger chamber, but extending to the limit of the wheel of TGC chambers in the pivot plane. The mean angle of these muons to the beam-line is 60 mrad. The flux of muons from the simulation at the entrance to the ATLAS cavern ($|z| \sim 26.5$ m) and in the pivot plane of the end-cap trigger counters ($|z| \sim 14$ m) is shown in Figure 14-32.The counting rate per detector plane in each end-cap from such muons is ~60 kHz.

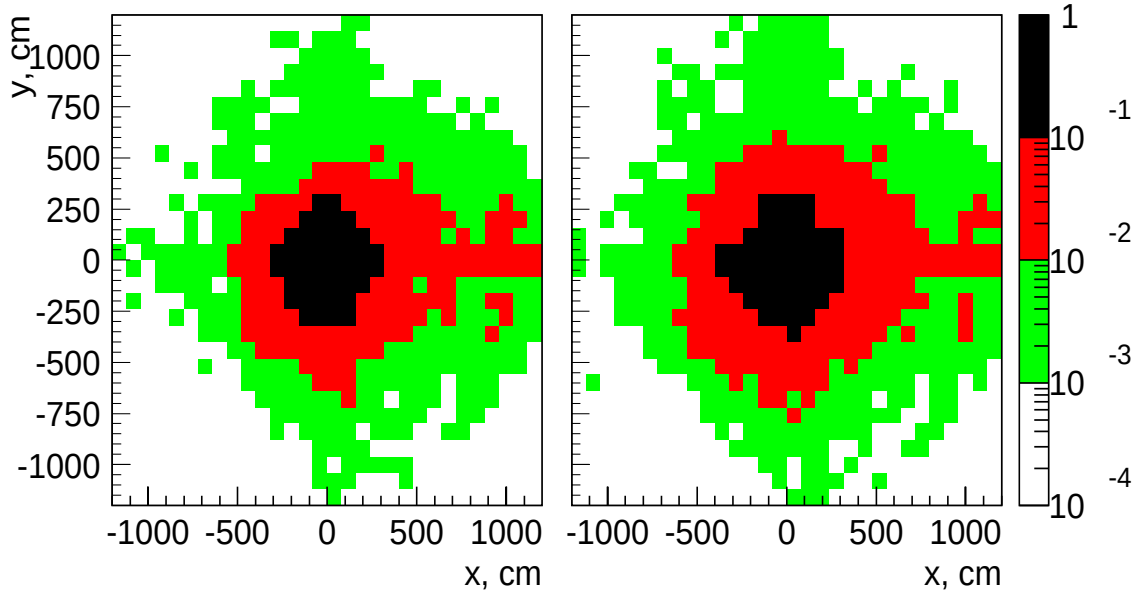

Figure 14-32 The flux of muons at (a) $|z| \sim 26.5$ m, the entrance to the experimental hall, and (b) $|z| \sim 14$ m, the position of the pivot plane of the TGC counters, as modelled by [14-15]. Units are muons cm^{-2}s^{-1}.

These simulated muons from beam-loss processes have been passed through the full ATLAS detector and trigger simulation to estimate the rate in the end-cap level-1 muon-trigger.

The estimated rates arising from the halo muon flux are 250 Hz and 16 Hz for the 6 GeV and 20 GeV low- and high-p_T triggers respectively. These rates are negligible in comparison with the rates from interaction products, and can contribute significantly to the trigger rate only if the halo rate here is underestimated by a factor of ten or more; in this instance the rate would still be quite tolerable.

14.6.3 Dimuon-trigger rate

14.6.3.1 Two-muon physics reactions

The rate of dimuons from the process $bb \rightarrow \mu\mu X$ has been extracted using the Pythia Monte Carlo. This process dominates all other sources of dimuon production where both muons have $p_T < 20$ GeV. For $p_T \sim 20$ GeV Z decays have a similar dimuon production cross-section and above this point Z products dominate. For the purposes of providing a b-physics trigger, demanding two muons that both pass the nominal low-p_T threshold (6 GeV) yields genuine dimuon events almost entirely from b-decays. The integrated cross-section for the process $bb \rightarrow \mu\mu$ in the region $|\eta| < 3$, where both muons have $p_T > 6$ GeV, is 0.03 μb, in reasonable agreement with previous calculations made using ISAJET [14-17]. The total rate from genuine high-p_T dimuon b-events is thus very small compared with the single-muon rate for the same p_T.

14.6.3.2 Trigger rate from double-counted single muons

In the end-cap

When a muon traverses the overlapping edges of neighbouring trigger sectors, it can produce a trigger in each of the trigger sectors. In the end-cap trigger system, efforts to both minimize such double triggers and avoid inefficiency in the overlap regions have been made as described in Section 14.4. Despite this, small overlap regions remain. Figure 14-33 shows, as an example, the position in station 3 of 6GeV single muons which are doubly triggered by the end-cap low-p_T trigger system.

The vertical bands in the figure are due to the overlap of neighbouring trigger sectors by half a strip width. The horizontal band located around a radius of 437 cm is also due to the overlap of the trigger sectors in which the patch-panel logic can not be applied owing to the different chamber layout between the end-cap and forward regions. The current design has at most 3.5cm of overlap between neighbouring sectors, which has yet to be optimized by using specially modified strips in this region; such an optimization will reduce double triggers.

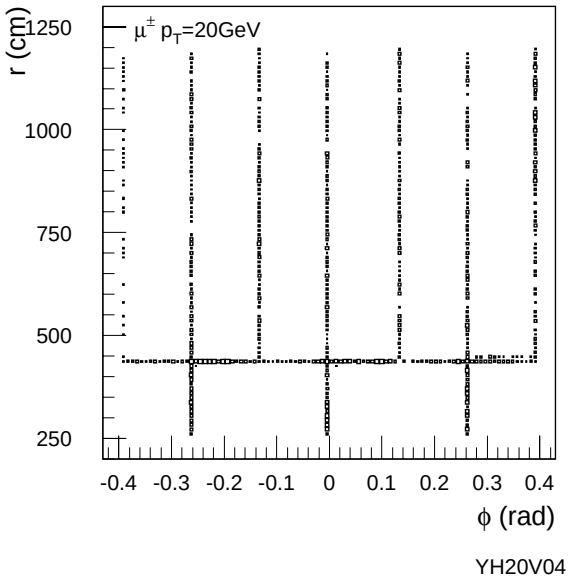

Figure 14-33 Hit positions of doubly triggered muons in the end-cap pivot station for 20GeV p_T muons. The double triggers are due to overlaps between neighbouring trigger sectors.

The fraction of the double-triggered muons to muons generated in an $|\eta|$ region of 1.2–2.35 is summarized in Table 14-5. These fractions imply a rate at nominal LHC luminosity ~4 kHz. Efforts are being made to find strategies to reduce the double-counting rate further.

Table 14-5 The fraction of doubly triggered muons in the low-p_T system, as a fraction of muons generated in the region $1.2 < |\eta| < 2.35$ in the end-cap muon-trigger.

	6GeV	20GeV
μ^+	1.6±0.1%	2.2±0.1%
μ^-	1.9±0.1%	2.4±0.1%

The fraction of the geometrical overlap to the total area of the end-cap trigger counters is estimated to be about 2.5%. The difference between the estimate and the results of the simulation is due to the incident angle of muons to the counters. The double triggers in the transition region between the barrel and end-cap are not shown in Figure 14-33, but are solved by logic in the muon central trigger processor interface, MUCTPI (see Chapter 13).

In the barrel

In the barrel the major source of double-counted single muons is the overlap in the r–ϕ projection of the large and small sectors. This overlap is necessary to provide a level-1 muon-trigger for tracks crossing both large and small chambers, which are used for the calibration of the alignment system of the muon spectrometer. During normal operation of the experiment an appropriate masking of the ϕ strips located in the overlap regions will be applied. The optimization study of the single-muon-trigger acceptance against the rejection of double-counted single muons is in progress.

14.7 Fake trigger rate

A large background flux is expected at LHC due to the interaction of hadrons arising from the proton–proton interactions with the forward elements of the ATLAS detector, the shielding system and machine elements such as the beam pipe and collimators to produce neutrons in the experimental hall. Thermal neutrons thus produced are absorbed by nuclei and emit photons of an energy 10–1000 keV. Hadrons and muons are also produced in cascade processes. The particles thus produced can induce high counting rates in the muon-trigger chambers. Here we estimate the rate induced in the ATLAS level-1 trigger from such backgrounds.

The background flux and rate seen in the level-1 trigger have been evaluated using the FLUKA Monte Carlo program [14-18]. This Monte Carlo program is not that used for the standard simulation of high-energy physics processes; it can track particles at very low kinetic energies. In particular, neutrons are propagated down to thermal energies. The shielding layout of the ATLAS detector used here is version ‘TP32‘, as discussed in [14-19]. The shielding layout envisaged by the TP32 scheme has changed significantly in the current design; however the study and optimization of the new shielding scheme compatible with the most recent changes to the muon system layout has shown neutron and photon fluxes very close to those seen with TP32. The statistics accumulated with charged particles is, however, still insufficient. For this reason we have used the higher statistics results available from TP32 for the study of the accidental trigger rate.

Background can be classified into three classes:

- Background particles that produce a hit in a single trigger counter.

  Background originating from soft Compton electrons ($E < 2$ MeV) and neutron-induced soft protons is produced in the chamber wall or the gas of the detector and absorbed by

the material of the same detection unit. This does not introduce correlations between hits detected in different muon-trigger counters. The counting rate from this background at nominal LHC luminosity is 5.2 Hz/cm^2 in the barrel and 4.5 Hz/cm^2 in the end-cap. Figure 14-34 shows typical fluxes in the second station of the precision muon system.

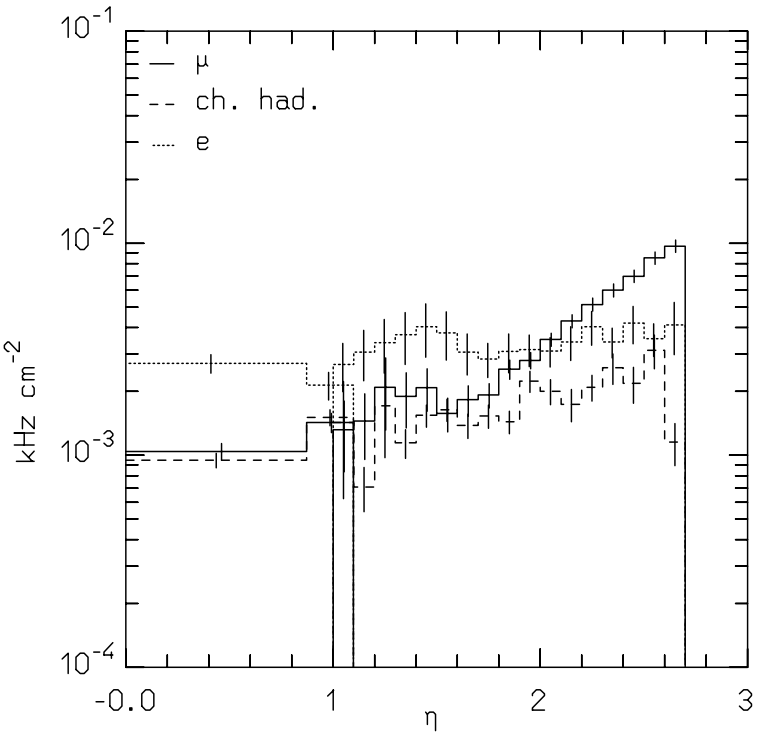

Figure 14-34 Charged-particle background fluxes, for muons, charged hadrons and electrons, as a function of rapidity at the second station of the MDT chambers, close to the muon-trigger counters. This data is from the FLUKA simulation.

- Background particles that can produce hits in two adjacent trigger counters.

  Hard Compton electrons with energy above 2 MeV and soft protons: These particles are absorbed in ~3g/cm^2 of the typical material of the trigger counters. This background produces correlated hit patterns in trigger counters belonging to the same trigger station. Counting rates from this background at nominal LHC luminosity are 2.0 Hz/cm^2 in the barrel and 3.1 Hz/cm^2 in the end-cap.

- Background particles that can produce correlated hits in more than one muon-trigger station.

  Hard protons, pions and muons having momentum ~100 MeV, sufficient to cross not only neighbouring trigger planes within a station, but traverse closely spaced planes of a trigger station, such as the two planes of the low-p_T system. Such a particle can produce a correlated hit pattern simulating a true muon-trigger. Detector simulation by FLUKA predicts a substantial rate for these relatively energetic particles: ~1.9 Hz/cm^2 and 3.0 Hz/cm^2 in the barrel and end-cap, respectively, where approximately 60% of the rate comes from muons.

The following section will present a study of the first two types of background, i.e. those that produce hits in at most one trigger station. Section 14.7.2 will discuss fake trigger rate produced by penetrating particles.

14.7.1 Trigger rate from background hits in a single trigger station

A simplified version of the trigger-counter geometry and the trigger logic has been used to investigate the rate arising from soft ($E < 2\,\mathrm{MeV}$) background particles. Each of the trigger subsystems (barrel and end-cap) was divided into projective towers. The tower dimension is given by the counter size in the pivot plane. The size of a tower in the second and third trigger planes was determined by projecting from the nominal interaction point to the pivot plane counter. The calculation of the rate proceeds as described in [14-20].

In the rate calculation the coincidence of correlated hits from a single track with other background hits has been considered, accounting for the appropriate majority logic. It is seen that the rate due to this is much larger than that due to accidental coincidence of uncorrelated hits. However, only correlations in one trigger station are included and not the rate due to tracks that leave correlated hits in more than a single trigger station (see Section 14.7.2).

The rate calculation in the barrel system has been performed considering six counter planes (three doublets). It assumes two different scenarios for the majority logic in the outermost doublet: both one-out-of-two and two-out-of-two. Whilst the two-out-of-two coincidence provides a more robust logic and thus better rate performance, the one-out-of-two logic has been considered for the study because of its greater efficiency.

In the end-cap system seven planes of trigger counters are considered with the innermost station consisting of a triplet unit. This innermost station is used in the high-p_T trigger where the majority logic used is two-out-of-three in the r-plane, and one-out-of-two in ϕ. The low-p_T trigger logic is the same as in the barrel and the calculation proceeds in the same way.

In the calculation of the rate expected in the high-p_T systems (for both barrel and end-cap) the coincidences of prompt muons with background hits are also taken into account. This effect introduces a linear dependence of the fake trigger rate on the background flux and a quadratic dependence on the luminosity.

Figure 14-35 compares the fake trigger rate (for both the nominal background flux and the flux multiplied by a factor of five and ten, for the low luminosity, and ten, for the high luminosity) with the rate expected from prompt muons and in-flight decays of π and K mesons as a function of luminosity. The result is presented for the low- and high-$p_T$triggers (at nominal thresholds) for the barrel and the end-cap.

The dominant contribution to the fake low-p_T trigger rate in both barrel and end-cap is due to the coincidence of a pair of hits from a penetrating particle in one of the low-p_T stations, with one or more hits deposited by any other particle. The fake high-p_T trigger rate is dominated by a low-p_T trigger in coincidence with any other hit (or track) in the high-p_T station of the barrel or end-cap.

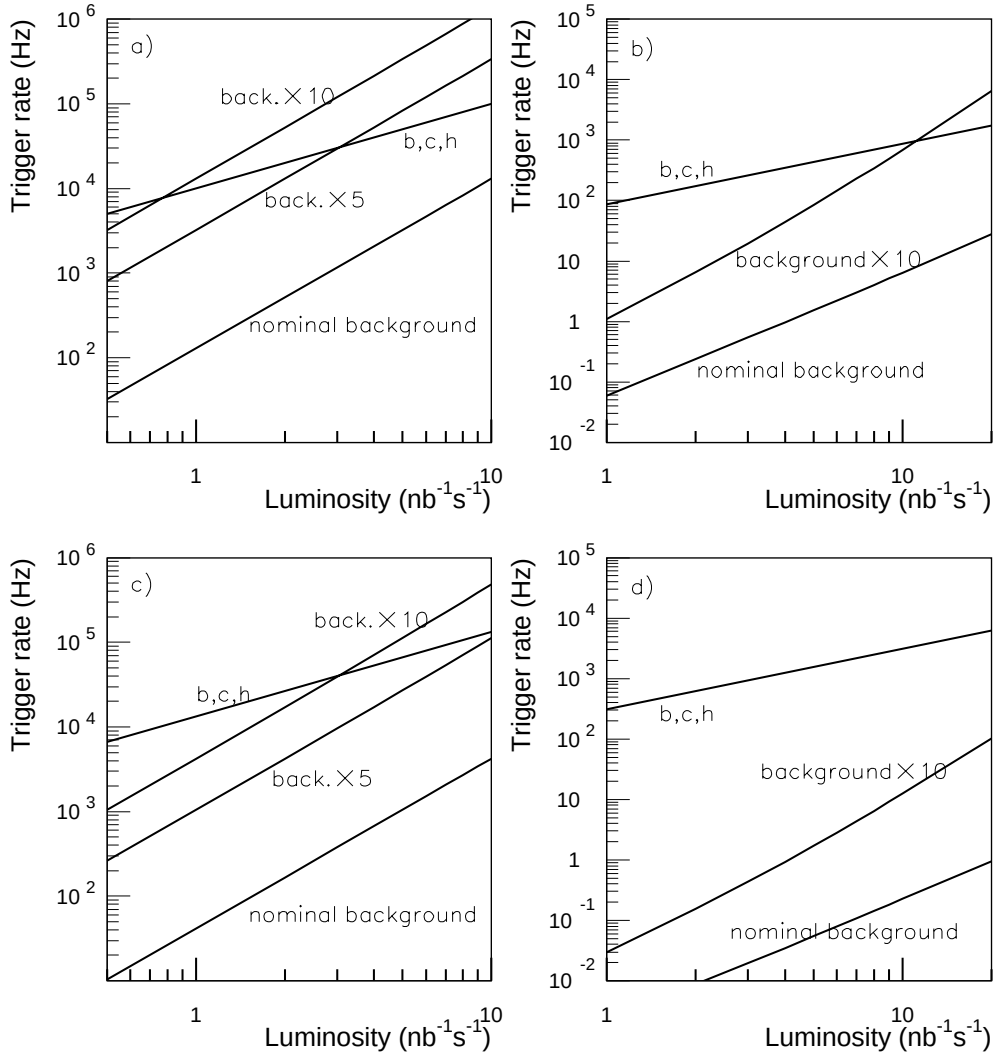

Figure 14-35 The accidental trigger rate from soft background particles compared with the rate expected from prompt muons (b, c, h) as a function of luminosity, for a) barrel trigger, low-p_T 6 GeV threshold; b) barrel trigger, high-p_T 20 GeV threshold; c) end-cap trigger, low-p_T 6 GeV threshold; d) end-cap trigger, high-p_T 20 GeV threshold.

14.7.2 Trigger rate from hard background particles

Monte Carlo studies of hard (E > 10 MeV) charged particles in the ATLAS cavern predict that whilst the rate of electrons with energy large enough to cross two muon stations is very small, there is a substantial flux of muons, charged pions, protons and charged kaons. The expected flux at nominal LHC luminosity of muons and pions is 1.1 Hz/cm^2 and 0.5 Hz/cm^2, respectively, in the barrel and 2.1 Hz/cm^2 and 0.6 Hz/cm^2 in the end-cap. This flux leads to a significant background rate in the level-1 muon-trigger system. The momentum of these particles is in the range 50 to 200 MeV, Figure 14-36. Such background particles can penetrate trigger-detector components and produce a trigger coincidence consistent with that from a

prompt muon, in particular in the low-p_T system, where only a two-plane coincidence is required.

The processes in the Monte Carlo program that yield this particle flux as well as their angular distribution are not yet known and a dedicated analysis is ongoing. In the absence of detailed information about these particles a few possible scenarios as to their origin and angular distribution were assumed in a simulation study that has been performed to assess possible trigger rates due to this flux.

Preliminary results of this simulation study indicate that the level-1 trigger rate from this background is potentially problematic. This study is preliminary and it will be repeated when more information is available about the origin and nature of this background. In parallel to this study, the muon-trigger group is contemplating scenarios to make the trigger more robust against such backgrounds.

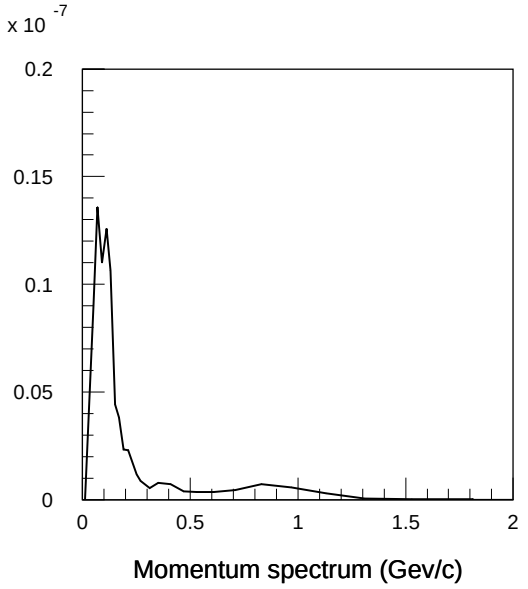

Figure 14-36 The momentum spectrum of the background responsible for most of the rate induced.

Several solutions are being considered:

- For the single low-p_T muon-trigger (at low luminosity), add a loose coincidence with the plane used in the high-p_T trigger to confirm the origin of the muon track from the vertex. This can be implemented in the “non-bending” projection for the barrel and in both projections for the end-cap with no change in the electronics design of the muon-trigger. The consequence (to be verified) is a possible increase of the lowest p_T threshold that can be set. Initial studies of this solution have given encouraging results.
- Demand an additional coincidence of the muon in the level-1 muon system with associated calorimeter cells. The tile calorimeter [14-21] has demonstrated good separation of the muon signal from the electronic and physics noise of the detector, and this could be used in conjunction with the level-1 muon system in the barrel. In particular, the third sampling appears very attractive for this purpose, although the modifications to the design of the front-end electronics that would be required to make available the additional signals have not yet been studied in detail. However, the tile calorimeter only extends to $|\eta| \sim 1.5$.
- Utilizing the end-cap calorimeter for muon confirmation in the level-1 trigger, although not impossible, appears quite difficult and would imply a considerable redesign of the end-cap calorimeter electronics.
- For the level-1 trigger in the end-cap muon system, another possibility is to demand the additional coincidence of the innermost TGC station, situated between the forward calorimeter and the end-cap toroid. This inner station is not nominally used in the trigger, being foreseen only to measure the second coordinate of muon tracks with $|\eta| > 1$. This solution implies a significant extension to the present design of the forward trigger electronics, and can be considered only after a detailed and complete study of the hard charged-particle background. The study must consider the trigger rate for the present

forward trigger design and for that including the innermost TGC station, to evaluate the overall benefit of this approach.

However, by far the most important step to be performed by the ATLAS community is a very detailed analysis of the physics processes which lead to the production of such high-momentum background particles, in conjunction with the identification of the regions where these processes mostly occur. Further optimization of the shielding system of the ATLAS muon system must take place to identify solutions that reduce the hard charged-particle rate. Finally, experimental checks of the background rate in the LHC-like environmental conditions can help to quantify the reliability of the radiation-level estimates in the ATLAS hall.

14.8 References

14-1 *LVL1 Muon Trigger User Requirements Document (Draft version 1.4)*, ATLAS working document, ATL–DA–ES–0002, March 1998.

14-2 *ATLAS Trigger Performance Status Report*, CERN / LHCC / 98–15, June 1998.

14-3 *GEANT 3 — Detector description and simulation tool*, CERN program library long write-up W5013, 1993.

14-4 *DICE manual*, ATLAS internal note SOFT–NO–10, 1994.

14-5 *ATRIG — ATLAS Trigger Simulation User Guide Revision 0.00,* ATALS note SOFT–94–017, 1994.

14-6 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97-22 ATLAS TDR 10, May 1997.

14-7 L. Chevalier et al., *AMDB_SIMREC: A Structured data base for the ATLAS Spectrometer Simulation Program*, ATLAS note MUON–NO–148, 1997.

14-8 O. Jinnouchi and O. Sasaki, *Study of the muon level-1 trigger scheme for TGC,* ATLAS note DAQ–NO–82, 1998.

14-9 A. Nisati, *Preliminary timing studies of the barrel muon-trigger system*, ATLAS note DAQ–NO–83, 1998.

14-10 A.Rimoldi, *Geometrical acceptance for the trigger chamber system and study of the reaction* $H \rightarrow ZZ^* \rightarrow$ *four muons from the acceptance point of view,* ATLAS note DAQ–98–119.

14-11 T. Sjöstrand, *PYTHIA 5.7 and JETSET 7.4 Physics and Manual*, CERN–TH.7112 / 93, 1993.

14-12 J. Ranft, *DPMJET version II.3 and II.4: sampling of hadron-hadron, hadron-nucleus and nucleus-nucleus interactions at cosmic ray energies, according to the Dual Parton Model description of the model and code manual*, INFN-AE-97-45.

14-13 P. Eerola, *The inclusive muon cross-section in ATLAS*, ATLAS note PHYS–NO–120, November 1997.

14-14 M. Mangano, *Heavy quark production*, LHC98 workshop, CERN, February 1998.

14-15 Particle Data Group, *Review of particle properties*, Phys. Rev, D54, 122, 1996 and references therein.

14-16 Drozhdin, et al., *Accelerator related background in the CMS detector at LHC*, Nucl. Instrum. Methods, A381, (1996) 531-44.

14-17 A. Nisati, *Muon rates at the LHC*, Proc. Large Hadron Collider workshop, Vol. II, eds. G. Jarlskog and D. Rein, CERN 90-10, ECFA 90-133, December 1990.

14-18 A. Fasso et al., *FLUKA: present status and future developments*, Proc. 4th International Conference on Calorimetry in High-energy Physics La Biodola, Elba, Italy; September 1993, edited by A. Menzione and A. Scribano.

14-19 *The ATLAS Technical Coordination Technical Design Report* (in preparation).

14-20 A. Nisati, *Fake level-1 muon-trigger rates in the ATLAS experiment*, ATLAS note DAQ–98–123, July 1998.

14-21 *The ATLAS Tile Calorimeter Technical Design Report*, CERN / LHCC / 96-42, December 1996.

15 Level-1 central trigger processor

15.1 Requirements analysis and functional specification

The role of the Central Trigger Processor (CTP) in the LVL1 trigger is to combine information from the calorimeter and muon trigger processors, as well as from other sources such as the special calibration triggers, and to make the final LVL1 accept decision (L1A). This decision is distributed via the Timing, Trigger and Control (TTC) system to the front-end electronics of the detector systems, where it initiates the readout. The CTP also sends summary information regarding the trigger type to the front-end systems, via the TTC.

The input data to the CTP from the calorimeter trigger and the muon trigger processors consist of the following: for electrons / photons, taus / hadrons, jets and muons the CTP receives multiplicity values for a variety of p_T thresholds (and isolation criteria in the case of electrons / photons and hadrons / taus); for missing-E_T and total-E_T the CTP receives flags (possibly encoded) indicating the result of comparing the missing-E_T and total-E_T values with a variety of thresholds.

Provision is also made for trigger inputs from sources other than the calorimeter and muon processors. These inputs might be used, for example, for luminosity triggers or for other special triggers. External inputs are also needed for certain calibration and test triggers for detector systems. However, for stand-alone tests and calibration runs the detector systems can introduce their own triggers via the TTC system (see Chapter 16).

The data arriving at the CTP from different sources will not generally be aligned in time (i.e. data from the same bunch crossing but from different sources will arrive at the CTP at different times). The CTP therefore has to include programmable elements to delay the signals that arrive first to be in time with those that arrive last.

The algorithm used by the CTP to combine the different trigger inputs allows events to be selected on the basis of menus. An event is selected if it satisfies the criteria of one or more menu items. Each menu item consists of the logical combination of a number of criteria, typically multiplicity requirements for electrons / photons, hadrons / taus and jets, and threshold requirements on missing or total transverse energy. The CTP makes provision for prescaling certain kinds of triggers; there is an individually-programmable prescale factor for each menu item. In total, the CTP allows 128 bits of input data to be combined in up to 96 menu items. Menu items may be enabled / disabled for certain bunches in the LHC, for example during the gaps in the bunch train when calibration pulsers may be fired.

The CTP is responsible for introducing deadtime as required by the front-end systems [15-1]. Deadtime is normally generated internally using algorithms that are described in detail in Section 15.2.5. In summary, there is a dead period of four bunch crossings following each LVL1 trigger (required by some front-end systems to avoid overlapping time frames for readout) and a more complicated algorithm that introduces deadtime when it is predicted that buffers in the front-end system are nearly full. The more complicated algorithm allows low-priority and high-priority triggers to be defined; for low-priority triggers, deadtime is introduced well before the front-end systems become full. Whilst high-priority triggers are vetoed if necessary, the deadtime fraction is smaller than for low-priority ones. Deadtime can also be caused by an external VETO signal, but this is used only as a last resort.

In addition to providing the L1A signal and trigger-type information to the TTC system, the CTP provides information to the DAQ and LVL2 trigger systems. The data sent to the DAQ are recorded on permanent storage for events that are retained by the LVL2 trigger and event-filter systems. These data can be used for offline monitoring of the trigger system (and of experimental conditions such as beam quality) and also for the calculation of trigger efficiencies. The data sent to the DAQ may also be used for online monitoring of the LVL1 trigger. As far as the DAQ is concerned, the CTP acts as a readout driver (ROD); in this respect, it must satisfy the rules given in Ref. [15-1]. A time frame is read out around each triggered bunch crossing, providing information on the status of the trigger in the bunch crossings just before and just after the one that triggered.

The data that are sent to the LVL2 trigger, via a separate path from the one to the DAQ, are used to guide the LVL2 processing. These data provide a detailed description of why the CTP selected the event. They are used, together with information sent to the LVL2 system by the LVL1 calorimeter and muon trigger processors, to distinguish between 'primary' RoIs that may have contributed to the selection of the event at LVL1, and 'secondary' RoIs that are flagged by LVL1 but which played no part in the selection of the event. Information about the requirements placed on the CTP by the LVL2 trigger can be found in Ref. [15-2]. For example, the information must be made available to the LVL2 trigger system within a specified latency.

The CTP contains several sets of scalers which monitor rates for different input conditions and for each of the trigger menu items. The menu items are monitored before and after the deadtime and prescaler logic, providing a measurement of the live-time fraction for use in offline analysis. Most of the scalers integrate over all filled bunches and are disabled for empty bunches. However, limited provision is also made for monitoring rates separately for each bunch in the machine.

There are numerous requirements concerning the design and operation of the CTP. A very important design requirement is that the latency of the CTP should be no more than four BCs (100 ns). The operational requirements relate to control (e.g., setting up the trigger menus), monitoring that the CTP is working as specified, test and diagnostic facilities, reliability and error rates, and maintenance. The CTP system contains numerous programmable timing parameters to synchronize and align input data, etc., and a detailed strategy is required to set up these parameters in order to meet the various run requirements (for beam–beam collisions, cosmic-ray triggers, test and calibration triggers, etc). The requirements placed on the CTP are given in detail in Ref. [15-3].

15.2 Hardware design specification

15.2.1 Overview

A block diagram of the CTP is shown in Figure 15-1. A total of 128 bits of input data are received by the CTP as differential electrical signals. These data are from the muon and calorimeter trigger processors, and from other sources. The CTP must synchronize these data with the local board clock and align them in order to compensate for the spread in latencies on the different inputs.

Once aligned, the input data are combined using programmable combinatorial logic to form 96 trigger items. These items can be masked and prescaled individually before they are combined

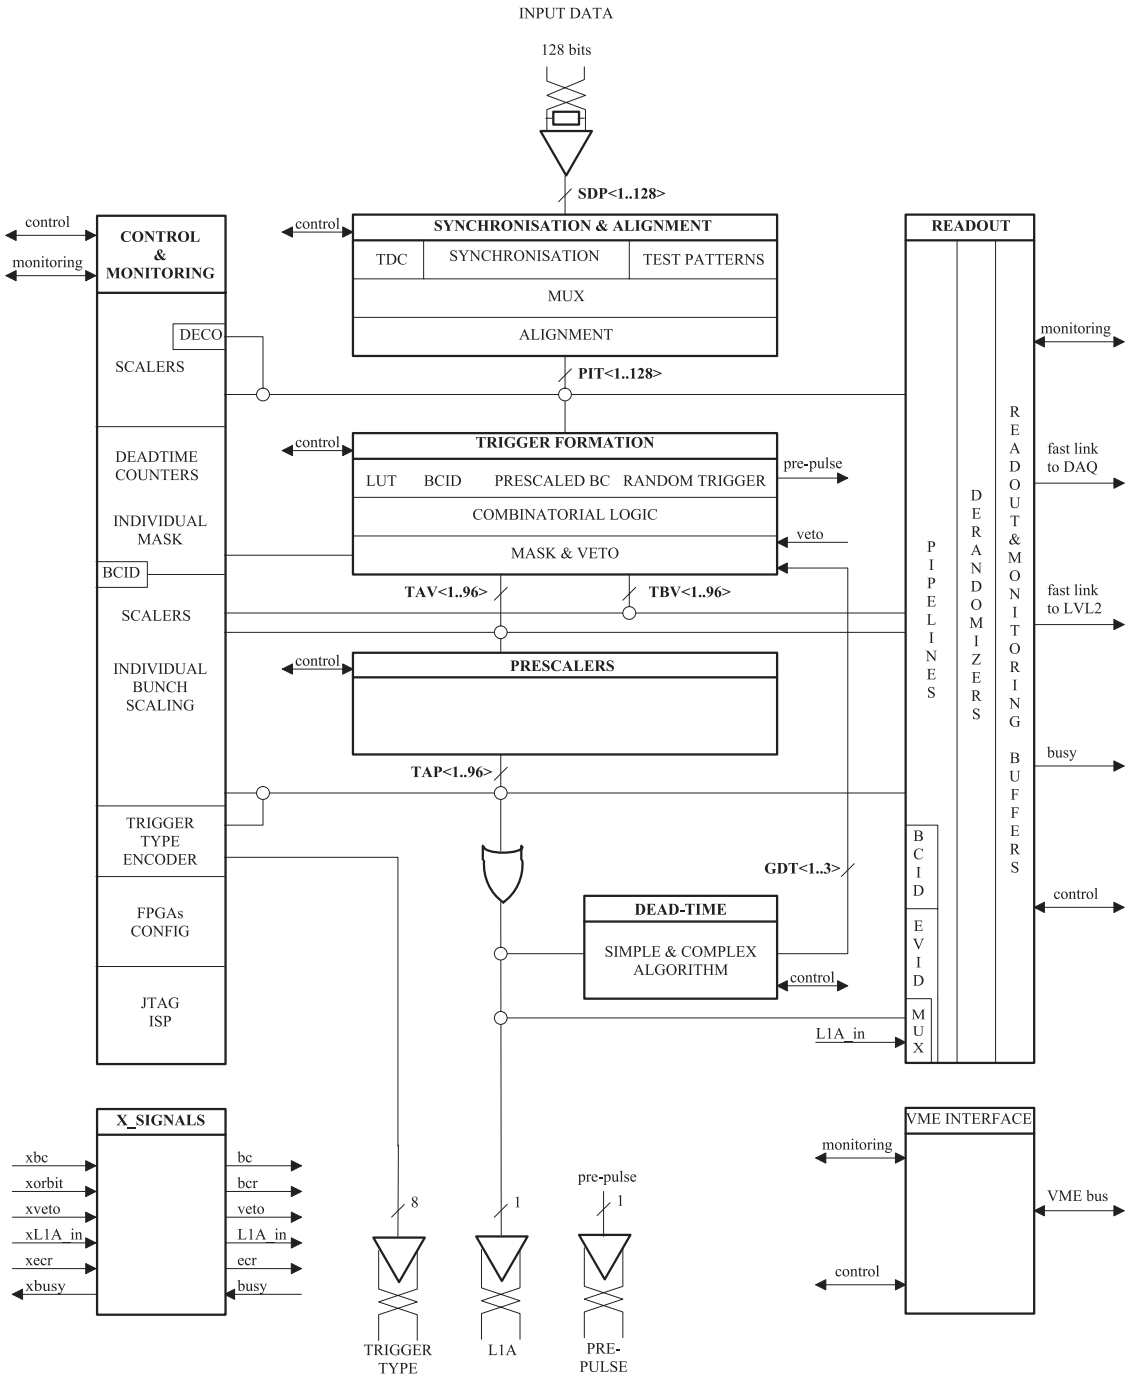

Figure 15-1 Block diagram of the CTP.

(using a logical OR) to form the overall LVL1 trigger decision. Deadtime is introduced to prevent the LVL1 trigger accepting events at a rate greater than that which can be handled by the detector front-end and readout systems.

When the CTP accepts an event it transmits to the front-end electronics, via the TTC system, the L1A signal and an 8-bit trigger-type word containing encoded information about the 96 trigger items. Detailed information about the selection criteria for each accepted event is transmitted from the CTP to the DAQ and LVL2 systems, via separate paths.

A context diagram, showing the interfaces between the CTP and external systems, is shown in Figure 15-2. In addition to the interfaces with the TTC, DAQ and LVL2 trigger systems mentioned above, there is a monitoring path over which the data from the scalers can be read out; these scalers are used to monitor the rates and the deadtime. There is also a link to the control system which is used to configure and test the system. The CTP accepts an external VETO signal derived from the ROD_BUSY signals of the front-end systems (see Section 20.6). This VETO signal is used to inhibit triggers should buffers in the front-end systems become full.

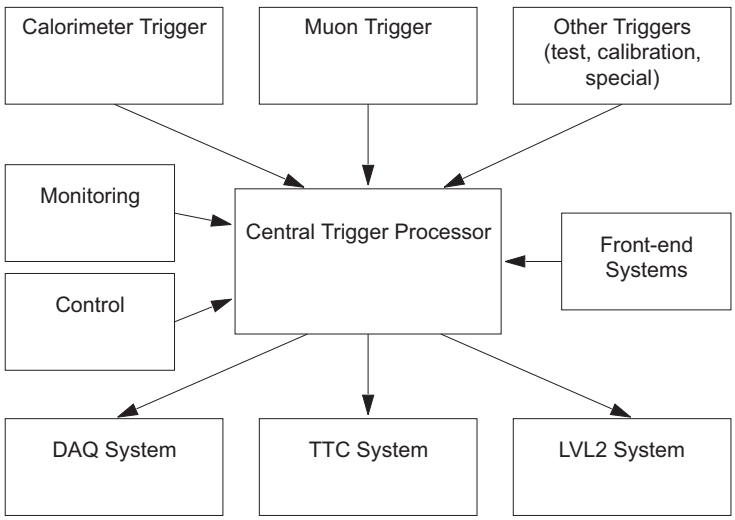

Figure 15-2 Context diagram for the CTP.

The main parameters of the CTP design are summarized in Table 15-1. The sections below describe the function and implementation of the various CTP logic blocks. Included is a description of all logic on the time-critical trigger-decision path, and that associated with readout and monitoring. A more detailed specification of the CTP is given in Ref. [15-4].

Table 15-1 The main parameters of the CTP design.

CTP Parameter	Value
Data input	128 bits
Trigger menu items	96
Maximum prescale factor	16×10^3
Latency	< 100 ns (4 BCs)
Deadtime	0–16 BCs after each trigger
	1–32 triggers in 0–1.7 ms
Number of boards	1–2

15.2.2 Synchronization and alignment

As can be seen from Figure 15-1, the first block of logic in the CTP data path is the synchronization and alignment block. The functions performed by this logic are:

- the synchronization of the input data with the local CTP clock;
- the alignment of input data so that all data originating from the same bunch crossing reach the subsequent logic in the same clock cycle;
- the introduction of internal test data to the CTP data path.

The synchronization of the input data with the local clock is necessary because although the full LVL1 trigger is driven by the same global clock, derived from the LHC machine clock, the phase varies in different parts of the system. Synchronization is achieved by using either the leading or trailing edge of the CTP clock to sample the input data. As indicated in Figure 15-3, this is implemented by including or excluding a flip-flop, clocked by the inverted clock signal, in the data path. It should be noted that the logic shown in the figure is repeated for each of the 128 inputs. A TDC is included in the design of the CTP in order to facilitate the timing-in of the system. This TDC measures the relative phase of transitions in the input data and the local clock.

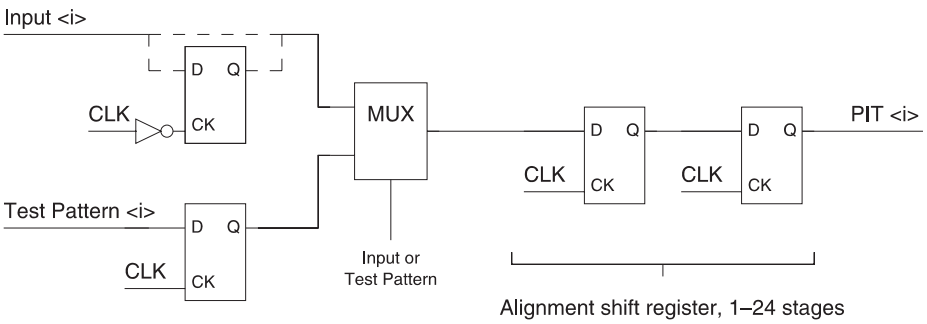

Figure 15-3 Synchronization and alignment logic.

Once the input data have been synchronized with the local clock they must be aligned in time, so that all of the data from one physical bunch crossing reach the subsequent combinatorial logic in the same clock cycle. This alignment is performed using variable-length shift registers. The input that arrives latest is passed without delay to the subsequent logic. Other inputs are passed through the number of delay elements necessary to bring them into alignment with the last one.

Test data can be introduced into the CTP data path via the multiplexer shown in Figure 15-3. These data are generated internally in the CTP, in the logic described in Section 15.3.5.

The synchronization and alignment block is implemented using FPGA-based logic. The insertion of synchronization flip-flops and the length of the alignment registers is set by *in situ* configuration of the FPGAs. As the timing parameters will be stable once the system has been set up, the need to download a new configuration in order to change them does not present problems.

15.2.3 Trigger formation

The trigger-formation logic is responsible for the following tasks: the formation of the 96 trigger-menu items, the gating of these 96 items with various mask and veto signals and the generation of a ‘prepulse’ signal. The logic used to perform each of these tasks is described below.

15.2.3.1 Combinatorial logic

The combinatorial logic forms the 96 trigger-menu items implemented by the CTP. This is done in two stages, as shown in Figure 15-4. In the first stage, a bank of 512-kbyte SRAMs acting as Look-Up Tables (LUTs) processes groups of inputs and produce intermediate results. The distribution of the 128 bits of input data between the LUTs is optimized in order to allow as much of the trigger processing as possible to be carried out at this stage, where there is no limitation on the combinatorial logic that can be implemented. The output of each LUT is an 8-bit encoded word.

The data produced by the LUTs are correlated in the second stage of combinatorial logic, which is implemented using high-density CPLDs. The result of this second stage of processing is a pattern of 96 bits indicating which menu items have been satisfied. Each menu item formed in the CPLDs can use up to 36 inputs. Four of these inputs are reserved for the inhibit signals described in Section 15.2.3.2. The remainder can be selected from the following data:

- the output of the LUTs (80 bits);
- random signals (2 bits);
- prescaled clocks (2 bits);
- bunch-crossing mask (8 bits).

The random signals, prescaled clocks and bunch-crossing mask used here are all formed locally. The random signals allow physics conditions to be used in conjunction with random triggers. The prescaled clocks, which are formed from the LHC clock prescaled by up to 16×10^6, allow the component conditions of single trigger items to be prescaled at different rates. The bunch-crossing mask, which is formed by feeding the local BCID number into a LUT, allows any BC within the LHC machine cycle to be masked for any trigger item. This could be used to mask out noisy bunches, for example, or to allocate individual calibration triggers to specific windows within the machine cycle.

Using current technology, four CPLDs are required to implement the second stage of combinatorial logic. Each CPLD can receive up to 72 inputs. These inputs must include all of the data required to build the trigger items formed in that CPLD, and the mask and veto signals described in Section 15.2.3.2.

Although it is not thought to be necessary, more than 96 trigger conditions can be defined by the combinatorial logic if different trigger conditions are combined into single menu items using a logical OR function. It should be noted, however, that such combined trigger conditions cannot be individually masked, or monitored, or assigned different deadtime parameters. Furthermore, individual prescaling requirements for such combined conditions cannot be handled by the prescaler block (see Section 15.2.4).

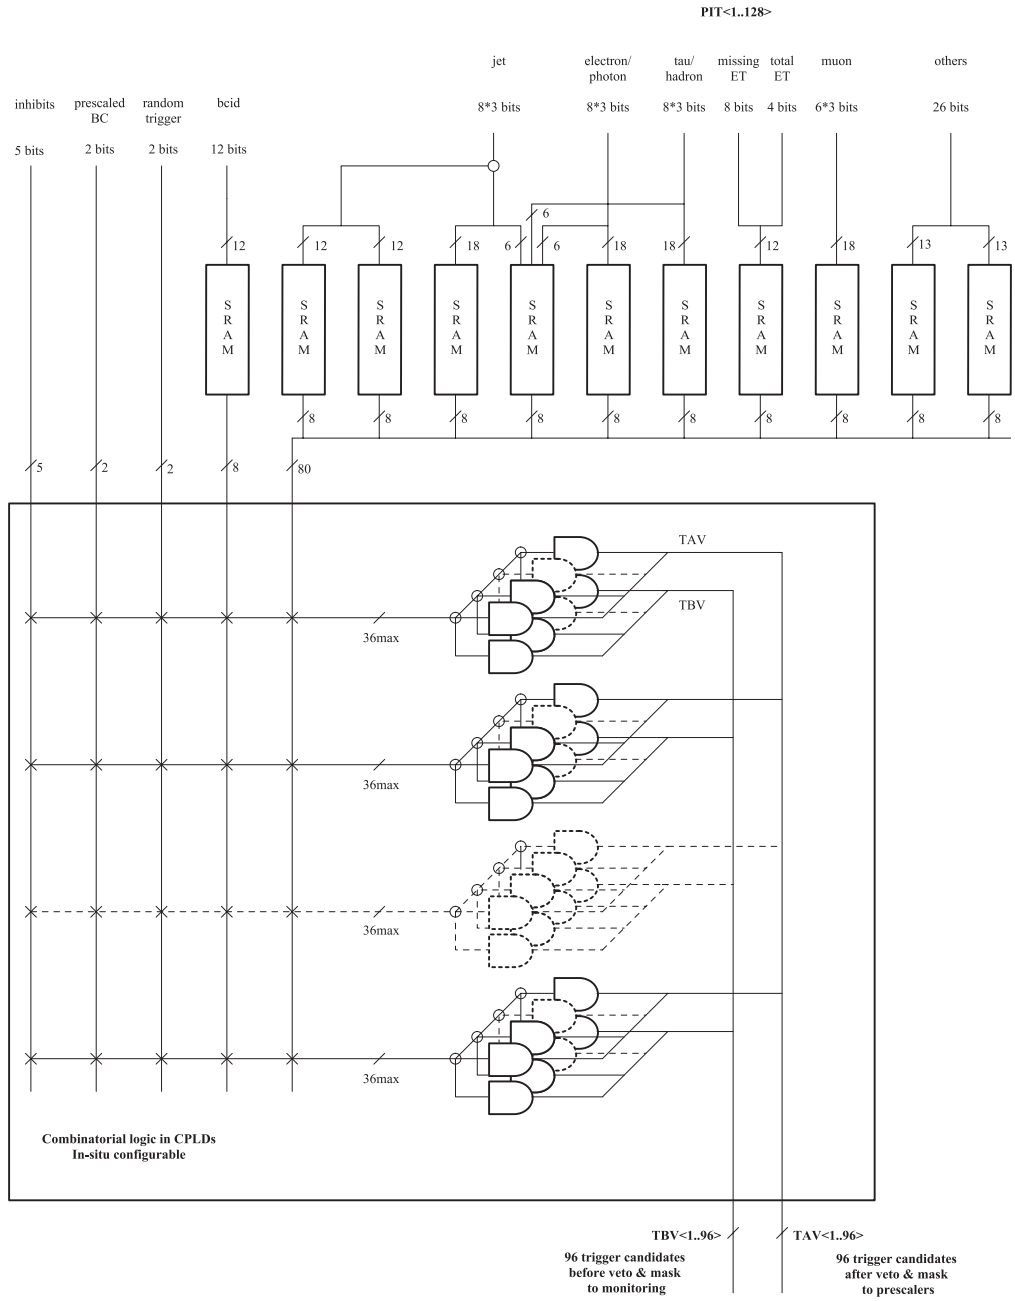

Figure 15-4 CTP trigger formation logic.

The CTP can be used to trigger on the scalar sum of the E_T of jets (Section 4.6.2). The jet inputs to the CTP from the calorimeter trigger provide a histogram of jet multiplicity versus E_T threshold. This can be used to form various estimates of the jet E_T sum, for example:

$$\sum E_{\mathrm{T}}^{\mathrm{jet}} = N\left(j_{\max}\right) E_{\mathrm{T}}^{\mathrm{thresh}}\left(j_{\max}\right) + \sum_{j=2}^{j_{\max}} \left(N\left(j-1\right) - N\left(j\right)\right) \frac{E_{\mathrm{T}}^{\mathrm{thresh}}\left(j-1\right) + E_{\mathrm{T}}^{\mathrm{thresh}}\left(j\right)}{2}$$

where $N(j)$ is the jet multiplicity for the jet threshold $E_T(j)$ and j = 1–8 is the threshold index.

As shown in Figure 16-4, the jet inputs to the CTP are fanned out to two sets of LUTs. This is necessary to avoid excessive use of the internal resources of the CTP (i.e. the number of connections between the LUTs and the CPLDs). One set of LUTs can be used to implement the 'standard' jet triggers (requirements on jet multiplicity and E_T values), while the other set can be dedicated to the calculation of the jet E_T sum.

15.2.3.2 Mask and veto

Each of the 96 trigger items formed by the combinatorial logic is gated with the following signals:

- an individual masking bit for that trigger item;
- a global VETO signal produced by the external OR of the ROD_BUSY and external veto signals (Section 20.6);
- the simple deadtime veto produced by the deadtime logic (Section 15.2.5);
- one of the two complex deadtime vetoes (either high- or low-priority) produced by the deadtime logic (Section 15.2.5).

The gating of the trigger items with the mask and veto signals is performed in the CPLDs that implement the final stage of combinatorial logic (Section 15.2.3.1). The data output from these CPLDs to the next stage of processing logic (the prescalers) are 96 trigger items with masks and vetoes applied. For monitoring purposes only, the CPLDs also output copies of the 96 trigger items with no masks and vetoes applied (see Section 15.2.4).

15.2.3.3 Prepulse

For calibration triggers the CTP should be able to generate a 'prepulse' signal which precedes the L1A by a known duration. This signal could be used to fire pulser systems in detectors, so that the data generated are captured when the associated L1A initiates an event readout. In order for the prepulse signal to precede the L1A, it must be generated by the trigger-formation logic and it must bypass all subsequent logic, such as the prescalers. The full implementation of the prepulse signal has to be investigated.

15.2.4 Prescalers

The prescaler logic allows each of the 96 trigger candidates to be prescaled individually by a programmable factor. The logic is implemented in FPGAs and the maximum prescale factor available depends upon the level of FPGA technology available. Using current technology a prescale factor of 1–4096 can be implemented for each trigger item. Using technology available at the time of construction, it is thought that a maximum prescale factor of 16×10^3 can be achieved. If greater prescaling factors than this are required, they can be implemented at the trigger-formation stage by using the prescaled-BC or random-trigger inputs in the trigger item definition (Section 15.2.3).

15.2.5 Deadtime

The deadtime logic produces two different kinds of deadtime veto, labelled ‘simple’ and ‘complex’. The simple deadtime veto is a signal of programmable duration (0–16 BCs) which is introduced after each LVL1 trigger. A simple deadtime veto of four BCs is required to prevent overlapping readout time-frames for some front-end systems.

The purpose of the complex deadtime veto is to limit the rate at which LVL1 triggers are generated and thus prevent the derandomizers of some front-end electronics systems filling up. The complex deadtime veto is introduced so that no more than N triggers are generated within any period of duration T, where N is programmable within the range 1–32 and T is programmable within the range 0–1.6 milliseconds, in units of 25 ns. A ‘leaking bucket’ algorithm is used to implement this. The bucket, which leaks at a constant rate, R, has its contents increased by a given value, X, every time a LVL1 trigger occurs. When the bucket is full deadtime is introduced. The values of the programmable parameters X and R directly determine the values of the parameters N and T given above.

The deadtime logic provides two, independently-programmable, complex vetoes and each trigger item is assigned to one of these. This allows high-priority and low-priority trigger items to be defined. In order to make best use of the total permissible LVL1 trigger rate, low-priority trigger items can be given a maximum rate lower than that of high-priority trigger items.

15.2.6 Readout

The CTP readout logic is responsible for the following tasks: generation of the local bunch-crossing-identification (BCID) and event-identification (EVID) numbers, transfer of data to the ATLAS DAQ for permanent storage if an event is accepted by LVL2 and the event filter, and transfer of data to LVL2 for RoI-building. The logic responsible for each of these tasks is described below.

15.2.6.1 BCID and EVID generation

The local BCID and EVID numbers are generated in 12-bit counters. The BCID counter is reset by the bunch-counter reset (BCR) signal, which is generated on receipt of the LHC ORBIT signal from the TTC (see Chapter 16).

15.2.6.2 Transfer of data to DAQ and LVL2

As the CTP runs, the following information is copied from the data path to pipeline memories:

- the 128 bits of input data after synchronization and alignment;
- the prescaled-clock, random-trigger and BC-mask inputs to the trigger formation logic;
- the 96 trigger items before any vetoes have been applied;
- the 96 trigger items after all vetoes have been applied;
- the 96 items after prescaling;
- the EVID and BCID numbers.

Each time an L1A is generated these data are transferred from the pipeline memories to a derandomizing buffer. This is done for a programmable window of BCs before and after that which produced the L1A (0–5 BCs in total). The data are then transferred to the ATLAS DAQ using the standard ATLAS event format and readout links. When the CTP's derandomizing buffers are nearly full a ROD_BUSY signal is generated. This signal is handled in the same manner as all other ATLAS ROD_BUSY signals: it is fed into the module described in Section 20.6 which generates the VETO signal for the CTP (Section 15.2.3.2).

The same data that are transferred to the DAQ are also transferred to LVL2, although for only the trigger BC. An S-link can be used for this purpose. Less than 500 ns are needed to transmit all of the data from one L1A, although the time between the occurrence of the L1A and the availability of these data in LVL2 depends upon the instantaneous L1A rate.

The L1A used to initiate the transfer of data to the DAQ and the LVL2 trigger is normally generated internally by the CTP. It is possible, however, for the CTP to accept an external L1A signal and use it to initiate the transfer process.

15.2.7 Control and monitoring

The monitoring and control logic contains the following elements: scalers, trigger-type encoder, individual trigger mask, FPGA configuration logic and JTAG control logic. These elements are described below.

15.2.7.1 Scalers

The scalers included in the control and monitoring logic count the number of bunch crossings during which various data bits are set. Two types of scalers are provided: those which count the data for all filled bunch crossings, and those which count the data for one programmable bunch crossing within the LHC machine cycle. For every filled bunch crossing, the following data are counted:

- the 128 bits of input data after synchronization and alignment (with encoded data being decoded to ensure that sensible quantities are being scaled);
- the 96 trigger items before any vetoes and prescaling have been applied;
- the 96 trigger items after all vetoes and prescaling have been applied;
- the overall trigger result;
- the three deadtime vetoes generated by the logic described in Section 15.2.5;
- the external, global VETO signal.

For any one programmable bunch crossing within the LHC machine cycle the following data can be counted:

- the 128 bits of input data after synchronization and alignment (with decoding performed where necessary);
- the 96 trigger items before any vetoes and prescaling have been applied;
- the 96 trigger items after all vetoes and prescaling have been applied.

Upon receipt of a reset signal the content of the scalers is loaded to associated readout registers and the scalers are cleared. The scaler data can then be read from the registers, via VME, in the period before the next reset is issued. The frequency with which the scalers must be read out and cleared depends upon the maximum size of scaler which can be implemented using the available FPGA technology and board space.

The scalers which need to have the largest capacity are those that monitor all filled bunch crossings. Using currently-available technology, 12-bit scalers can be provided for this purpose. These scalers can handle any data rate for one LHC machine cycle (88 μs). It is believed, however, that advances in FPGA technology will make 16-bit scalers possible at the time of construction. If 16-bit scalers are used to monitor every filled bunch crossing, and 4-bit scalers are used to monitor individual bunches, the scalers only need to be read and cleared once every 16 LHC cycles.

15.2.7.2 Trigger Type Encoder

This logic encodes the 96 trigger items, after all vetoes and prescaling have been applied, into an 8-bit trigger-type word. If an event is accepted by the CTP, this word is transmitted with the L1A to the TTC system.

15.2.7.3 Individual Trigger Mask

The 96-bit individual trigger mask, which is used in the mask-and-veto logic described in Section 15.2.3.2, is stored here in a VME-accessible register.

15.2.7.4 FPGA configuration logic

This block contains all of the logic necessary to configure the FPGAs included in the CTP design. To configure the devices the configuration data are first downloaded, via VME, to an on-board RAM. The data are then transferred from this RAM to the FPGAs themselves, under the control of a state-machine. The state-machine, which is implemented in a CPLD, generates all of the required clock, status and control signals.

15.2.7.5 JTAG control logic

A JTAG chain links all of the FPGAs and CPLDs used in the CTP. The JTAG control logic generates all of the signals necessary to control that chain. In addition to testing the interconnections of the logic, the JTAG chain can be used to perform *in situ* programming of the CPLDs on the board.

15.2.8 Test and diagnostic facilities

To aid testing and fault-diagnosis, the CTP has diagnostic memories which can capture data at strategic points in the system. Such memories are situated at all data interfaces to the CTP and at either end of any cable or backplane links within the CTP.

The CTP also includes two internal sources of test data — a RAM, which can be loaded with test data, and a pseudo-random pattern generator for much longer test sequences. Data from either

of these sources can be introduced to the CTP data path via the mulitplexer in the synchronization and alignment block (Section 15.2.2). The length of test sequences available from the RAM and the pseudo-random pattern generator will depend upon the board-space available to implement this logic.

15.2.9 Implementation

The CTP will be implemented using field-programmable logic: FPGAs and CPLDs. SRAMs will be used to store those CTP parameters that will need to be changed throughout the lifetime of ATLAS, for example the trigger menus. Changing these parameters will thus require only a simple VME access. Those parameters that should remain stable after the initial calibration of the CTP, for example the length of the alignment pipelines (Section 15.2.2), will be written into the configuration files of the field-programmable logic. They will still be programmable, but changing them will require the generation of a new design file.

It is anticipated that the CTP will be implemented as a single 9U VME board, although a complete design is needed before this can be confirmed. This evaluation is based on knowledge gained from the CTP demonstrator programme and on the predicted evolution of FPGAs and CPLDs. If it proves necessary to split the design of the CTP over multiple boards, it is desirable that the main data path (from the input data through to the L1A signal) should be implemented on a single board, in order to reduce the latency. As shown in Section 15.2.10, however, this is not a requirement.

15.2.10 Latency

A timing diagram for the CTP is shown in Figure 15-5. This diagram is based on the worst-case design scenario for the CTP latency. It assumes that the trigger-formation logic cannot be implemented on a single board and that some inter-board communication is therefore necessary at this stage. A total of 50 ns are allowed for trigger formation (from ‘input data after alignment’ to ‘trigger menu items’ in the diagram). The resulting latency of 100 ns for the CTP is the upper limit of that expected for the final design.

It should be noted that the latency of the CTP is the duration between the arrival of the latest bit of input data and the generation of the L1A. The alignment pipelines used to delay earlier signals do not, therefore, contribute to the latency.

15.2.11 Status of the CTP design

The full specification of the CTP is not due to be finalized until 1999 (see Section 15.7). However, much is already known about the design and implementation of the CTP. For example, the feasibility of using FPGAs and CPLDs to implement the core CTP logic with the required speed and density has been proven in the CTP demonstrator (Section 15.6). Much of the logic designed for that demonstrator will be used in the final CTP, either directly or with some degree of modification. Where extensive modification or expansion of the logic is required, for example in the trigger-formation block, studies have already been conducted to determine the implementation of the proposed design.

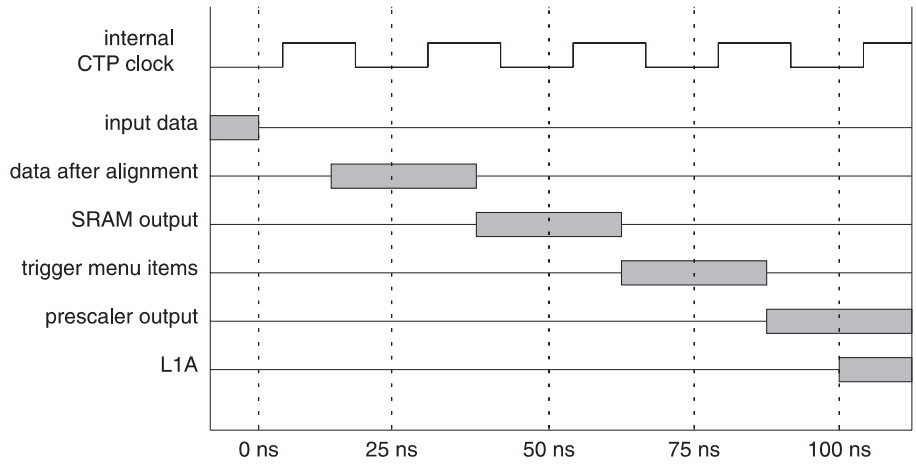

Figure 15-5 Timing diagram of the CTP.

15.3 Software Design Specification

The CTP requires software to perform calibration, control, monitoring, readout, testing and fault-diagnosis. The following sections describe the software proposed for each of these tasks.

15.3.1 Calibration

The parameters which need to be calibrated for the CTP are those associated with the synchronization and alignment logic (Section 15.2.2). For each input signal, the phase of the sampling clock and the length of the alignment pipeline need to be calibrated. As setting these parameters to new, arbitrary values requires the generation of an FPGA configuration file, the calibration process cannot be fully automated; the calibration software can determine the required parameter values but not set them.

The phase of sampling clock required by each input signal can be determined locally, using the TDC included within the CTP (Section 15.2.2). To determine the required lengths of the alignment pipelines, the calibration software must correlate the data input to the CTP from different sources. This could possibly be done offline. See Section 15.4 for more details on this subject.

15.3.2 Control

The CTP is controlled via the ATLAS DAQ run-control system [15-5]. The control software must be able to load and read all of the register-based parameters of the CTP and to configure all FPGA- and CPLD-based logic.

Software is required to translate between the actual control and configuration data used by the CTP and a more user-friendly format of data that can be understood by non-experts. It should be possible to program new trigger menus, for example, without any knowledge of how those

menus are implemented. Graphical user interfaces should be used where possible. For trigger menus, however, the large number of degrees of freedom may make this is impractical. A better solution might be to provide a text-based equation editor and a compiler to translate between these equations and the actual data downloaded to and from the trigger-formation logic (Section 15.2.3).

Generating configuration data for FPGAs and CPLDs is beyond the scope of the control software. It can only chose between the configuration files available and new files, if required, must be generated using proprietary software. This should not present a problem as those parameters implemented in FPGAs and CPLDs should remain stable after the initial commissioning phase of the experiment.

15.3.3 Readout

The CTP is read out by the ATLAS DAQ. It complies fully with the DAQ requirements on front-end systems given in Ref. [15-1], and therefore places no special requirements on the DAQ readout software.

15.3.4 Monitoring

Local monitoring software is required to read out and analyse the data from the scalers described in Section 15.2.7.1. It is foreseen that this software will run on microprocessors situated in the CTP crate itself. VME Direct Memory Access (DMA) can be used to read out the data at the required rate, provided scalers of 12 bits or greater are used. Real-time analysis of the data should be performed, and any potential problems should be reported to the run-control system. The data should be written, via the run-control system, to permanent storage in the run log.

The ATLAS DAQ, which has access to the readout data described in Section 15.2.6, should also be able to monitor the operation of the CTP and to report any potential problems.

15.3.5 Test and Diagnosis

To allow the CTP to be tested and any faults to be diagnosed, software is required which can perform the following tasks:

- provide local control of the CTP;
- test all programmable memories on the CTP, for example, control registers and LUTs;
- test any cable and backplane connections within the CTP by comparing the contents of diagnostic memories situated before and after the interconnections;
- test the data links to and from the CTP by comparing the contents of diagnostic memories situated at either end of each link;
- test the CTP logic by comparing the output generated for test data with the results predicted by a software model of the CTP;
- initiate and analyse the results of a JTAG boundary scan.

The control and diagnostic software for the CTP should be available locally. It should provide a user-friendly, graphical interface which can be used by non-experts, and it should allow the tasks described above to be run interactively or as part of one or more automated test procedures. These automated test procedures should thoroughly exercise the CTP and they should be available to the DAQ test manager [15-5].

In order to test the data links to and from the CTP, it is necessary for the CTP test and diagnostic software to interface with the test and diagnostic software of the LVL1 muon processor, LVL1 calorimeter processor, LVL2 and TTC systems. Cooperation is therefore required between the groups responsible for these systems. Cooperation is also desirable, within the LVL1 trigger community in particular, in order to share skills and to minimize duplicated effort.

15.4 Strategy for setting up the timing of the trigger

This section outlines the procedures necessary to set up the timing of the CTP. A full strategy for this process will be formulated and documented at a later date.

The first task required to set up the timing of the CTP is the synchronization of the input data with the CTP clock. Test data supplied by the input systems can be used for this purpose (provided the test data have the same timing characteristics as physics data). The TDC included in the CTP logic can be used to measure the phase of all input data, under the control of local software. The FPGAs responsible for the synchronization of the input data can then be reconfigured, according to the results obtained from the software. See Sections 15.2.2 and 15.3.1 for more details.

Once synchronized, the input data must be aligned so that all data from one physical bunch crossing reach the CTP trigger-formation logic in the same clock cycle. For this to be accomplished it is necessary to correlate the latencies of all subdetectors and subtriggers that supply the CTP with data. This could be done by triggering on real data from beam–beam collisions or cosmic ray events, for example, and observing the spread in latencies in the data recorded by the CTP. The ability of the CTP to record a time-frame of bunch-crossings around each L1A is essential for this task.

In normal running mode, the readout of the CTP is initiated by the internally-generated L1A. In this mode, therefore, no timing calibration is necessary to ensure that the correct bunch-crossing is read out; only the BCID number must adjusted to match the FE_BCID (see Chapter 19). It is possible, however, for the readout of the CTP to be initiated by an external L1A signal. The timing issues associated with this feature of the CTP need to be investigated.

15.5 Trigger menu implementation

The capacity of the CTP in terms of number of inputs, number of trigger combinations and the availability of internal resources has been checked by implementing an example of a complete trigger menu. This menu, shown in Table 15-2, is based on the high-luminosity trigger menu shown in Table 3-1 of Section 3.1. However, it includes, in addition, a large number of prescaled trigger items. Typically, there are several triggers for each object type, with different threshold values and prescale factors.

The menu uses a total of 115 bits of input information, compared to the available total of 128 bits. These are:

- $6 \times$ 3-bit muon multiplicities;
- $8 \times$ 3-bit EM cluster multiplicities;
- $8 \times$ 3-bit hadron / tau multiplicities;
- $8 \times$ 3-bit jet multiplicities;
- 8 bits of information on missing-E_T;
- 4 bits of information on total-E_T;
- 4 bits reserved for information on forward energy;
- 12 bits used for detector calibration triggers;
- 1 test trigger input.

It should be noted that it would be possible to combine the calibration and test trigger inputs externally to the CTP, for example encoding the calibration / test trigger type in four bits. The number of physics-trigger inputs is determined by the design of the muon and calorimeter trigger processors.

The menu contains a total of 85 items, including twelve calibration triggers and a test trigger, compared to the available 96 trigger items. It should be noted that groups of menu items that have the same prescale factor and trigger priority could be consolidated, as described in Section 15.2.3.1.

The use of internal resources has also been assessed. The most heavily used part of the system is that related to jet triggers, where 30% of the available number of combinations at the output of the lookup table are used.

It should be noted that the menu in Table 15-2 is shown only to demonstrate that the CTP is able to accommodate complicated sets of trigger conditions. The details of the trigger conditions and threshold settings that are shown for prescaled triggers are for illustration only. They do not represent an optimized choice.

Table 15-2 Example Trigger menu used to check internal resources of the CTP. The trigger conditions and thresholds have not been optimized. However, the complexity of the menu is representative of that required.

Trigger Condition	Comment
Muon triggers	
Single muon, p_T > 5 GeV	Prescaled
Pair of muons, p_T > 6 GeV	
Single muon, p_T > 6 GeV	Prescaled
Single muon, p_T > 10 GeV	Prescaled
Pair of muons, p_T > 10 GeV	High priority
Single muon, p_T > 10 GeV AND single isolated EM cluster E_T> 15 GeV	
Single muon, p_T > 15 GeV	Prescaled
Single muon, p_T > 20 GeV	
Single muon, p_T > 25 GeV	High priority
Electromagnetic cluster triggers	
Single isolated EM cluster, E_T > 7 GeV	Prescaled
Single isolated EM cluster, E_T > 10 GeV	Prescaled
Single isolated EM cluster, E_T > 15 GeV	Prescaled
Single isolated EM cluster, E_T > 20 GeV	Prescaled
Pair of isolated EM clusters, E_T > 20 GeV	
Single isolated EM cluster, E_T > 20 GeV AND missing E_T > 20 GeV	
Single isolated EM cluster, E_T > 30 GeV	
Pair of isolated EM clusters, E_T > 30 GeV	High priority
Single isolated EM cluster, E_T > 40 GeV	High priority
Single isolated EM cluster, E_T > 50 GeV (relaxed isolation)	High priority
Single EM cluster, E_T > 60 GeV (no isolation)	High priority
Jet triggers	
Single jet, E_T > 30 GeV	Prescaled
Four jets, E_T > 30 GeV	Prescaled
Three jets, E_T > 30 GeV	Prescaled
Single jet, E_T > 50 GeV	Prescaled
Four jets, E_T > 50 GeV	Prescaled
Three jets, E_T > 50 GeV	Prescaled
Single jet, E_T > 70 GeV	Prescaled
Four jets, E_T > 70 GeV	Prescaled
Three jets, E_T > 70 GeV	Prescaled
Single jet, E_T > 90 GeV	Prescaled
Four jets, E_T > 90 GeV	
Three jets, E_T > 90 GeV	Prescaled
Single jet, E_T > 100 GeV	Prescaled
Three jets, E_T >100 GeV	Prescaled
Single jet, E_T > 100 GeV AND missing E_T > 100 GeV	
Single jet, E_T > 130 GeV	Prescaled
Three jets, E_T > 130 GeV	
Single jet, E_T > 200 GeV	Prescaled
Single jet, E_T > 200 GeV AND missing E_T > 200 GeV	High priority
Single jet, E_T > 290 GeV	
Tau / hadron triggers	
Single tau / hadron candidate, E_T > 20 GeV	Prescaled
Single tau / hadron candidate, E_T > 30 GeV	Prescaled
Single tau / hadron candidate, E_T > 40 GeV	Prescaled
Single tau / hadron candidate, E_T > 60 GeV	Prescaled
Single tau / hadron candidate, E_T > 60 GeV AND missing E_T > 60 GeV	
Single tau / hadron candidate, E_T > 80 GeV	

Table 15-2 Example Trigger menu used to check internal resources of the CTP. The trigger conditions and thresholds have not been optimized. However, the complexity of the menu is representative of that required.

Trigger Condition	Comment
Tau / hadron triggers (continued)	
Single tau / hadron candidate, E_T > 100 GeV (relaxed isolation)	
Single tau / hadron candidate, E_T > 120 GeV (more relaxed isolation)	
Single tau / hadron candidate, E_T > 140 GeV (even more relaxed isolation)	
Missing energy triggers	
Missing E_T > 20 GeV	Prescaled
Missing E_T > 30 GeV	Prescaled
Missing E_T > 40 GeV	Prescaled
Missing E_T > 50 GeV	Prescaled
Missing E_T > 60 GeV	Prescaled
Missing E_T > 100 GeV	Prescaled
Missing E_T > 150 GeV	Prescaled
Missing E_T > 200 GeV	
Total energy triggers	
Total scalar E_T> 1000 GeV	Prescaled
Total scalar E_T> 1200 GeV	Prescaled
Total scalar E_T> 1400 GeV	Prescaled
Total scalar E_T> 1600 GeV	
Other special triggers	
SUM(E_T^{JET}) (PRESCALED)	Prescaled
SUM(E_T^{JET}) (PRESCALED)	Prescaled
SUM(E_T^{JET}) (PRESCALED)	Prescaled
SUM(E_T^{JET})	
RANDOM 40 kHz	Prescaled
RANDOM 10 Hz	
PRESCALED BC	Prescaled
FORWARD ENERGY (PRESCALED)	Prescaled
FORWARD ENERGY (PRESCALED)	Prescaled
FORWARD ENERGY (PRESCALED)	Prescaled
FORWARD ENERGY	
Calibration triggers	
PIXEL-1	Calibration trigger
PIXEL-2	Calibration trigger
SCT-1	Calibration trigger
SCT-2	Calibration trigger
TRT-1	Calibration trigger
TRT-2	Calibration trigger
Liquid-argon-1	Calibration trigger
Liquid-argon-2	Calibration trigger
Tile-calorimeter-1	Calibration trigger
Tile-calorimeter-2	Calibration trigger
Muon-1	Calibration trigger
Muon-2	Calibration trigger
TEST TRIGGERS	
EXTERNAL TRIGGER GENERATOR	

15.6 The CTP demonstrator programme

15.6.1 Motivations for the demonstrator programme

A demonstrator for the CTP — the CTPD — has been built. The motives for building CTPD were the following:

- to allow the implementation of the CTP in field-programmable logic to be evaluated;
- to allow a realistic estimate of the CTP latency to be made;
- to enable the CTP to be demonstrated together with other elements of the LVL1 trigger;
- to gain design experience.

A brief summary of the CTP demonstrator programme is presented below. The full specification of the CTPD is given in Ref. [15-6] and more detailed descriptions of the testing of the CTPD are given in Ref. [15-7] and [15-8].

15.6.2 Comparison of the CTPD with the CTP

The CTPD implements the core functionality foreseen for the final system. Crucially, the CTPD implements all of the functionality that lies in the critical time path, which means that a reliable estimate of the latency of the final system can be made from that of the demonstrator. The main difference between the CTPD and the proposed final system is size, as shown in Table 15-3. The demonstrator also features a simplified readout system. Rather than implement the full ATLAS readout system, in the event of a LVL1 trigger the demonstrator transfers data to an 8k-deep FIFO which can be read out via VME. This process is well suited to operation in a test-beam environment.

Table 15-3 Comparison of CTP Demonstrator with Final System.

	Final System	**Demonstrator**
Input data	128 bits	32 bits
No. triggers formed	96	32
No. boards	1–2	1
Readout	Derandomizer, readout driver	8k FIFO, VME
Latency	< 100 ns	62.5 ns (measured)

15.6.3 Implementation

The CTPD is implemented as a 40 MHz, fully synchronous, pipelined processor. It is housed on a single 9U VME board with J1, J2 and JAUX connectors. The board is an eight-layer PCB, four layers of which are signal planes. Extensive use is made of field-programmable technology and the core logic of the CTPD is contained within a total of eight FPGAs (ALTERA EPF81500ARC240-2) and eight CPLDs (ALTERA EPM7128EQC160-10P).

15.6.4 Latency

Considerable effort was made during the design of the CTPD to minimize the latency of the module. The three main areas of work which have contributed to minimizing the latency are:

- optimizing the clock phase at different parts of the circuit;
- optimal partitioning of functions between FPGAs to minimize communication between devices;
- careful routing of circuits, both within FPGAs and at board level;

The result of this work is a latency of 62.5 ns for the CTPD (compared to initial predictions of about 125 ns).

15.6.4.1 Testing the CTPD

The CTPD was initially tested in isolation, using data supplied either by the module's built-in test RAM or a digital signal generator. These tests proved that the CTPD functioned correctly at full speed (40 MHz).

In October 1997 the CTPD was run in an ATLAS test beam with the Tile calorimeter module-0 and the LVL1 calorimeter-trigger demonstrator system. This was the first time that a full LVL1 processing chain had been run together. Figure 15-6 shows a block diagram of the system that was used.

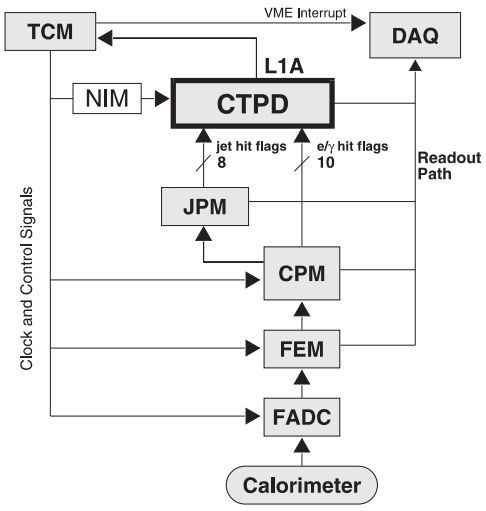

Figure 15-6 Block diagram showing the combined LVL1 calorimeter trigger and CTP demonstrator system.

The modules of the calorimeter-trigger demonstrator are described in detail in Chapter 7. The key operations performed on the data before they arrived at the CTPD were:

- analogue-to-digital conversion (FADC);
- bunch-crossing identification (TXM / FEM);
- electron / photon cluster-finding (CPM);
- jet finding (JPM).

The calorimeter-trigger demonstrator processed a 6×6 array of calorimeter towers. This resulted in the following data being available to the CTPD:

- 72 electron / photon hit flags: 36 electron / photon-cluster cells $\times$ 2 sets of thresholds;
- 8 bits of jet results: 4 jet cells $\times$ 2-bit words containing encoded information for three jet thresholds.

These data were of a different format to those expected for the final CTP as cluster counting, which will be performed by the final calorimeter trigger processor, was not implemented in the demonstrator. Therefore, rather than receive multiplicity values the CTPD received hit flags directly. The implementation of the trigger-forming logic in the CTPD is sufficiently flexible, however, for it to process these data.

One consequence of receiving hit flags at the CTPD rather than multiplicity values was that it greatly increased the amount of data to be processed. Of the 80 bits of data available to the CTPD, only 32 bits could be processed. The number of electron / photon and jet windows which were actually illuminated by the test beam at any one time, however, produced less data than this. Generally, the CTPD received a subset of data which consisted of eight electron / photon hit flags and eight bits of jet data.

One important difference existed between the CTPD and most of the other demonstrator modules present in the system. Whilst the CTPD was designed to run continuously and be read out in parallel with data-processing, all of the other modules except the FEM were designed to be halted during readout. It was therefore necessary to create a veto signal to prevent the CTPD triggering on spurious data received from the electron / photon and jet modules whilst they were halted. NIM logic was used to generate an appropriate CTPD global veto from existing system control signals.

15.6.5 CTPD Test Results

Figures 15-7 and 15-8 show two oscilloscope plots obtained during the stand-alone testing of the CTPD. In Figure 15-7, the lower of the two signals shown is an input bit to the CTPD, whilst the upper signal is the L1A (active low). The CTPD has been programmed to generate an L1A when the input bit is set, and it can be seen that it does this with a latency of 62.5 ns. In Figure 15-8, the lower signal is an input bit to the CTPD which is toggling with the same frequency as the clock. If no deadtime parameters were set, an L1A would be generated with every BC. It can be seen, however, that the following deadtime parameters have been set:

- Simple algorithm: 2 BC deadtime after each L1A.
- Complex algorithm: $\leq$ 3 L1As in 3.2 μs.

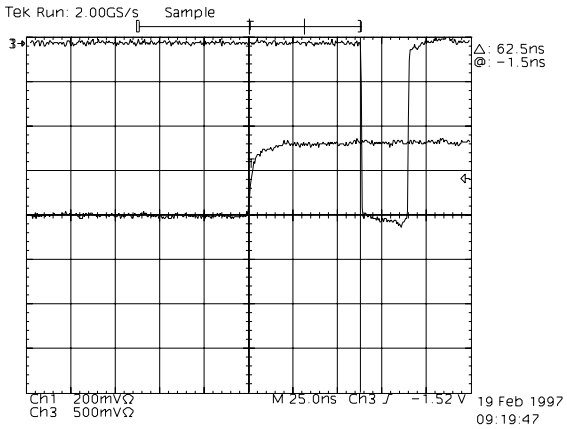

Figure 15-7 Oscilloscope plot showing the latency of the CTPD.

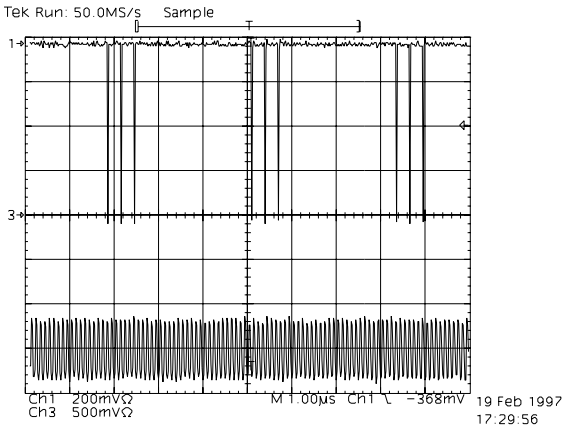

Figure 15-8 Oscilloscope plot showing the operation of the simple and complex deadtime algorithms on the CTPD.

Figure 15-9 shows data recorded during the combined beam tests for the calorimeter-trigger and CTP demonstrator systems. The data shown here are taken from a single event. Plot (a) shows data recorded at the earliest stage in the calorimeter trigger — a tile calorimeter pulse as captured in one FADC channel. This channel contributes to the cluster window whose output, for one set of trigger thresholds, is shown in plot (b). It can be seen that this cluster window produces a positive hit flag with a duration of one sample. The data shown in plot (b) are taken from the CPM output memory. Plot (c) shows the same data after they have been transmitted to the CTPD. The apparent discrepancy between the timing of the data in plots (b) and (c) is due to the different time-frames recorded by the CPMs and the CTPD for this event.

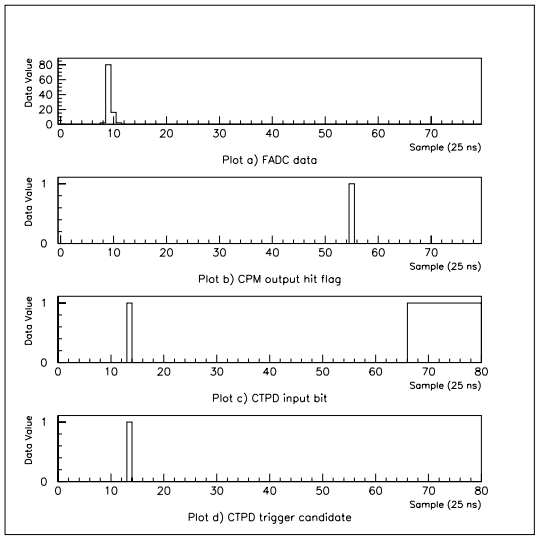

Figure 15-9 Data from a single event recorded at various stages in the LVL1 trigger demonstrator.

At sample 66 in plot (c) the data bit shown goes high for an extended duration. This marks the point at which the DAQ system has frozen the calorimeter-trigger system for readout — owing to the internal design of the CPM the hit flags are set high for this period. Plot (d) shows one of the CTPD candidate triggers from the processing stage immediately prior to L1A formation. This candidate trigger requires three or more electron / photon hits from the calorimeter trigger. A positive trigger decision has been produced for this event as electron / photon hits occur in a total of four channels, including that shown. Note that whilst the beam particle produces a trigger, none are produced by the electron / photon-hit flags set high during readout. This indicates that the CTPD global-veto input, which is asserted by the DAQ during readout, is working correctly.

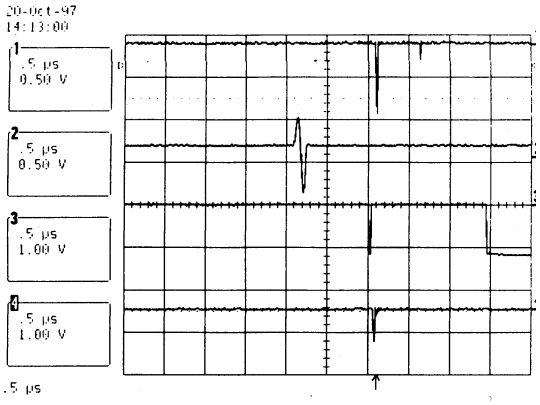

Figure 15-10 Oscilloscope plot showing the DAQ of the demonstrator system being triggered by the CTPD. Traces shown are (2) an analogue pulse input to a FADC; (3) a CPM hit output caused by this pulse; (4) the resulting CTPD L1A which is used as (1) the DAQ trigger.

Figure 15-10 shows an oscilloscope plot made during the beam tests. Here, L1As from the CTPD are being used to trigger the data-recording process, so the demonstrator system is acting as a real trigger, recording only those events which meet the trigger criteria. The FADC input in this case is provided by an analogue signal generator rather than the tile calorimeter and the CTPD trigger criterion is $\geq$ 1 electron / photon hit.

15.6.6 Conclusions drawn from the demonstrator programme

The CTP demonstrator implements the core functionality foreseen for the final CTP, including all functionality which lies on the critical time path. The demonstrator makes extensive use of field-programmable logic and has been optimized for minimum latency.

The CTPD works correctly at full speed (40 MHz) with a latency of 62.5 ns. This has been shown both in isolation and in conjunction with the LVL1 calorimeter-trigger demonstrator system. These combined tests were the first demonstration of a full processing path through the ATLAS LVL1 trigger.

The CTP demonstrator programme has produced the following conclusions for the final CTP:

- it is feasible to construct the CTP using field-programmable technology;
- the CTP can be implemented with a latency of $<$ 100 ns (where this upper limit assumes that the CTP is a two-board system).

15.7 Schedule, construction and assembly procedure

A schedule for the CTP project is shown in Figure 15-11. The CTP specification should be finalized by March 1999, allowing one year for design, production and stand-alone testing before the full CTP system is ready to be tested in an ATLAS test beam. All test and diagnostic

software should be ready prior to the completion of the CTP design. The milestones for the full LVL1 trigger are summarized in Section 23.4.

The construction of the CTP will follow the usual rules, namely:

- the PCB(s) will undergo electrical tests before assembly;
- after assembly the CTP board(s) will be ‘burnt in’ by being left under power without cooling for two days;
- the fully-assembled board(s) will be tested.

A serial number will be given to each CTP board (i.e. the one or two boards that comprise the complete CTP and any spare boards) and a database will be set up to store information relating to each board. Such information will include the origin of the components that populate the board, the results of tests, a history of failure and the location of the board.

In order to facilitate the test and the maintenance of the boards, the JTAG test bus will be used as much as possible.

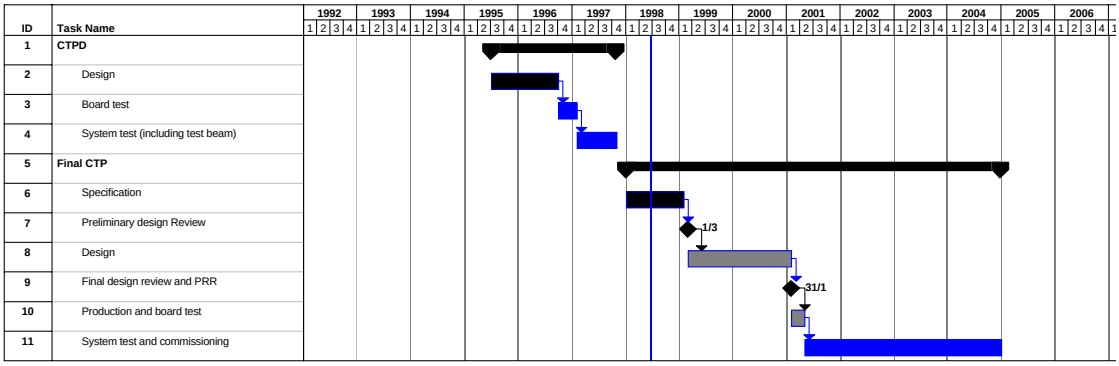

Figure 15-11 CTP project schedule.

15.8 Quality Assurance

The CTP project will be run under the ISO9001 quality management system. The design and specification of the CTP will be subject to at least two review exercises — a Preliminary Design Review (PDR) and a Final Design Review (FDR) — with the possible addition of Interim Design Reviews (IDR). All of these reviews will be conducted by the CTP design team and invited experts from the systems which interface with the CTP.

The PDR will have the following goals: to examine and assess the full requirements and specifications for the CTP, to identify any missing functionality, to ensure full compatibility with all connecting systems and to determine the overall feasibility. Detailed written specifications will be supplied to the review group in advance of the review itself.

IDRs may be held at any time during the design phase, but most usefully at the completion of schematic capture when many engineering issues (timing margins, interfaces to other systems, detailed latency calculations, etc.) can be explored.

The FDR will be held before the CTP design is sent for manufacture. It is intended to catch any design or engineering errors before budgetary commitment. This review will necessarily be of a more technical nature than the initial review, but there should be few problems to detect by this stage. It will be merged with the ATLAS production readiness review (PRR).

A user requirement document (URD) has already been written for the CTP [15-3] and this has undergone an initial review by representatives of each interfacing subsystem.

15.9 References

15-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirements Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

15-2 *ATLAS level-2 trigger: user requirements document*, DAQ–NO–79, March 1998.
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL2/DOCUMENTS/urdv1_2up.ps

15-3 *LVL1 Central Trigger Processor: (User) Requirements Document (version 1.1)*, ATLAS working document, ATL–DA–ES–0003, February 1998.

15-4 *Central Trigger Processor Specifications (Draft 0.1)*, ATLAS working document, March 1998.
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/ctpttc/CTPspec980310.ps

15-5 *ATLAS DAQ Back-end software User Requirements Document (Draft 3)*, ATLAS working document, May 1996.
http://atddoc.cern.ch/Atlas/DaqSoft/document/draft.ps

15-6 G. Schuler, *Central Trigger Processor Demonstrator (CTPD) specification*, RD27 working document, August 1996.
http://hepwww.rl.ac.uk/atlas/L1/1997_slice/ctpd.ps

15-7 I Brawn et al., *A Demonstrator for the ATLAS Level-1 Central Trigger Processor*, CERN / LHCC / 97–60, Third Workshop on Electronics for LHC Experiments, London, UK, September 1997.

15-8 I Brawn et al., *A Demonstrator for the ATLAS Level-1 Central Trigger Processor*, ATLAS internal document, March 1998.
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/ctpttc/CTPD980309.ps

16 Timing, trigger and control

16.1 Introduction

The ATLAS readout elements, such as the front-end electronics, the readout drivers (ROD) and possibly the readout buffers (ROB), need the bunch-crossing signal (BC clock) and the level-1 accept signal (L1A). As specified in Ref. [16-1], each readout element has to produce an event identifier (EVID) and a bunch-crossing identifier (BCID) with each event. It is therefore necessary to make available some synchronization signals, in order to maintain coherence between these identifiers across the experiment. These synchronization signals are the bunch counter reset (BCR) and the event counter reset (ECR). During test and calibration periods, the readout electronics also need to receive test and calibration signals.

The Timing, Trigger and Control (TTC) system allows the timing and trigger signals to be distributed to the readout electronics. The timing signals comprise the LHC clock (BC clock) and the synchronization signals (BCR, ECR). The trigger signals include the L1A, test and calibration triggers. The TTC allows the timing of these signals to be adjusted. The requirements on the TTC are defined in Ref. [16-2].

The ATLAS TTC system is based on the optical fan-out system developed within the framework of RD12 [16-3], which allows signals to be distributed from one source to up to 1024 destinations. The system can be partitioned and subdetectors can be run with the central ATLAS timing and trigger signals, or independently with their own specific timing and trigger signals. The TTC system receives the LHC 40 MHz clock (BC clock) and the ORBIT signal from the LHC, the L1A signal from the central trigger processor (CTP), and commands and data from the CTP and subdetector-specific electronics. Encoding allows this information to be transmitted on a single optical link which is fanned out to a maximum of 1024 destinations. At the receiving end, an ASIC decodes the incoming signal and makes available the BC clock, the L1A signal, the ECR and BCR signals, the EVID and BCID numbers, and the user commands and data. Provision is made to adjust the timing of all the signals. The context diagram of a partition is shown in Figure 16-1.

Some of the subdetectors (e.g. the liquid-argon calorimeters) will use the TTC system to bring signals directly to the on-detector front-end electronics; others (e.g. the inner tracker) will use the TTC system only as far as the RODs, at which point they will change the functionality and the protocol of the signal transmission. Figure 16-2 summarizes the different configurations.

The way a subdetector will use the TTC system will depend on its specific requirements. Most of the subsystems will use more than one partition to allow the concurrent running of different parts of the subdetector in different trigger modes during commissioning or calibration periods. The TTC team will work closely with the subdetector communities in order to define the optimum number of partitions for each subsystem, and to determine how the TTC should be used to set up the timing.

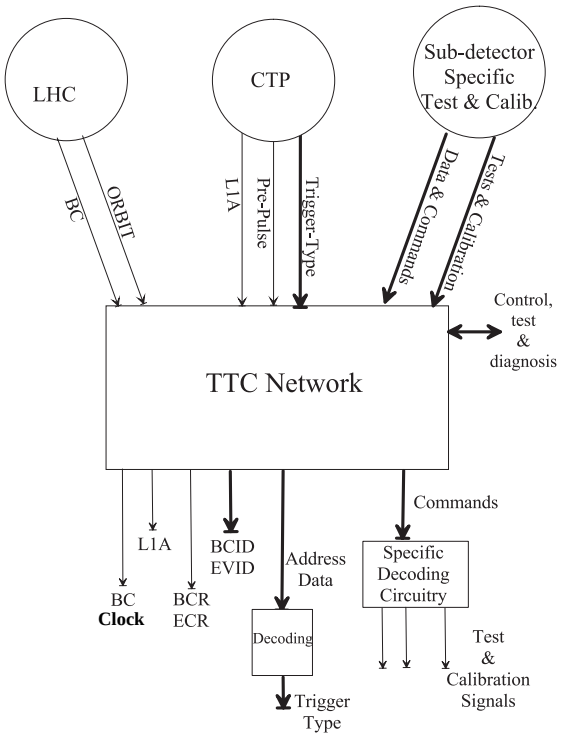

Figure 16-1 TTC context diagram.

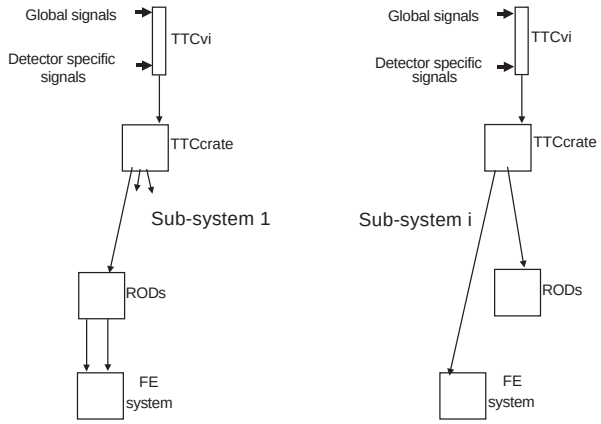

Figure 16-2 TTC configurations.

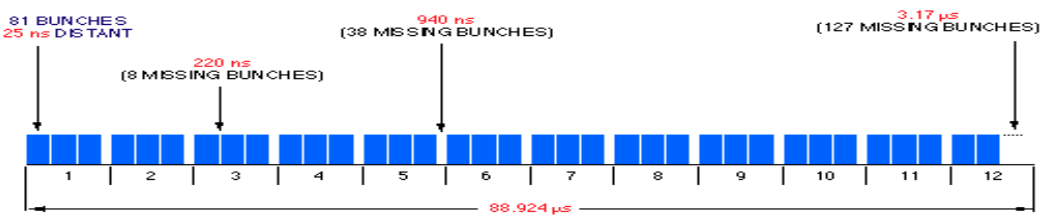

Figure 16-3 LHC bunch structure.

In ATLAS, the TTC system is used in different ways:

- In normal running, each TTC partition receives its clock from the LHC and the L1A from the CTP. The BCR is derived from the LHC ORBIT signal. After each L1A, an 8-bit trigger-type word is forwarded to the destinations as well as (optionally) a 24-bit event ID. The trigger-type word is formed in the CTP and contains information on what gave rise to an L1A, while the 24-bit event ID is formed in the TTC VME interface (TTCvi). The TTC system can also transmit specific subdetector data and commands without introducing deadtime, e.g. test pulses when there are no bunches (LHC gap), and front-end parameters (e.g. delay values).
- During commissioning and for test and calibration runs, triggers can be injected locally in each TTC partition.

The specifications of the different components of the backbone system are given in Sections 16.2–16.4. A description of the functionality of the system is given in Section 16.5.

16.2 Interfaces

The TTC system has interfaces to the following: the LHC machine, the ATLAS CTP, the specific test and calibration logic of the subdetectors, and the front-end and readout electronics of the subdetectors (see Figure 16-1). These interfaces are described below.

16.2.1 Interface to the LHC machine

The LHC provides a continuously running clock signal (BC clock), although there are empty bunches as shown in Figure 16-3. The LHC also delivers an ORBIT signal at a fixed time within the LHC cycle allowing synchronization within this cycle. The BCR signal is derived from this ORBIT signal.

Both BC and ORBIT signals are fanned out electrically to the TTC partitions.

16.2.2 Interface to the central trigger processor

The CTP provides to the TTC:

- the L1A signal which has to be broadcast to all of the readout elements;
- a prepulse (PPS) signal that can be issued a predefined number of clock cycles before a test trigger is generated. This signal can be used to initiate the transmission of a test pulse to the front-end electronics at the correct time with respect to the arrival time of the test L1A;
- an 8-bit trigger-type word with each L1A. This word is transmitted to the readout elements and allows different types of events (e.g. physics events, test events, calibration events) to be distinguished.

All of these signals are electrically fanned out from the CTP crate to the TTC partitions.

16.2.3 Interface to subdetector test and calibration electronics

The TTC system receives from the subdetector-specific electronics, timing and trigger signals which can be used during commissioning, test and calibration periods. These signals are transmitted to the front-end and readout electronics as TTC commands.

16.2.4 Interface to the readout electronics

The TTC system provides to the readout electronics the following signals or data:

- the BC clock;
- the L1A signal;
- the BCR and ECR signals;
- the BCID and EVID words with each L1A;
- the trigger-type word with each L1A;
- commands and data, including test and calibration pulses.

The phase of the BC clock is adjustable in steps of 100 ps within 25 ns. L1A, BCR and commands can in addition be delayed by up to 15 BC clock cycles.

16.3 TTC backbone components

The TTC backbone components are:

- the TTC crate and fibre network [16-4];
- the TTC VME interface (TTCvi) [16-5];
- the TTC receiver chip (TTCrx) [16-6].

The TTC crate receives electrical signals from the TTCvi and the LHC machine, and performs the encoding and the electrical-to-optical conversion. An optical tree network with optical passive fan-outs distributes the encoded optical signals to up to 1024 destination nodes. These nodes consist of an optical-to-electrical conversion device followed by a receiver ASIC (TTCrx) which decodes the incoming frame and makes available all of the timing, trigger and control signals. The TTCvi is a VME module which provides the trigger and control signals to the TTC crate in the form of two encoded signals (A-channel and B-channel). This module allows the user to select the trigger source and to generate commands at known times. A TTC

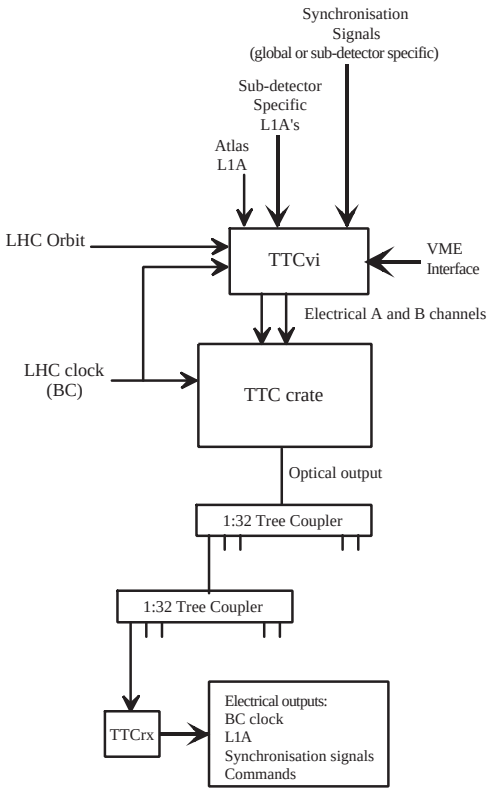

Figure 16-4 A TTC partition.

partition consists of a TTCvi, a TTC crate, an optical network and as many receivers as necessary. An example is shown in Figure 16-4.

16.3.1 The TTC crate

The TTC crate receives two electrical signals from the TTCvi and performs the final encoding and the electrical-to-optical conversion. The TTC crate includes a high-power laser transmitter incorporating a 1310 nm laser head module. At this wavelength the chromatic dispersion of normal optical fibre is negligible.

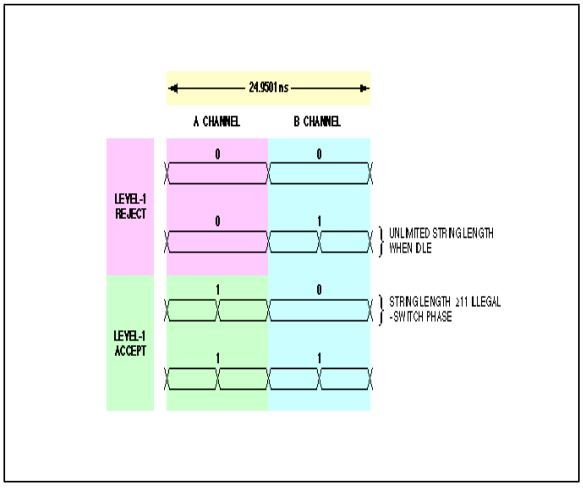

Figure 16-5 TTC two-channel encoding.

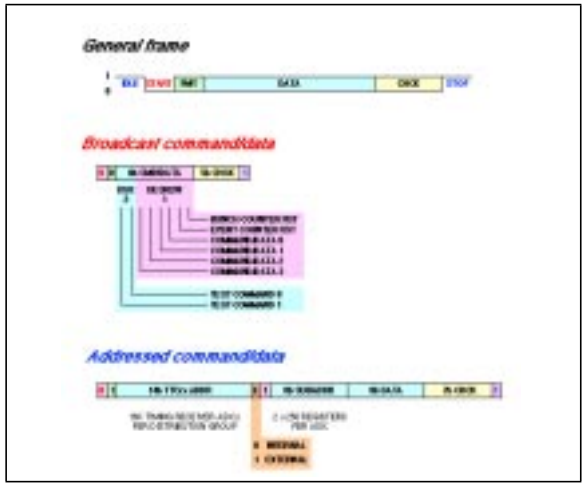

Figure 16-6 TTC encoding format.

The crate incorporates a 160.32 MBd biphase mark encoder which time-division multiplexes the two input signals (A-channel and B-channel, Figure 16-5) using a balanced d.c.-free code. The primary phase-locked loop (PLL) plus encoder jitter when transmitting data is about 7 ps r.m.s.. A narrow-bandwidth PLL with a low-noise voltage controlled oscillator (VCXO) having low subharmonic feedthrough allows an r.m.s. output jitter of less than 10 ps to be maintained with an input clock reference jitter of several hundred picoseconds. The low-latency A-channel is dedicated to the transmission of the L1A signal. The B-channel transmits framed and formatted broadcast and individually-addressed commands and data using forward error-correction for high reliability. The format supports up to 16,000 receivers per distribution zone (partition), each of which may be associated with up to 256 internal and external subaddresses. Figure 16-6 gives the encoding format.

16.3.2 The TTC VME interface (TTCvi)

The complete specification of the TTCvi module is available in Ref. [16-5]. A simplified block diagram of this module is given in Figure 16-7.

The TTCvi delivers the A-channel and B-channel signals to the TTC transmitter crate. These two signals carry timing, trigger and control information.

The TTC A-channel is used only to transmit the L1A signal.

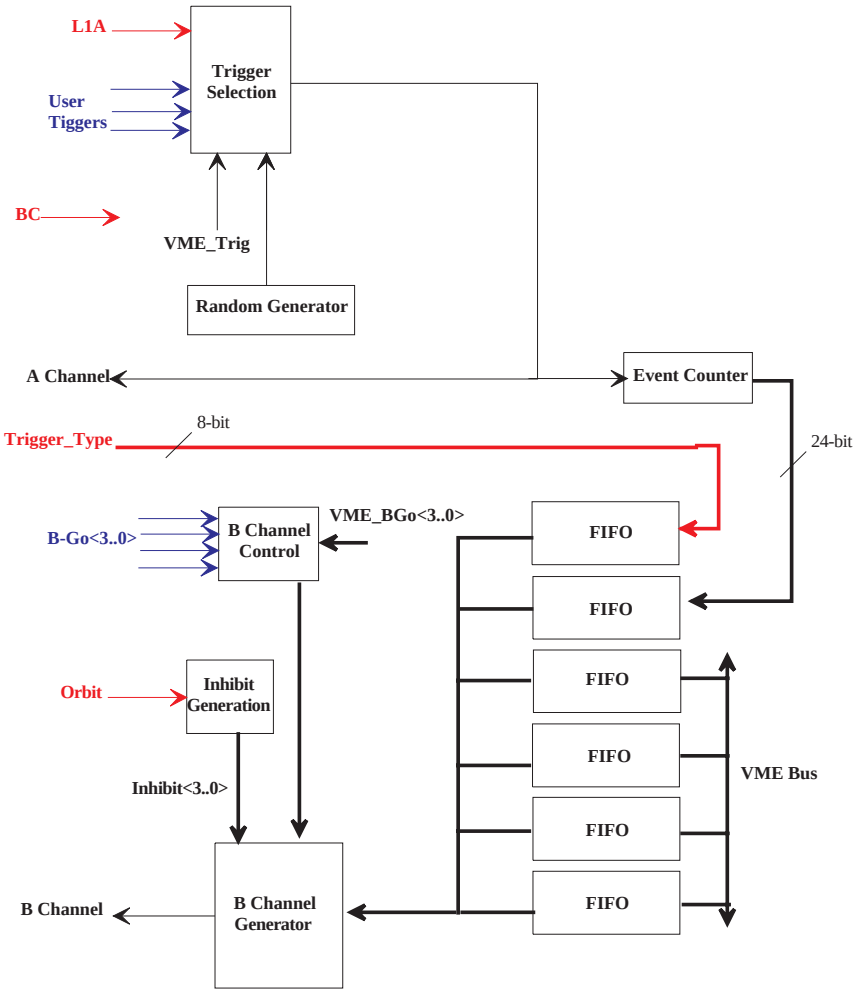

Figure 16-7 TTCvi simplified block diagram.

The TTC B-channel is used to transmit framed and formatted commands and data. These can be either:

- Short-format synchronous or asynchronous broadcast command / data cycles. If synchronous, the timing of these cycles relative to the LHC orbit is controlled precisely. Such cycles are used for the broadcasting of the BCR signal and for the transmission of other fast synchronous broadcast controls and test commands. These can be decoded at the receiving end (after the TTCrx) and used to produce test pulses, calibration pulses, etc.
- Long-format asynchronous individually-addressed or broadcast command / data cycles. The timing of these cycles with respect to the LHC orbit is indeterminate and they are not individually deskewed in the TTCrx ASICs. They are used for the transmission of parameters and non-time-critical commands. For instance, they are used to transmit the trigger type after each L1A signal.

The TTCvi is the master of a TTC partition. It is fully controlled through a VME interface by the partition user. It is at this level that the user decides to include this partition in the global ATLAS readout system or to run it in independent test or calibration runs. In this last case it uses subdetector-specific timing and trigger signals. When the partition is in the global ATLAS readout system the TTCvi is controlled by the DAQ run control.

16.3.2.1 Trigger inputs

Although in normal running mode the trigger input to the module is the L1A signal provided by the CTP, the TTCvi allows other trigger sources to be selected for test or calibration purposes without modifying the cabling. The main L1A sources are:

- The CTP L1A. The latency introduced by the module on the L1A coming from the CTP is minimized and is about 3 ns.
- Three additional local L1A inputs for subdetector-specific purposes.
- An internal random trigger generator. The number of L1A signals per unit of time follows a Poisson distribution with a mean rate programmable from about 1 Hz to 160 kHz.

16.3.2.2 ORBIT input

The ORBIT signal is a square wave of period 88.924 s which is received from the LHC machine and used for the generation of signals which are synchronized to the LHC orbit. Adjustment of the phase of the ORBIT signal permits a global control of the timing of the entire TTC system relative to the LHC bunch structure.

16.3.2.3 INHIBIT signals

In order to be able to send B-channel signals at known times within the LHC cycle, four independently-programmable timing signals called INHIBIT are generated within the TTCvi module. Each INHIBIT signal starts a programmable number of clock cycles after the start of the ORBIT signal and has a duration of a programmable number of clock cycles. Transmission of the associated synchronous command commences at the end of the INHIBIT signal duration.

One INHIBIT is used to initiate the transmission, during the long gap in the LHC cycle, of a broadcast command containing the BCR. The three other INHIBIT signals are available for the generation of other synchronous commands.

16.3.2.4 Generation of B-channel cycles

The TTCvi permits synchronous and asynchronous short- and long-format B-channel command / data cycles to be generated in a number of different ways:

Short- and long-format asynchronous cycles.
Asynchronous cycles may be initiated by writing the required data (a single byte for short-format or two 16-bit words for long-format) to specified TTCvi VME addresses. Normally, short-format cycles are used for broadcast commands or data, while long-format cycles are used for individually-addressed commands or data. However, a broadcast of 16 bits of data can be made with long-format cycles if TTCrx address zero is chosen. The timing of these cycles is not synchronized with the LHC orbit.

Preloaded synchronous or asynchronous cycles.
Four VME-addressable FIFOs are provided which may be preloaded with commands and data to be transmitted by B-channel cycles. For each of the four channels, the actual transmission of the preloaded information is initiated by a signal called B-Go. If synchronous mode is selected, the B-channel cycle is generated at the end of the corresponding INHIBIT signal. If

asynchronous mode is selected, the B-channel cycle is generated when the B-Go signal occurs. In addition, a special mode has been implemented to allow the download of a large amount of data (e.g. at the beginning of a run if the TTC system is used to download parameters). In this mode, the B-channel cycle is initiated as soon as the selected FIFO contains data.

Event-number and trigger-type cycle.
After each L1A is transmitted, the contents of a 24-bit event counter in the TTCvi can optionally be broadcast together with an 8-bit trigger-type word received from the CTP via a front panel connection. This broadcast is made asynchronously and takes about 4.4 μs if the B-channel is free. Internally, the event number and the trigger type are stored in FIFOs to avoid any losses due to the random time of arrival of L1A.

16.3.3 TTC receiver chip (TTCrx)

The TTCrx (Figure 16-8) is a custom ASIC that receives control and synchronization information from the central TTC system (after optical-to-electrical conversion) and makes it available to the readout electronics. The TTCrx can be programmed to compensate for particle times of flight and for propagation delays associated with the detectors and their electronics.

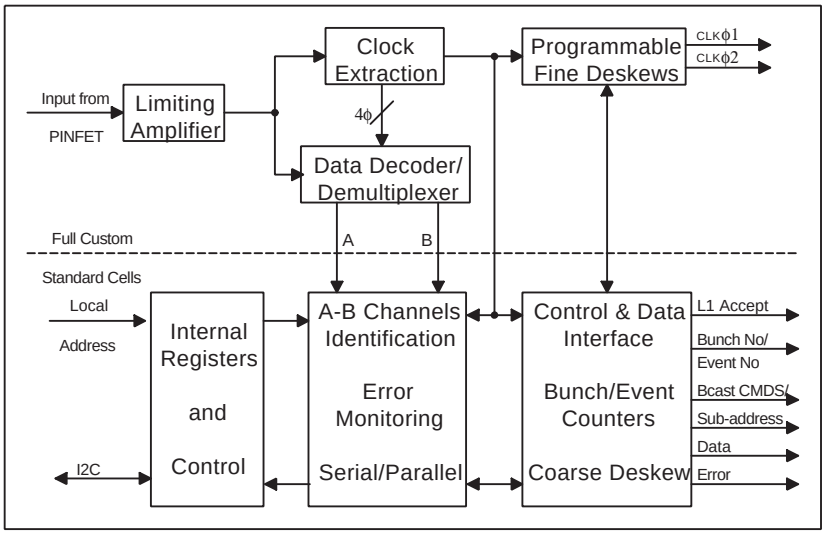

Figure 16-8 TTCrx diagram.

One of the main functions of the TTCrx is to recover and distribute the 40.08 MHz BC clock with minimal jitter. It also makes available to the readout electronics the level-1 trigger-accept decision (L1A signal) and its associated BCID and EVID numbers.

Each TTCrx IC is identified in the distribution network by a unique 14-bit channel-ID number. The ASIC control logic identifies the A- and B-channels, deserializes the data in the B-channel and continuously monitors it to look for the presence of its channel-ID number.

The TTCrx IC has already been prototyped in a standard 1 μm CMOS digital process. The ASIC footprint is 20 mm^2 and has been packaged in a 100-pin BGA package. A new version is being developed in a radiation-hard technology. A prototype should be available in spring 1999.

16.3.3.1 TTCrx characteristics

Bunch crossing
The BC clock is recovered with minimum jitter of the order of 40 ps r.m.s.. The recovered clock is then fed to a programmable fine deskew unit where two different clock phases are generated. The phases of the two clocks can be controlled independently in steps of 104 ps between 0 and 25 ns.

Level-1 accept, bunch-crossing identifier and event identifier
The TTCrx decodes the input frame and extracts the L1A signal with minimal latency (three BC clock cycles). With each L1A, the TTCrx delivers a 12-bit BCID number and a 24-bit EVID number on a 12-bit bus. Hence, three clock cycles are necessary for the readout electronics to get these numbers. The L1A signal can be delayed in the TTCrx by up to 15 BC clock cycles.

Bunch counter reset and event counter reset
The BCR and ECR are decoded in the chip and made available. The BCR signal can be delayed in the TTCrx by up to 15 BC clock cycles.

Broadcast commands
The other broadcast-command bits are made available as an 8-bit word. External decoding of these commands is necessary (for instance to generate a test pulse). The command word can be delayed by up to 16 BC clock cycles.

Set-up
There are four registers used to define the following delays in the TTCrx:

- fine deskew for BC clock 1 (0 to 25 ns with 104 ps step);
- fine deskew for BC clock 2 (0 to 25 ns with 104 ps step);
- L1A coarse deskew (0 to 15 BC clock cycles);
- command word (including BCR and ECR) coarse deskew (0 to 15 BC clock cycles).

These delays can be programmed via commands on the B-channel of the TTC network or through a local I^2C interface for the new version.

Error detection and correction
Provision is made for error detection on the incoming frame, and internal error registers are incremented each time an error is detected. These error registers can be accessed through the I^2C port.

16.4 Partitioning

As mentioned in Section 16.1, a TTC partition consists of a TTCvi, a TTC crate, an optical network and as many TTCrx chips as necessary (up to 1024). Any signal, data or commands transmitted on a particular partition reach only the TTCrx chips belonging to that partition. Partitioning the TTC system allows different subdetectors to work concurrently in the following instances:

- during test or calibration periods with independent timing and trigger signals;
- during the setting up of a physics run when the TTC network may be used to download parameters in the front-end and readout electronics.

Not only is it necessary to have independent partitions per subdetector, but also to have further subpartitions within each subdetector. This is for the following reasons:

- During assembly and commissioning, a subdetector is divided into independent parts.
- One could wish to run different parts of a subdetector in different modes (e.g. test mode and calibration mode).
- The fibre length from the TTC crate to the readout electronics may vary a lot from one part of the detector to the other and hence the delay compensation foreseen in the TTCrx chip may not be sufficient.

The exact number of subpartitions needed per subdetector is still being discussed. Table 16-1 shows the numbers according to present assumptions. This model leads to a total of 35 TTC partitions.

Table 16-1 Number of partitions per subdetector.

Subdetector	Number of partitions	Remarks
Pixels	2	
SCT	4	Barrel right & left, end-cap right & left
TRT	4	Barrel right & left, end-cap right & left
EM liquid-argon calorimeter	4	Barrel right & left, end-cap right & left
Hadronic end-cap calorimeter	2	End-cap right & left
Forward calorimeter	2	End-cap right & left
Tile calorimeter	4	Barrel right & left, extended barrel right & left
TGC	2	End-cap right & left
RPC	2	Barrel right & left
MDT	4	Barrel right & left, end-cap right & left
CSC	2	End-cap right & left
Level-1 calorimeter trigger	1	
Level-1 muon trigger	1	
Level-2 / DAQ	1	

It has to be noted that the current model assumes a modest number of partitions per subdetector (of the order of four or fewer). If a subdetector wants more partitions, then the model may not be cost-effective and a specific solution will have to be found. One could, for instance, think of adding an optical switch in between the TTC crates and the distribution network, as shown in Figure 16-9.

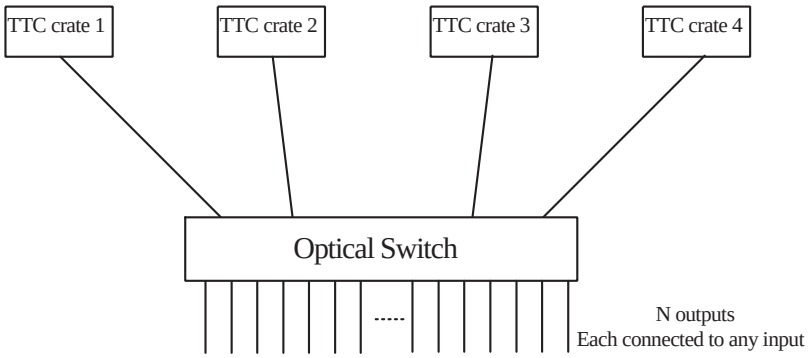

Figure 16-9 Possible use of an optical switch to get a high number of TTC partitions.

16.5 Functionality

This paragraph describes how the TTC system will be used, what functionality is available and how it is controlled. It does not touch upon the question of how the timing is set up; this is addressed in Chapter 19.

Figure 16-10 shows the different delay adjustments and their location in the TTC system.

16.5.1 Initialization

At the start of a run or a test or calibration period, one has to:

- define whether a particular TTC partition is to run independently or is to be part of the global ATLAS data-taking;
- select the signals to be used accordingly (clock, trigger source, etc.);
- configure necessary parameters in the TTCvi driving the TTC partition;
- set all the programmable delays in the TTCrx chips.

The first three operations are done through the VME interface of the TTCvi while the fourth one can be done either through the TTC network or through the I^2C interface of the TTCrx chips. This last method is more convenient when the TTCrx chips are located at the ROD level. The TTCrx registers are loaded in the same way as the other registers of the readout chain (and can be read back and checked).

The initialization is performed within the framework of the DAQ/run-control system (and not the DCS) as, from the DAQ viewpoint, there is a direct correlation between a TTC partition and a DAQ partition.

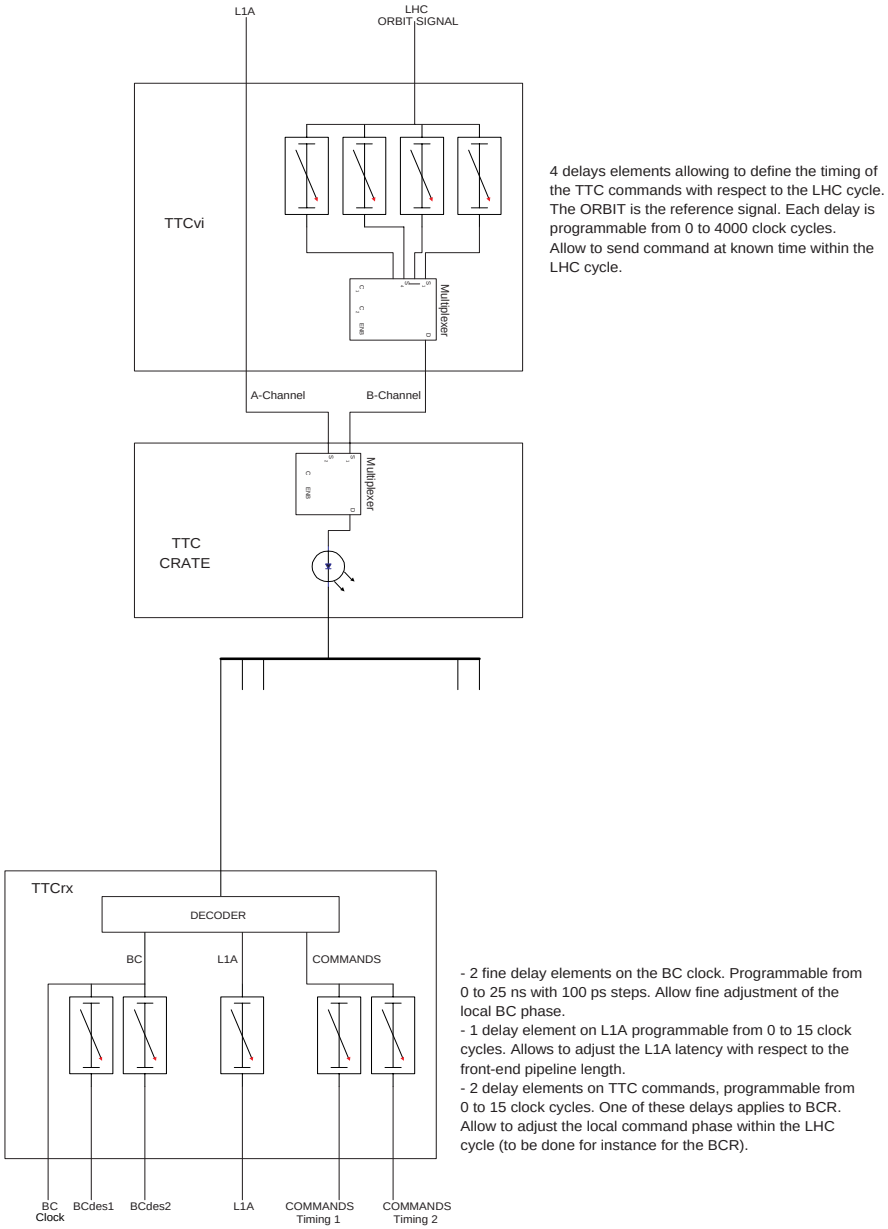

Figure 16-10 Different delay elements in the TTC system.

16.5.2 Synchronization

With regard to synchronization, one has to consider the following:

- synchronization with respect to the LHC cycle;
- maintaining a consistent BCID across the complete experiment;
- maintaining a consistent EVID across the complete experiment.

16.5.2.1 Synchronization within the LHC cycle

The LHC cycle is shown in Figure 16-3. The LHC machine provides a synchronization signal (ORBIT) which exhibits a transition edge at a known time in the cycle. This signal is used by the TTCvi to build up to four internal synchronization signals (INHIBIT). These signals allow the generation of commands (including BCR) at known times within the LHC cycle.

16.5.2.2 BCID synchronization

All the necessary tools to maintain a coherent BCID across the experiment are available. In the TTCvi, one of the INHIBIT signals is devoted to the BCR command. The phase of this command with respect to the LHC cycle is properly adjusted in the TTCvi. At the receiving end (TTCrx) one can locally adjust the phase of the BCR signal in the range 0–400 ns. This allows a misalignment of up to ±200 ns between the TTCrx's of a given partition; this is deemed to be enough to cover differences in cable or fibre lengths or differences in time of flight in the detector. BCR will be issued at the end of each LHC cycle in order to reset the bunch-crossing counters in the TTCrx. BCID 0 will be allocated to the last BC of the 3.17 μs gap at the end of the LHC cycle so that the LHC cycle starts with BCID 1. This choice may be subject to change if it gives rise to problems.

16.5.2.3 EVID synchronization

The ECR signal is necessary to reset the local event number counters if necessary (e.g. if a de-synchronization has been detected). The system must guarantee that there is no L1A during the time this signal is broadcast to all the TTCrx's. This can easily be achieved by introducing dead-time at the CTP level.

16.5.3 Test and calibration signals transmission

During test and calibration periods one has to be able to send pulses and triggers to the readout electronics. The TTCvi allows this to be done in different ways, without requiring any hardware intervention. In addition to the ATLAS L1A signal coming from the CTP, three other trigger inputs are available. The test signals can be transmitted as synchronous commands on the TTC and any one of the available B-Go signals can be used (in this case the relative phase between the test pulse and the trigger is to be adjusted externally), or sequences of commands with timing defined by the internal INHIBIT signals can be used. In this latter case the user has to decode the TTCrx command outputs to locally generate test and trigger signals. The INHIBIT signals are available as NIM signals on the TTCvi front panel and can be used to fire external logic if needed. The command output of the TTCrx can be resynchronized with one of the two deskewed clocks available (Bcdesk1 and Bcdesk2). This allows the readout electronics to have, for instance, a test signal with an adjustable timing with respect to the clock.

The prepulse (PPS) signal that can be issued by the CTP a known time before a random L1A is asserted, will be connected to one of the B-Go inputs and used to transmit a test pulse to the front-end electronics.

16.5.4 Trigger type transmission

It is necessary to have available, with each L1A, some information on the type of the trigger (for example, calibration, physics). This information may be used in the RODs. The TTCvi receives an 8-bit trigger-type word from the CTP and transmits it to the TTCrx. This transmission cannot be in phase with L1A, as there is not enough bandwidth available to transmit this word within five BC (the BC of the L1A followed by the four-BC dead time). The trigger type is transmitted in an asynchronous way when the TTC system is available (i.e. when there is no transmission of synchronous commands such as BCR). The trigger type is available at the output of the TTCrx 1–2 μs after the L1A has been issued, assuming there were no L1As in the previous 2 μs. If consecutive L1As occur within a short period of time, the trigger type transmission is queued and the trigger type of a given L1A will become available later at the output of the TTCrx. This is deemed to not be a problem as the trigger type is used only at the ROD level where the event data are subject to similar fluctuations in arrival time.

16.6 Latency

The TTC system introduces latency (delays) on the L1A signal at different places:

- The L1A signal has to be fanned out and transported from the CTP to the TTCvi.
- The L1A has to be transformed in the TTCvi to provide the A-channel signal to the TTC crate.
- The signal has to be transmitted (cable) from the TTCvi to the TTC crate and, in the TTC crate, it has to be encoded and converted from electrical to optical.
- The light has to travel through the optical network.
- At the receiving end, there is an optical-to-electrical conversion and the TTCrx has to decode the incoming frame and make available the L1A signal.

The time to transport the L1A signal from the CTP to the TTCvi has been estimated to be 50 ns, which includes a possible fan-out stage after the CTP and about 8 m of cables. This is based on the assumption that the TTCvi's are in the vicinity of the CTP. This scheme may not be optimal in terms of practicality as a lot of subdetector-specific electronics is connected to the TTCvi and subdetector groups may wish to have their TTC electronics close to their readout electronics.

Table 16-2 TTC latency.

Element	Latency (ns)
CTP to TTCvi	50
TTCvi	3
Cable to TTC crate	3
TTC crate	22
Fibre (≤ 80 m)	400
TTCrx	75
Total	**553 (23 BC)**
Total without cables and fibres	*100 (4 BC)*

The part of the latency due to the optical network depends on the fibre lengths and is very difficult to evaluate today. A worst-case length of 80 m has been taken. This corresponds to the maximum distance between any on-detector front-end electronics element and the furthest crate in the underground control room (USA15).

The final latency will have to be computed for each sub-system when the physical locations of the different elements are known. It should also include the latency introduced by a change of protocol on the line as is done, for instance, for the inner tracker.

The latency in the TTCvi-TTC crate has been optimized in three ways. First, there is no re-synchronization of the L1A signal with the BC clock used by the TTCvi; second, only three ECL stages are necessary for the TTCvi to produce the A-channel signal; and third, the phase of the BC clock used in the TTC crate is adjusted in such a way that its rising edge is properly set with respect to the L1A phase. The total latency of the TTC system is summarized in Table 16-2.

16.7 Test and monitoring

There is a need for a global timing adjustment of the experiment (Chapter 19) which has to be done regularly and at least before each run. This timing adjustment sequence is a very good test of the TTC network itself as it requires the transmission to be fully efficient and all the timing characteristics of the delivered signals to be good.

The monitoring of the system is done in the following ways:

- The information carried by the TTC is part of the event readout data (BCID number, EVID number, trigger-type word) and this information is checked at every stage of the readout chain. This allows the basic transmission on the TTC network to be monitored.
- The TTC includes an error detection mechanism and each TTCrx maintains an error register. These registers will be read out regularly.
- The timing characteristics of the delivered signals (jitter on the BC clock, delay stability, etc.) will require a fine analysis of the readout data.

The test of the TTC system is done via VME for the TTCvi and via the I^2C interface for the TTCrx.

16.8 Experience from prototyping

Each part of the TTC system has been prototyped and tested and some experience acquired. The system works, but there are area where performance needs improvement. The main problems discovered up to now concern the receiver chip, for which some nonlinearity in the delay elements and some sensitivity to temperature variations have been discovered. These problems have been solved either by a design modification or by the development of a special small ball-grid-array (BGA) package.

There is still a lack of experience in the use of the system in a readout environment, but this will be addressed during the years 1998 and 1999 as most of the subdetector groups are planning to use this system for beam tests.

16.9 Construction, assembly procedure and schedule

As mentioned in Section 16.8, the prototype of the full system has not yet been validated by extensive use in a real data-taking environment. For this reason the specifications of the different parts of the system are not yet finalized. It is expected that after two years of utilization in the test-beam environment, the final specifications will be issued and the final design will start in 2000. The production of the TTC crate and TTCvi will then be subcontracted to external firms as the total amount of modules or elements needed is high.

The TTCrx chips will be ordered centrally by CERN from the selected foundry and made available to the community.

A schedule for the TTC system production is shown in Figure 16-11.

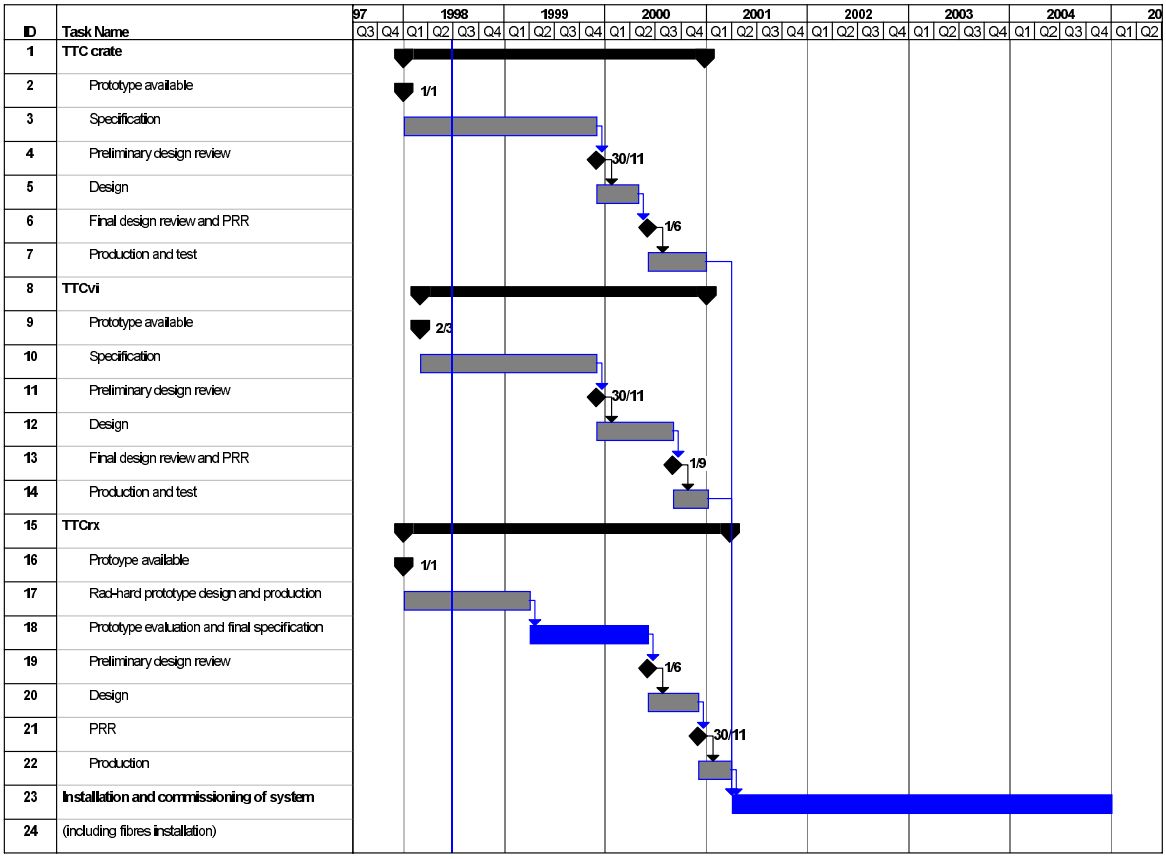

Figure 16-11 Schedule for the TTC system production.

16.10 Quality assurance and review procedures

A set of procedures and rules must be followed to assure the required quality of the system.

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

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

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

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

This review work has already been done for the prototyping. A User Requirement Document (URD) for the TTC system has been written and reviewed for a first time by representatives of each subsystem.

Prototypes of each element exist and will be used by the subsystems in the coming year. During that time, the final specifications will be written. The PDR will review them as well as the performance and functionality of the prototype system.

The construction of the TTC crates and of the TTCvi will follow usual rules, namely:

- Electrical tests of the PCB before assembly.
- Burn in of the boards after assembly by leaving them under power without cooling for two days.
- Test of the full boards.

Each element will receive a serial number and a database will be set up to store the information relative to each board: origin of the components which populate the board, results of tests, history of failure and location.

In order to facilitate the test and the maintenance of the boards, the JTAG test bus will be used as much as possible.

16.11 References

16-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

16-2 *Timing, Trigger and Control System (TTC) (User) Requirements Document (Draft version 0.2)*, ATLAS working document, ATL–DA–ES–0004, February 1998.

16-3 *RD12 Status Report*, CERN / LHCC / 97–29, April 1997.

16-4 TTC crate information, RD12 web page.
http://www.cern.ch/TTC/intro.html#Transmitters

16-5 *TTC-VMEbus Interface: TTCvi, Revision 1.5*, RD12 working document, November 1997.
http://www.cern.ch/TTC/TTCviSpec.pdf

16-6 *TTCrx Reference Manual, version 2.2*, RD12 working document, July 1997.
http://pcvlsi5.cern.ch:80/MicDig/ttc/MANUAL22.PDF

17 Computing

17.1 Introduction

The ATLAS Level-1 Trigger will require a large body of accompanying software at all phases of its design, development, installation, and running. We will need to build a well-documented and easily maintainable system that will last through the inevitable changes in personnel during the long running period of ATLAS. This chapter aims to give an overview of most aspects of our software and computing requirements, from module diagnostics to data acquisition and online monitoring, calibration and offline simulation and reconstruction. In many cases it summarizes experience and plans which have been presented in greater detail earlier in this document, in particular Sections 7.8, 8.6, and 11.2.

17.2 Physics simulation

Over the last few years we have performed wide-ranging and in-depth physics simulation of the performance of a number of proposed trigger designs. The methods used are described in the accompanying ATLAS Trigger Performance Status Report [17-1].

Initially, two levels of simulation for the calorimeter trigger were implemented: one level used smeared four-vectors with a only simplified detector model, but included the full time-history of each pulse. This was a fast simulation and was essential for studying effects of BCID, pile-up etc. The other level performed full GEANT tracking and digitization, but did not include the pulse history. This was more useful for understanding effects of resolution etc. We have also written some other stand-alone programs to examine certain requirements of the system, such as BCID, in greater detail. Recently, these approaches have been unified in a new program which has been used to produce many of the calorimeter simulation results presented in this document. This program will be developed further in the near future for continuing detailed design studies.

Similarly in the muon trigger, many early studies used simplified models of the detector and the physics. However, the acceptance of the muon trigger is much more complex, not least because of the highly nonuniform magnetic field, and so the motivation to use the full GEANT simulation was earlier and greater. A detailed description of the detector geometry and response of the muon trigger system has been simulated with the ATLAS Monte Carlo program DICE (GEANT based). The simulated data have been processed through a detailed simulation of the trigger logic, independently for the barrel and endcap subsystems.

Independent stand-alone simple simulation programs have been used to check the impact on physics performance of the more important physics effects of the muon trigger system. Also, the accidental trigger rate induced by the cavern background has been studied with dedicated numerical and Monte Carlo programs, to evaluate the robustness of the trigger system against possible severe environmental background conditions.

In the longer term, a suitable software model of the trigger will need to be implemented in the framework of the object-oriented offline software. This will be required both for future

simulations, for developing trigger menus, and in the reconstruction program for offline calibration and monitoring of the trigger performance.

17.3 Diagnostic, test and calibration software

During the prototyping phase, all components of the trigger will need to be tested. Each group will be providing its own test and diagnostic software for this purpose. There would be some benefits arising from a common approach, but the logistical difficulties involved may make this impossible. The experience of designing object-oriented software gained during the calorimeter trigger demonstrator programme may however be of use to the other groups.

The calibration procedures required by the various trigger components have been described elsewhere. It is envisaged that many of these will be implemented in software as test modes of the DAQ system.

17.4 DAQ software

We will need software running in a data acquisition framework on various timescales: from laboratory and beam tests of prototype modules to the full ATLAS environment. We are closely following the progress of the DAQ group ‘prototype –1’ project and hope also to learn from their experience in designing and implementing object-oriented software systems.

17.4.1 DAQ for prototype testing

Prototype modules will be tested in the laboratory and probably also in the ATLAS test beam. At this stage we would preferably use software integrated into the prototype ATLAS DAQ system to start gaining experience with that. However, it is possible that not all the DAQ components we require will be ready on the timescale we need, so we may have to keep evolving some of our existing DAQ software in the short term.

17.4.2 DAQ for final ATLAS system

Most of the basic DAQ infrastructure which we had to develop for the laboratory and test-beam environments will, in ATLAS, be provided by the DAQ group. We will still have to provide some software to run in or to control whatever processor(s) are chosen for the local crate controllers and RODs. We will also need to provide software, callable by the Run Control component of the ATLAS DAQ, to initialize the modules from their power-up state (e.g. load the trigger) and to prepare them for run start, etc.

For debugging and some calibrations, we will want to run the DAQ with the trigger as a stand-alone partition. Other calibrations will require both the trigger and either the muon or calorimeter detectors to be controlled together. We will need to provide the software to run in these modes, though the principle of partitioning the DAQ is envisaged as part of the ATLAS central DAQ system.

A vital feature of the online system will be real-time monitoring of the integrity of the trigger and the quality of the triggered data. We expect to monitor the system at various levels. Some monitoring may be provided in hardware, for example on the PPMs in the calorimeter trigger, which would give an untriggered ‘level-0’ view of what is happening in the detector. The next level would be implemented in the readout chain of each system to provide an immediate check of the performance of the electronics. An overall indication of the performance of the trigger system would be provided by online analysis of a fraction of complete events.

17.5 Detector Control System

We may need software to interface with the Detector Control System (DCS). We may want to report the status of our own crates on local displays and we will need to respond to certain DCS commands. We will also need to write the software for the local part of the DCS which is our responsibility. Details of the interaction with the DCS are still to be defined.

17.6 Platforms, languages, tools, and training

The trigger software will be run on a variety of platforms. DAQ, online monitoring, and some test and diagnostic will run in processors housed in VME (or other) crates. Parts of the diagnostic software may be run on PCs or workstations. Calibrations will be performed both online and offline on workstations or work-group servers. For this reason, it is desirable to avoid manufacturer-specific features in the DAQ code, and premature relationships with a single hardware manufacturer should be avoided. The choice of hardware platforms will be guided by ATLAS standards, if and when they emerge.

One of the supposed advantages of OO software is that it will allow for reuse of software components. There may be some merit in considering the design of all the trigger software as a package, with some classes used only online, others only offline, but some shared by online and offline software. If this is desirable, a high degree of portability will be needed. This constrains the choice of languages to (probably) just C++, but possibly including Java as an alternative.

The advantages of Java are that it is a cleaner, safer language than C++; it could allow a natural integration of some applications with Web-based documentation. This might be desirable especially in the ATLAS control room. However, Java can also be compiled and run as stand-alone programs with their own graphical user interface. Compiled Java is only moderately slower than C++.

There is now some OO experience within the level-1 calorimeter trigger group. However, given the ATLAS timescales, some turnover of personnel is inevitable. It is therefore necessary to consider the training of new people in both OO techniques and in the existing software. The latter requires some discipline in both design and documentation. Previous experience in HEP experiments suggests that this self-discipline is very hard to achieve. Even though we tried to approach the design of the calorimeter trigger diagnostics software ‘properly’, we did not in practice spend enough time in the analysis and design phase, and we had at least one serious redesign of the software part way through the implementation phase.

For the calorimeter trigger diagnostic software, we did not use any CASE tools. Use of such tools in the future is highly desirable, both as a design discipline and to provide design documentation to people joining the project later.

The offline component of our software will clearly be part of the 'ATLAS Software Process' (ASP), as described in the *ATLAS Software Technical Proposal* [17-2]. Since the more online aspects will also be developed by widely-scattered people, some similar approach to project management and quality assurance may be required.

17.7 Summary

We have a considerable software task ahead of us which will require various skills in online and offline software. With extra OO training and suitable CASE tools, we should be able to complete the project in the time available.

17.8 References

17-1 *ATLAS Trigger Performance Status Report*, CERN / LHCC / 98–15, July 1998.

17-2 *ATLAS Software Technical Proposal*. CERN / LHCC / 96–43, December 1996.
http://atlasinfo.cern.ch/Atlas/GROUPS/SOFTWARE/TDR/TDR.bk.ps or
http://atlasinfo.cern.ch/Atlas/GROUPS/SOFTWARE/TDR/html/TDR-1.html

18 Summary of level-1 trigger latency

Latency is one of the key parameters of the LVL1 trigger that affects the design of all ATLAS front-end systems. Latency is defined (loosely) as the time taken to form and distribute the LVL1 trigger decision, measured from the time of the proton–proton interaction at the centre of the ATLAS detector until the time when L1A trigger signal is received in the detector front-end systems. It determines the required length of the pipeline memories within the front-end systems.

The maximum acceptable latency of the ATLAS LVL1 trigger is specified in a document [18-1] that has been endorsed by the detector system project leaders. All front-end systems must be able to accommodate a latency value of up to 2.5 μs. The latency is measured from the time of collision of the protons (assuming each proton is at the centre of its respective bunch), until the L1A signal is available as an electrical pulse at the output of the TTCrx ASIC (Chapter 16) in the front-end electronics.

Any detector that wishes to have one or more of the following:

- system-specific TTC distribution systems used in addition to the standard TTC backbone;
- non-optimal (from the point of view of latency) rack organization for the TTC crates in the USA15 counting room;
- non-optimal (from the point of view of latency) cable routing between the TTC crates in the USA15 counting room and the front-end electronics;

must be prepared to cope with latency beyond 2.5 μs. This flexibility in the latency requirements specification is necessary to allow each detector the possibility to make an overall optimization of their own sub-system. The target latency for the LVL1 trigger is 2.0 μs, leaving 500 ns as contingency. As discussed in the following, the present best estimate for the LVL1 latency is close to this target.

For the success of ATLAS it is vital that the LVL1 trigger and the detector systems respect the latency requirements that have been specified. A careful study of the LVL1 trigger latency has therefore been performed, as summarized in Table 18-1. Details of the calculations can be found elsewhere in this TDR, in the chapters dealing with the muon trigger, the calorimeter trigger, the CTP and the TTC system. It should be noted that, where necessary in the system, for example at the input to the CTP, signals that arrive relatively early are delayed to bring them into temporal alignment with the corresponding input that arrives last.

Typically, data are transmitted between LVL1 trigger sub-systems serially (or semi-serially) on fast optical or electrical links. The latency for each stage in the processing is calculated from the time when the last bit of serial data is available at the input until the last serial bit is available at the input to the subsequent subsystem. Hence, the latency for each subsystem includes the time required to transmit the data to the subsequent sub-system.

The following elements are included in the latency calculations:

- Time of flight of particles from the interaction point to the detector elements.
- Response time of the detectors and their associated analogue electronics.
- Propagation delays along cables used to transport analogue signals up to and within the detector front-end systems that are mounted on the detector.

Table 18-1 Summary of LVL1 trigger latency in bunch-crossing (BC) units

Item	Contribution to latency (BC)		Comment
Muon trigger	**54.0**		
RPC-specific part		36.0	Including worst-case 80 m fibres to USA15
TGC-specific part		46.0	Including worst-case 80 m fibres to USA15
Interface to CTP		8.0	
Calorimeter trigger	**56.1**		
Signal processing and cables up to input of trigger preprocessor		20.6	Including worst-case 60 m cables to USA15
Preprocessor (e/γ, τ/h)		15.0	
Preprocessor (jet,E_T)		17.0	
Electron/gamma finding		14.0	
Tau/hadron finding		14.0	
Jet finding		18.0	
Missing E_T		18.5	
Total scalar E_T		18.5	
CTP	**4.0**		
TTC	**3.1**		Includes cable from CTP
Fibres to FE electronics	**16.0**		Using worst-case 80 m fibres
TTC receiver (in FE electronics)	**3.0**		
TOTAL	**82.2**		

- In the case of the muon trigger, propagation delays along cables used to transport digital data (patterns of hits) from the detector front-end electronics to the trigger electronics that are mounted on the detector.
- Propagation delays along cables used to transport analogue signals (calorimeter trigger) or digital data (muon trigger) from the on-detector electronics to the USA15 counting room.
- In the case of the calorimeter trigger (liquid-argon calorimeter inputs), delays introduced by detector analogue electronics that receive the trigger-tower signals in USA15, as well as propagation delays for transmitting the signals on to the trigger front-end preprocessor.
- Propagation decays associated with any patch panels required to re-order cable groups for input to the trigger or to facilitate installation and maintenance.
- Time required to perform digitization, where relevant allowing for the need to adjust the clock phase locally in order to sample the analogue pulse optimally (e.g. to sample the peak of calorimeter pulses).

- Propagation delays for digital values along internal cabling, and over backplane and circuit-board connections, within the LVL1 system.
- Delays in the pipelined processing electronics, determined by counting the number of pipeline steps, taking into account the relative phase of the clock signals to different stages in the processing chain.
- Time required to fan out and transmit the L1A signal to the TTC crates associated with the detector sub-systems. This is based on a realistic model for the rack layout of the LVL1 system and the TTC crates that has been optimized to minimize the latency. As discussed above, detector systems may decide to organize their TTC crates differently, at the expense of increased-latency requirements on their front-end systems.
- Delays in the TTC electronic system and the associated electrical-to-optical conversion.
- Propagation delays along the fibres and over other optical elements that connect the TTC crates to the detector front-end systems. The fibre length assumed for the calculation (see Table 18-1) is 80 m, which is a conservative estimate. In the case of the calorimeter electronics, which uses analogue pipelines with critical latency requirements, the fibres could follow the same path as the cables that bring the trigger-tower signals to USA15 (length 60 m or less). The shortest possible fibre path to the inner-detector electronics would be even less (approximately 45 m), but here a different routing will be used in practice.
- As discussed above, detector sub-system groups may decide to route their fibres differently, at the expense of increased-latency requirements on their front-end systems. This will be the case for the inner detector.
- Delays in receiving and making available the L1A signal in the detector front-end electronics, including optical-to-electrical conversion and delay in the TTCrx ASIC (see Chapter 16).

The cable lengths for signal transmission between the detector and USA15 and for the return path have been calculated by the ATLAS Technical Coordination group, taking into account the constraints imposed by installation and access scenarios.

Given the critical importance of the LVL1 latency, a thorough check will be made for the detailed final LVL1 trigger design as part of the Production Readiness Review. Similar checks will be made in the PRRs for the front-end electronics of other ATLAS systems.

18.1 References

18-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

19 Strategy for setting up the timing of the experiment

19.1 Introduction

Several hundred thousand readout elements receive timing (BC clock), trigger (L1A), synchronization (BCR) and test (Test-pulse) signals. The phases of all these signals have to be adjusted in order to sample the input signals optimally, read out the proper event and maintain a coherent BCID number and hence synchronization over the whole experiment. There will be a very large number of delay elements to be adjusted and monitored, most of which are in inaccessible places, and these two tasks may require a lot of time.

Maintaining synchronization across a big experiment is always an issue even when the time between bunch crossings is large and the trigger rate is low (for instance in LEP experiments). In the case of the LHC, where the bunch-crossing period is very short, the trigger rate is very high and the number of detector channels is very large, making sure that all of the subdetector systems read out the same bunch crossing when a trigger occurs is a major issue.

In view of the above, the timing set-up is seen as critical and some automatic procedures must be defined to set up the timing for various modes of operation:

- beam–beam collisions;
- beam-halo triggers;
- cosmic-ray triggers;
- running with test pulses for test or calibration purposes.

Using the TTC system, the following timings have to be adjusted:

- BC clock phase at the front-end level (to ensure optimal sampling of the input signal);
- L1A arrival time at the front-end level (to ensure readout of the correct bunch crossing);
- Bunch Counter Reset (BCR) arrival time (to ensure a coherent BCID across experiment);
- test-pulse phase (test or calibration signal produced at the right time relative to the clock and the trigger).

Each subsystem must define a procedure to make these adjustments and this work is in progress in ATLAS. In this chapter, possible strategies are presented and discussed in general terms. In due course, each subdetector group will produce their own procedure which will be documented and formally reviewed.

This chapter presents example procedures for the following:

- timing set-up using beam for two representative front-end configurations;
- using test-pulses to make the timing set-up for beam;
- timing set-up for calibration with test-pulses.

The procedures for setting up the timing for cosmic-ray triggers and for beam-halo triggers have not yet been worked out in detail.

19.2 Timing set-up with beam

This section describes a possible automatic way to set up and check the coarse timing (i.e. everything except the fine BC clock phase adjustment) of the detector with beam. This can be applied only when beam is available; the next section will show that it is also possible to set up the timing with test-pulses, which will be of considerable interest during the commissioning phase of the experiment.

It is planned to make use of the LHC bunch structure and of the BCID values provided by the front-end electronics (FE_BCID), or the level-1 trigger electronics (BCID) as specified in Ref. [19-1]. Two cases are considered: in the first one the FE_BCID is formed before the level-1 pipeline and in the second one FE_BCID is formed after the level-1 pipeline. These two cases cover the vast majority of the ATLAS front-end systems.

The fine BC clock phase adjustment procedure is subdetector-dependent and can be done either before or after the coarse timing has been done (e.g. by looking at distributions of pulse-peak position or number of hits versus time). In the latter case, an iteration on the coarse alignment procedure might be necessary.

19.2.1 FE_BCID formed before the level-1 pipeline

Figure 19-1 shows a simplified block diagram of the front-end electronics circuitry when the FE_BCID value is formed before the LVL1 pipeline.

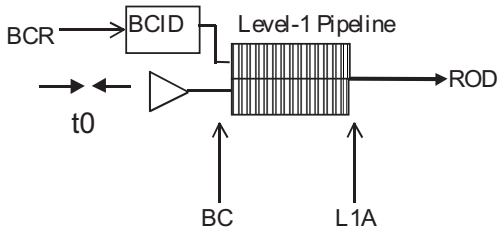

Figure 19-1 Block diagram of the front-end element with the BCID formed at the input of the level-1 pipeline.

The timing and trigger signals are transmitted to geographical regions of the subdetectors which house a large number of channels, and within which there is no need for fine timing adjustment. For instance in the TRT, such a region is 1/32 of a wheel in the end-cap, which corresponds to 200–400 channels depending on the type of wheel; in the liquid-argon calorimeter a region corresponds to a front-end board which houses 128 channels.

Figure 19-2 shows a flow chart of the timing adjustment procedure. This procedure is organized in two steps.

In the first step, random L1A signals are generated and data are read out. For each timing geographical region, a histogram is built of the number of hits (in the sense of the presence of a signal above threshold) versus the FE_BCID received from the front-end. This histogram should reflect the LHC bunch structure, but be shifted in time as shown in Figure 19-3. The BCR timing is then adjusted in order to get FE_BCID = 1 at the correct place.

In the second step, the system is run with L1A fixed at BCID = 1. The event readout will generally give a FE_BCID value not equal to 1 since the L1A timing has not yet been adjusted. The L1A timing is then adjusted in order to get the correct FE_BCID value.

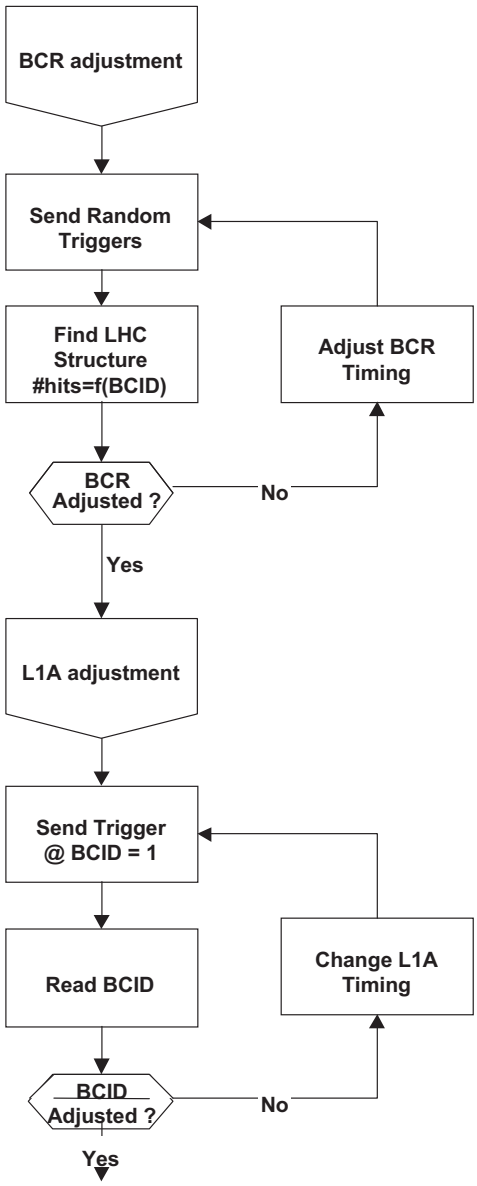

Figure 19-2 Simplified flow chart of the timing set-up with beam when FE_BCID is formed before the level-1 pipeline.

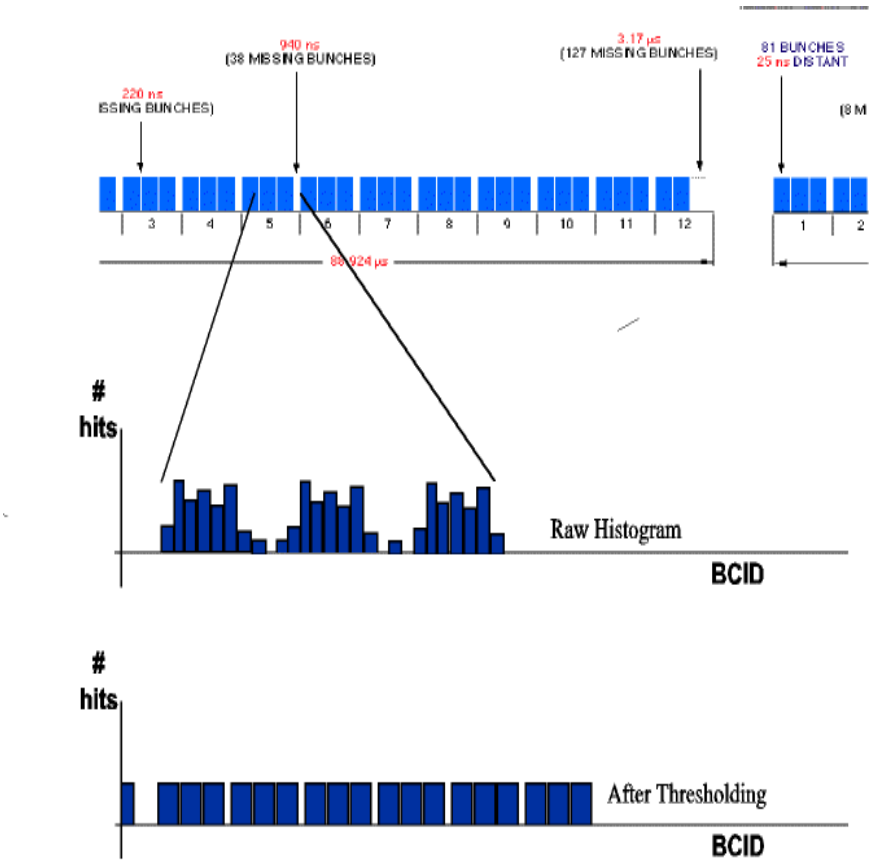

Figure 19-3 Hits versus BCID.

19.2.2 FE_BCID formed after the level-1 pipeline

In some front-end systems, the BCID number is formed after the Level-1 pipeline, when the L1A occurs (as shown in Figure 19-4). In this case, if one has bad timing of the L1A and BCR signals, the FE_BCID value will be wrong and equal to FE_BCIDtrue+ErrorL1A+ErrorBCR. Here the procedure described above in Section 19.2.1 has to be modified slightly, as shown in Figure 19-5.

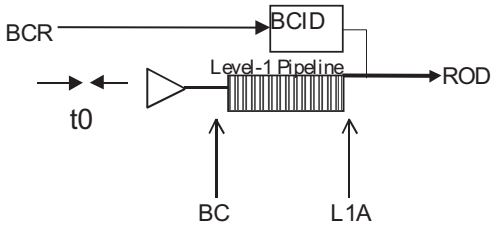

Figure 19-4 Block diagram of the front-end element with the BCID formed at the output of the level-1 pipeline.

In the first step, L1A signals are generated for each value of BCID and data are read out. For each geographical region, the histogram of the number of hits (sum of all the channels in the zone) versus the expected BCID number is built. This histogram should reflect the LHC bunch structure, but be shifted in time as shown in Figure 19-3. This histogram has to be shifted in one or other direction by a number of bunch crossings to coincide with the LHC bunch structure. The L1A signal delay is adjusted by this number of clock cycles. It has to be noted that until the fine BC clock phase adjustment has been done, an error of plus or minus one BC can be present.

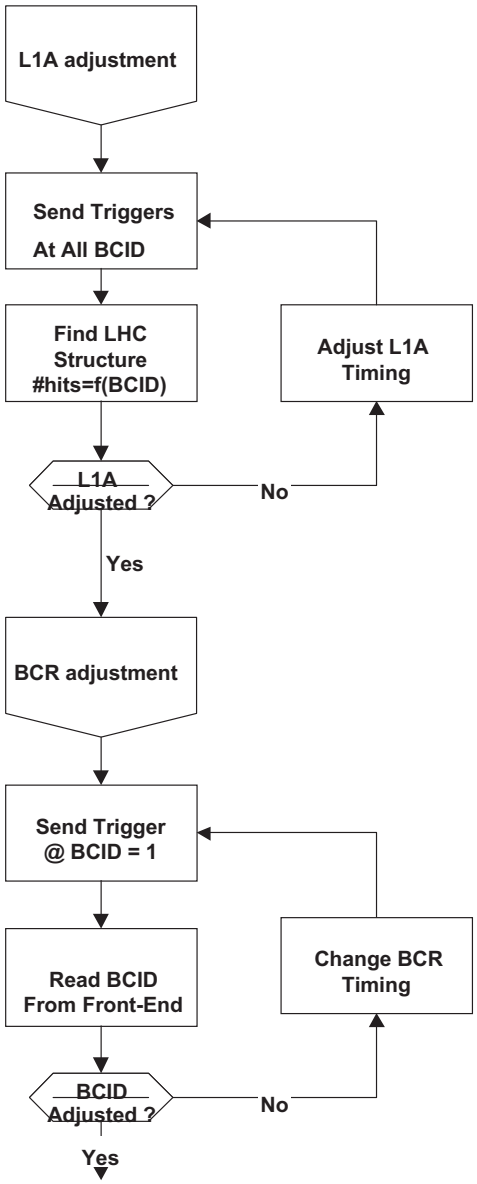

Figure 19-5 Simplified flow chart of the timing set-up with beam when the FE_BCID is formed after the level-1 pipeline.

In the second step, the system will be run with L1A fixed at BCID=1. The event readout will generally contain a FE_BCID different from 1 since the BCR signal timing has not yet been adjusted. The BCR delay is then adjusted in order to obtain the correct FE_BCID value from the front-end.

This method can be implemented very efficiently at high luminosity as the time needed to build the histograms will be very short. It can also be used to check the timing alignment at the beginning of a run for instance.

At low luminosity, the occupancy of the different subdetectors may not be sufficient to use this method efficiently, and, obviously, in the absence of beam this technique cannot be used.

19.3 Timing set-up for beam using test-pulses

It is extremely useful to be able to set up the timing ready for beam before the beam is available. Two objectives are being pursued:

- Set up the timing of the different parts of a subdetector so that only a single global adjustment is needed to align it in time with the other subdetectors.
- Set up the timing of each subdetector with respect to the others. This could be used during the commissioning phase of the detector to have as complete as possible data-taking tests.

This section shows that, at the expense of measuring the length of the fibres which distribute the TTC signals, both objectives can be attained.

Further work is needed to see whether it is possible to do even more with the test pulses, such as the fine BC clock adjustment. It could be envisaged, for instance, to tune the test-pulse timing to simulate the timing of the particles coming from the interaction point, including their time of flight and the detector response time. However, this part is very much subdetector-dependent and cannot be addressed here.

Two different methods of using test pulses are being considered. The first method consists of sending a test pulse followed by a L1A signal after a known delay (this functionality is available in the CTP). The second method consists of sending a test pulse which will produce a L1A signal through the normal level-1 trigger electronics. This can be used, for instance, in the calorimeter.

Figure 19-6 and Table 19-1 show the different delays involved under three conditions — test pulse followed by a test trigger from the CTP, test pulse generating a trigger, and beam collision generating a trigger. The timing adjustment with a test pulse will be considered as efficient if the adjustment of the delay applied to the L1A signal is the same in all three conditions.

In the case [Figure 19-6(c)] of a beam–beam interaction producing a signal at the input of the pipeline at the time t_0, the L1A will be generated at the time t_0 + *Latency* and will reach the front-end at the time t_0 + *Latency* + *d1* + *d2* + *d3* + *Adj1* which must be equal to t_0 + d5. Hence one must have *Adj1* = *d5* – *Latency* – *d1* – *d2* – *d3*.

In the case [Figure 19-6(b)] of a test pulse which generates a L1A, one will have the following if one calls t_1 the time at which the test pulse is initiated at the TTCvi level:
t_1 + *d2* + *d4* + *Adj2* + *Latency* + *d1* + *d2* + *d3* + *Adj1* = t_1 + *d2* + *d4* + *Adj2* + *d5*
i.e. *Adj1* = *d5* – *Latency* – *d1* – *d2* – *d3* which is identical to the value obtained with beam.

In the case where a pulse is sent by the CTP at the time t_2, *d6* microseconds before a L1A is generated, one needs to have:

Table 19-1 Delay elements

Delay	Meaning
d1	Cable delay from CTP to TTCvi
d2	Fibre length from TTC crate to TTCrx
d3	Latency on L1A (TTCvi + TTCrx)
d4	Latency on test-pulse (TTCvi + TTCrx)
d5	Pipeline length in time
d6	Delay between prepulse and L1A in the CTP
Adj1	Delay adjustment on L1A in TTCrx
Adj2	Delay adjustment on test-pulse in TTCrx
Latency	Level-1 latency up to the CTP output

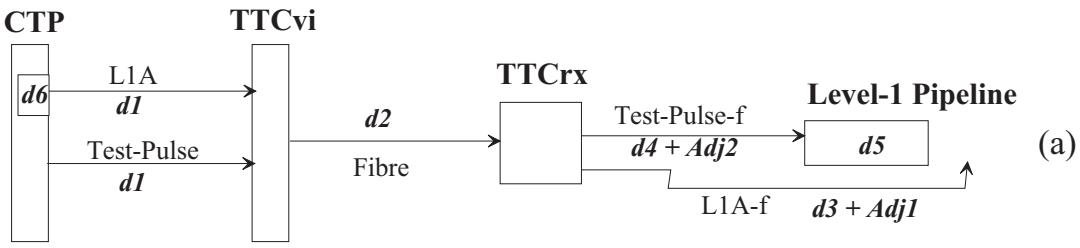

Test Pulse followed by a L1A.
Time adjustment: L1A-f and Test-Pulse-f in the TTCrx

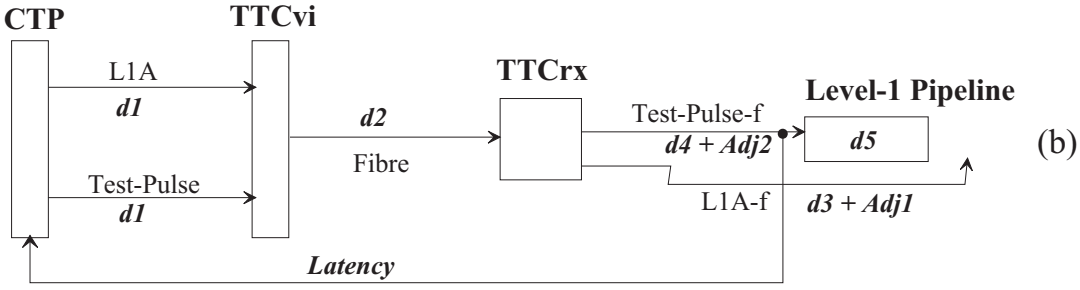

Test Pulse generating a L1A.
Time adjustment: L1A-f and Test-Pulse-f in the TTCrx

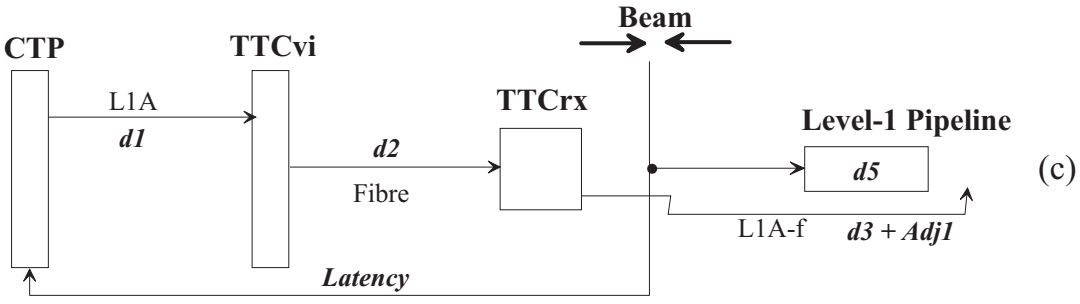

Beam generating a L1A.
Time adjustment: L1A-f in the TTCrx

Figure 19-6 Signal path and relevant timings (in italic).

$t_2 + d1 + d2 + d4 + Adj2 + d5 = t_2 + d6 + d1 + d2 + d3 + Adj$,
i.e.
$Adj1 = d5 – d6 – d3 + d4 + Adj2$.

Here, to have the same Adj1 value as with beam, one needs to have:
$Adj2 = d6 – d4 – d1 – d2 – Latency$

which can easily be achieved as all the different elements are known or easily measured. It only requires that the fibre length be measured at the installation time.

The procedure used to set up the timing is the following: the calculated *Adj2* values are written in every TTCrx and a sequence of test-pulse-followed-by-L1A is started. A scan on all the *Adj1* values is then done until the test-pulse signal is correctly seen in the readout data.

To check that the value used for the LVL1 latency is correct, all of the *Adj1* values are then written in every TTCrx and the system is run in the mode described in Figure 19-6(b). If the timing set-up has to be changed, the LVL1 latency is not correct.

The absolute BCR adjustment cannot be treated here as it requires the beam. Nevertheless, the relative BCR adjustment between different subdetector parts and between subdetectors can be done as it just requires an arbitrary BCID = 0 to be defined in the CTP.

19.4 Timing set-up for test or calibration with test pulses

For each subdetector one needs to be able to test the front-end electronics system and/or to calibrate it. Some timing set-up of test pulses and L1A signals are also necessary in this case, but in this mode of running the CTP is generally not involved and each subdetector (or part of a subdetector) can work independently. The test-pulse, as well as the associated trigger signal, is issued at the TTCvi level of a TTC partition. The different delay elements involved are shown in Figure 19-7.

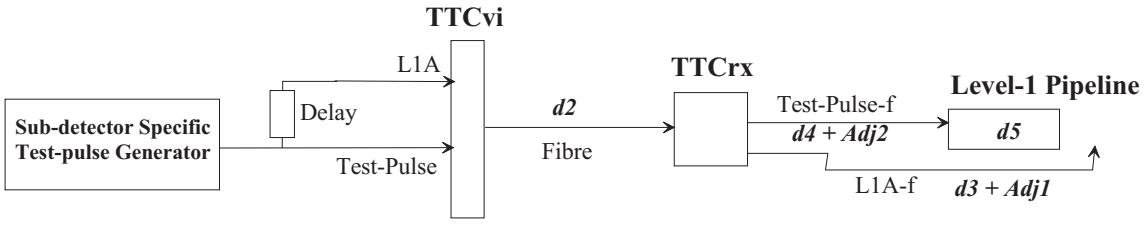

Test Pulse followed by a L1A.
Time adjustment: L1A-f and Test-Pulse-f in the TTCrx, Delay in the sub-detector generator

Figure 19-7 Signal path and relevant timings for sub-detector specific test-pulse/trigger generator.

The Adj1 and Adj2 values defined in Section 19.3 are generally not valid. Either they must be changed or some additional external delay elements must be introduced to make the local test-pulse / trigger timing sequence the same as the one obtained with the CTP.

19.5 References

19-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

20 Strategy for handling deadtime in the experiment

20.1 Introduction

The ATLAS front-end electronics and readout systems contain many levels of buffering. Information may be lost at any of a number stages of the readout chain if buffers become saturated. Different strategies can be adopted to handle this situation, the two extreme ones being:

- introduce deadtime to avoid uncontrolled information loss;
- accept information loss and build a readout system able to accept incomplete events and possible loss of synchronization.

The first of these strategies has been chosen and it has been decided to introduce deadtime in the Central Trigger Processor (see Chapter 15) in order to:

- easily control and monitor the deadtime of the experiment;
- have a relatively simple and safe readout system relying on the presence of data for every event;
- simplify the front-end electronics systems by imposing an upper limit on the event rate and a minimum time between consecutive events.

Nevertheless, the front-end and readout systems are being designed to be robust against buffer overflow and loss of synchronization. Note that a few front-end systems, such as the semiconductor tracker, have variable-length data and, for these, it cannot be excluded that buffers will occasionally overflow with some local loss of data.

As described in Ref. [20-1] and shown in Figure 20-1, there are buffers at different places in the readout chain:

- At the front-end level, a derandomizer buffer, located just after the level-1 pipeline, can store a few events. This is needed in order to match the limited bandwidth of the front-end links, designed to cope with the average LVL1 trigger rate, to the random arrival time of the Level-1 Accept (L1A) signal.
- In the Readout Drivers (ROD) there are buffers at the input stage before the data processing is done and / or at the output stage before the readout link.
- In the Readout Buffers (ROB) the data are stored until the LVL2 trigger has made a decision.

Any of these buffers can become full so some deadtime must be introduced if data loss is to be avoided. The deadtime is introduced in three ways:

- Systematic short deadtime of four bunch crossings after each L1A signal is introduced in the CTP to accommodate front-end electronics limitations.
- The CTP limits the number of L1A signals that can be generated within a given period of time to prevent derandomizer overflows.
- The CTP can be vetoed with an external signal to handle the occupancy of the ROD and ROB buffers.

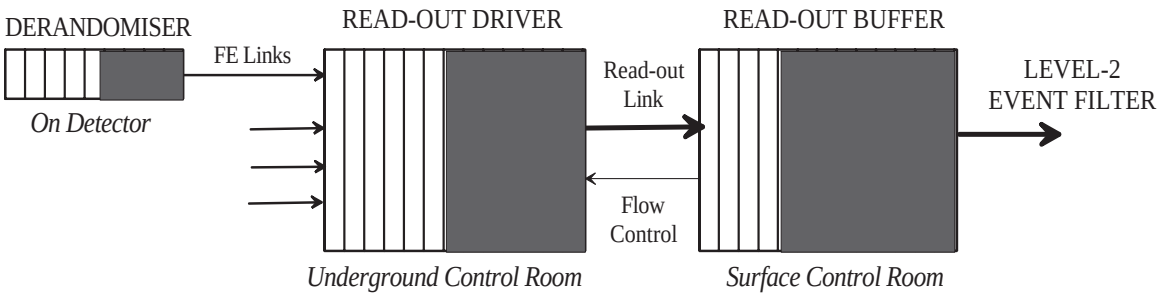

Figure 20-1 Buffers in the readout chain.

20.2 Systematic deadtime

Most of the subdetector front-end systems transmit the data of several consecutive bunch crossings on receipt of a L1A signal. This means that the data of a given bunch crossing could belong to more than one event and hence have to be read out several times. This complicates the readout control part of the front-end electronics when the level-1 pipeline is digital, and is almost impossible to achieve when the level-1 pipeline is analogue. The most demanding sub-detector in this respect is the liquid-argon electromagnetic calorimeter which has an analogue level-1 pipeline and reads out five consecutive bunch crossings per event. As a consequence, a systematic deadtime of four bunch crossings is introduced in the CTP after each L1A signal is issued.

20.3 Derandomizers

The derandomizers are part of the front-end electronics and must be large enough that no more that 1% of the events are lost due to deadtime when the L1A rate is 75 kHz [20-1]. There are several hundreds of thousands of such devices in ATLAS and their characteristics are sub-detector dependent. In some cases, zero-suppression has been applied on the data before they are stored in the derandomizer (e.g. SCT, MDT), while in other cases the full event is stored (e.g. TRT, LAr calorimeter). Some sub-detectors have small derandomizer size but high-speed front-end links, while others have large derandomizer size and low-bandwidth front-end links.

The level-1 trigger electronics must be throttled when the derandomizing buffers are nearly full, in order to avoid the loss of events in some places and hence the loss of synchronization between different parts of the detector. The classical scheme, where each buffer provides a signal to warn it is filling up dangerously, is not practical for two reasons. First, it would lead to a huge number of cables from the detector to the underground counting room where the CTP is located. Second, the time for a signal to transit from the on-detector electronics to the CTP is of the order of 500 ns (100 m of cable), and during that time up to four events can occur which corresponds to 30–50% of the derandomizer capacity.

The occupancy of the front-end derandomizers will be controlled by the CTP itself where an algorithm will be applied to make sure that the derandomizers never fill. This algorithm guarantees that there are no more than N L1A signals within any time window of M microseconds (see Chapter 15). The algorithm used is the ‘leaking bucket’ algorithm: each time an L1A occurs the content of the bucket is increased by a given value (X) and the bucket is

leaking at a constant rate (R). When the bucket is full, deadtime is introduced. The two programmable parameters (X and R) are used to control the deadtime. At the time of writing, the liquid-argon calorimeter electronics has the most stringent running conditions: the derandomizer contains up to 8 events, and 10 μs are necessary to transport an event on the front-end link. The CTP will be programmed so that there are no more than 8 events within any 80-μs time window.

This algorithm does not work when there is zero-suppression in between the level-1 pipeline and the derandomizer. In this case the event size is variable and the derandomizer occupancy is unpredictable. For these configurations, the subdetector electronics have to always have data for all the events and to produce a flag in the data to warn that there are data lost.

20.4 Readout drivers

In the current estimate, there are a total of about 1700 RODs in ATLAS. Most of these modules are located in crates in the underground control room (USA15). The RODs contain buffers which can fill up. These modules have to produce a signal (ROD_BUSY) when their buffer is close to being full in which case deadtime must be introduced. The logical OR of all these ROD_BUSY signals will be used to veto the CTP.

In order to avoid having to OR and monitor up to 1700 at the CTP level, a module will be provided to subdetector groups as described in Section 20.6.

20.5 Readout buffers

For the ROBs, an identical strategy to that of the RODs could be applied. However, as the ROBs are located in the surface counting room and the CTP is in the underground counting room, it is not deemed a very practical solution. Furthermore, deadtime could be introduced too early as a given ROB which is filling up does not know about the availability of buffer space in the ROD to which it is connected.

It is planned to use the flow control mechanism of the readout link to apply back pressure on the ROD. If a ROB has no available space in its buffer it disables the data transmission from the ROD. The buffer of the ROD may then fill up in which case the ROD_BUSY will be asserted in due course (if necessary).

20.6 Busy handling

20.6.1 BUSY module

As mentioned earlier, there will be about 1700 RODs in ATLAS, each of them providing a ROD_BUSY signal and capable of introducing deadtime. It is therefore very important to:

- monitor and control all of the ROD_BUSY signals so that any pathological ROD can be prevented from introducing deadtime;

- monitor the duration of these signals;
- gather the ROD_BUSY signals in a tree structure so that at the central-trigger-logic level only one BUSY signal per subdetector appears. An additional module will handle these signals and provide a VETO signal to the CTP.

A VME module handling up to 20 ROD_BUSY signals (i.e. corresponding to one crate full of RODs) will be made available. Its block diagram is shown in Figure 20-2.

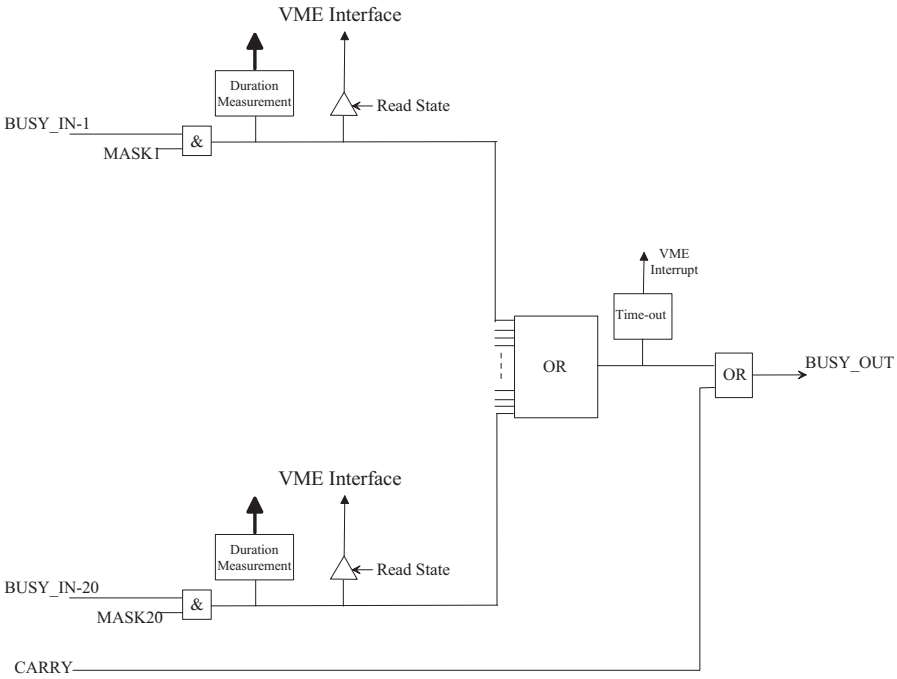

Figure 20-2 BUSY module block diagram.

Each input BUSY_IN signal can be individually masked and its integrated duration measured. A 16-bit counter clocked with the BC clock is used for this purpose. The duration measurement is available both in time and in number of BCs. It requires a readout of the counters every 1.6 ms, which is easily done by the processor controlling this module.

The logical OR of the unmasked BUSY_IN signals is made available as an output. The state of the BUSY_IN inputs can be read out through the VME interface for debugging purposes. A VME interrupt can be issued if the duration of one of the BUSY_IN signals exceeds a time limit in order to warn the processor controlling this module that an action has to be effected.

In order to be able to chain the modules in a tree structure (as shown in Figure 20-3), a CARRY input is made available. This input is connected to the BUSY_OUT of the previous module.

At the central-trigger-logic level, an additional module of the same kind is used to handle one ROD_BUSY signal per subdetector, and its output is used to veto the CTP.

Such a structure allows an efficient control of the deadtime in the experiment as well as an easy way to detect a faulty module introducing deadtime. It also allows the partitioning of the readout as a subdetector BUSY contribution can easily be removed from the CTP VETO. In the same way, additional subpartitioning within subdetectors can be implemented.

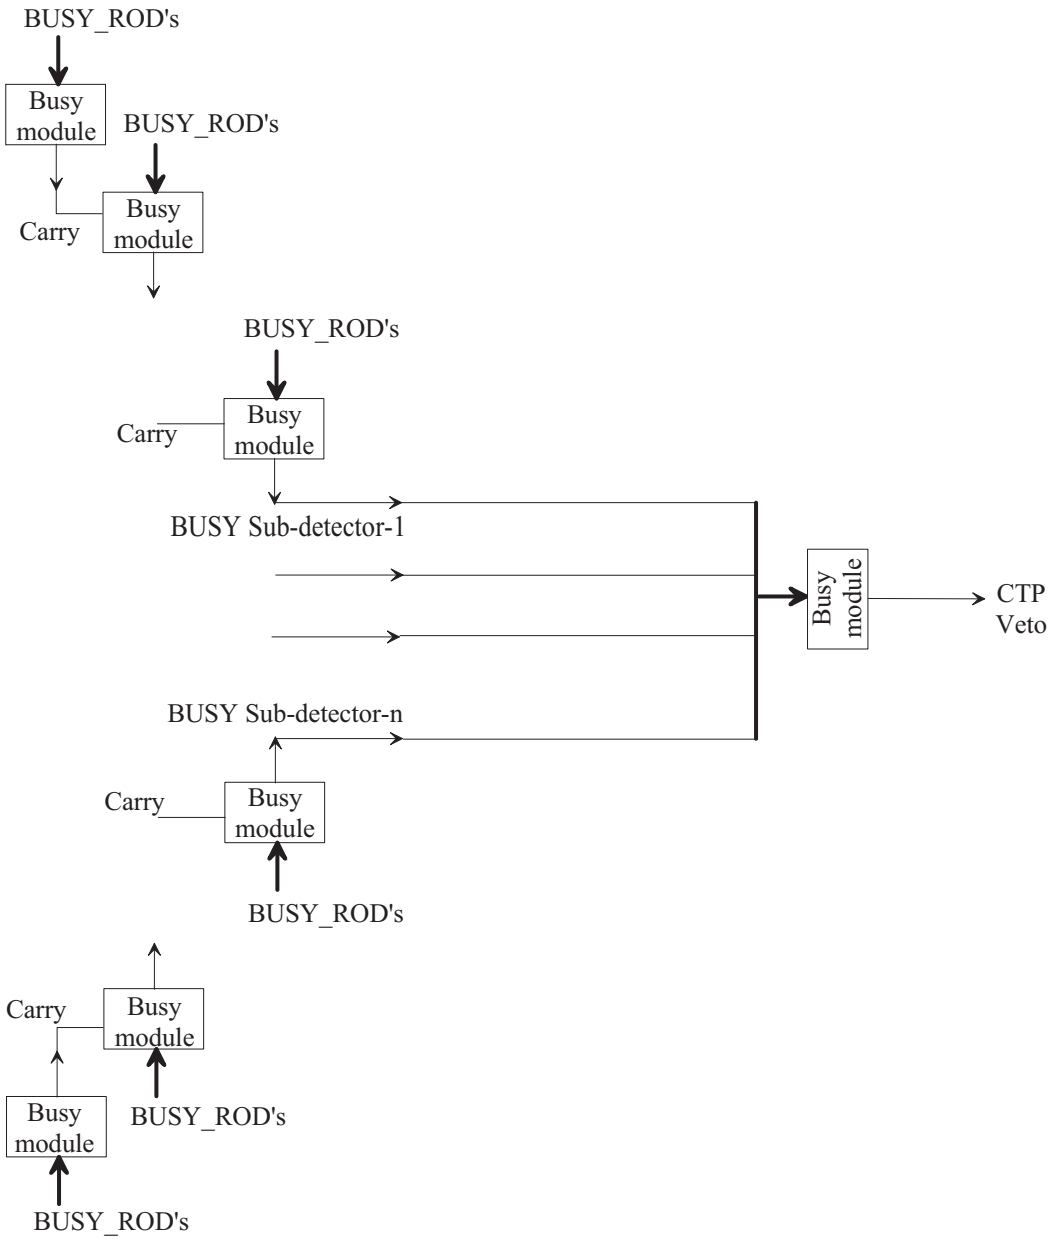

Figure 20-3 Chaining the BUSY modules.

20.6.2 Quality assurance and review procedure

Figure 20-4 shows the schedule for the design and production of the BUSY module.

After the specifications are defined, a preliminary design review (PDR) will take place. It will examine and assess the full requirements and specifications for the module, identify any missing functionality, ensure full compatibility with all connecting subsystems and determine overall feasibility.

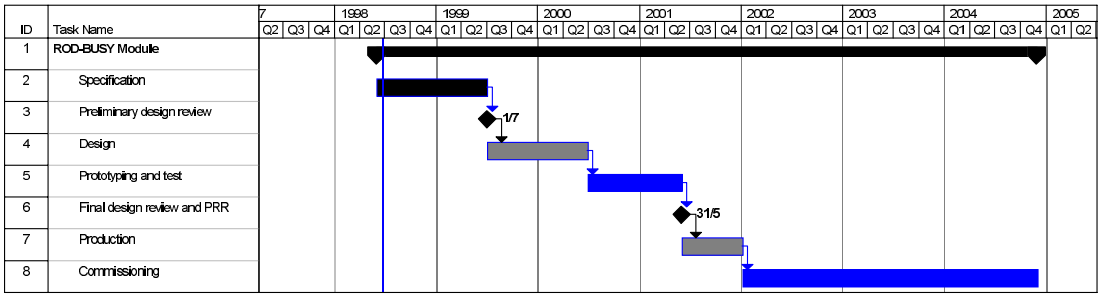

Figure 20-4 Schedule for the BUSY module production.

After the design is finished and before the production is launched, a final design review (FDR) will take place and be used as an ATLAS production readiness review (PRR).

The production of the boards will follow usual rules, namely:

- Electrical tests of the PCB before assembly.
- Burn in of the boards after assembly by leaving them under power without cooling for two days.
- Test of the full boards.

Each element will receive a serial number and a database will be set up to store the information relative to each board: origin of the components which populate the board, results of tests, history of failure and location.

In order to facilitate the test and the maintenance of the boards, the JTAG test bus will be used as much as possible.

20.7 References

20-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

21 Installation, access and maintenance

21.1 On-detector muon trigger electronics

Aspects related to the installation and maintenance of the on-detector muon-trigger electronics are to a large extent covered in the TDR for the muon system [21-1] and will be discussed in more detail in the forthcoming Technical Coordination TDR [21-2]. The on-detector muon-trigger subsystems require a significant amount of on-detector cabling, as well as cables between the detector and the USA15 counting room. Some of the electronics will be mounted on the chamber modules and some intra-module cabling will be carried out, prior to installation in the cavern. The remaining connections will be made at the time of installation.

21.2 General comments on electronics, cables and equipment

The electronics that are installed outside the cavern, either in the USA15 underground counting room or in surface buildings, will be mounted in standard racks. No special problems are foreseen with the installation or maintenance of this equipment. Nevertheless, the requirements analyses that have been performed have paid careful attention to issues of maintenance and repair.

The design of all of the trigger electronics, whether in the cavern or elsewhere, has to be such that faulty components (electronic modules, power supplies, backplanes, cooling units) can be replaced quickly. Other issues that have been addressed are the need for adequate supplies of spare parts, the availability at CERN of expert personnel during running periods and the need for test rigs at CERN.

Concerning the trigger-processor electronics, careful attention has to be given to the crate and rack layout, and to the routing of cables to, from, and within these electronics systems. A critical issue is the latency of the LVL1 trigger (see Chapter 18) for which propagation delays along electrical cables and optical fibres is a large contributor. In cooperation with the ATLAS Electronics Coordination and Technical Coordination groups, an optimal rack layout and cable routing has been devised that minimizes the overall latency. The LVL1 trigger team are working closely with the ATLAS Technical Coordination group on the routing of cables for electronics in the cavern and for connections between the cavern and USA15. The cable lengths that are used in the latency evaluations were calculated by the Technical Coordination team. The cable length calculations take into account the constraints that come from the need to open the detector for access without uncabling. For the connections between the cavern and USA15, the shortest available path is taken through the shielding wall, whilst making sure that the radiation levels in USA15 will be safe for unlimited access by personnel.

Another aspect of cable routing that needs to be considered is related to the issues of maintenance and repair — cables should be routed in such a way that they do not prevent easy replacement of electronic modules. In practical terms, this means that it should be possible to replace any module in the system without massive recabling of neighbouring modules in the system.

Apart from the custom electronics of the LVL1 trigger processors, there will be a significant amount of commercial electronics and computing equipment. This will include computers used to test, control, monitor and read out the trigger processors. No special problems of installation or maintenance are foreseen for this equipment which will be located in USA15 and also in surface buildings.

The issues discussed here, as well as the installation schedule for the cables and electronics, will be addressed in detail in the forthcoming Technical Coordination TDR [21-2].

21.3 References

21-1 *ATLAS muon spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

21-2 *ATLAS Technical Coordination Technical Design Report*, to be submitted to the LHCC.

22 Safety and environment

22.1 Introduction

A general description of the ATLAS safety aspects and the proposed risk-mitigation strategies is contained in the forthcoming Technical Coordination Technical Design Report [22-1]. This chapter deals only with those safety aspects particular to the LVL1 trigger system.

The LVL1 trigger system consists of electronics modules that are, with the exception of part of the muon trigger electronics, located in the USA15 underground counting room which is shielded against radiation.

The trigger system itself does not deal with large mechanical systems, high-voltage electrical equipment or flammable gasses. Part of the muon trigger system is mounted on the detector; the safety issues associated with the muon detector systems have already been documented [22-2].

The following sections discuss risks related to fire hazards, electricity and radiation, as well as associated considerations of access and cooling.

22.2 Fire safety

The LVL1 trigger system will contain a large number of electrical cables and optical fibres. Standard CERN rules [22-3] will be applied concerning the choice of cables. In particular, all cables will have a flame-resistant insulation that does not generate toxic fumes in case of fire.

The equipment located in the USA15 counting room will be mounted in standard CERN racks that include fire-detectors linked to the safety system in the experimental area. Damage to equipment in case of a localized fire (in cables, electronics modules or low-voltage power supplies) would be minimized by automatic fire-fighting equipment that would extinguish the fire using inert gasses. In such an eventuality, the automatic systems would also cut off the power supply to the associated equipment.

Issues relating to fire safety in the ATLAS cavern will be addressed in the forthcoming Technical Coordination Technical Design Report [22-1].

22.3 Electrical safety

Given the absence of high-voltage electrical equipment in the LVL1 trigger system, there are no particular electrical risks associated with personnel safety. Standard procedures and regulations will be followed concerning the use of equipment connected to the mains.

The LVL1 trigger systems, both in USA15 and in the cavern, require very high-current, low-voltage power supplies. As discussed above, these will be interlocked with the fire-protection systems. The Detector Control System, which will provide extensive monitoring of the low-

voltage systems (temperatures, voltages, currents) will detect problems and, if necessary, cut off the power supply in order to protect the equipment.

22.4 Radiation safety

A large part of the LVL1 trigger electronic system is located in the USA15 underground counting room that is well shielded against radiation. There are no radiation issues, either for personnel or for equipment, for this part of the system. The radiation levels will be sufficiently low to allow unlimited access to the USA15 area even when there is beam in the LHC machine.

Concerning the muon trigger electronics that is mounted in the cavern, and also the cables between the cavern and USA15, no radiation problems are expected from the point of view of personnel. This equipment is mounted on the outside of the detector, relatively well shielded from the interaction point, where there will be no significant activation. The overall CERN and ATLAS radiation protection policy will be followed.

Radiation issues do have to be taken into account in the design of the on-detector muon-trigger sub-systems as discussed in Chapters 9–14. In particular, all such electronics has to be tolerant to radiation up to the expected levels.

22.5 Access

As discussed in the Technical Design Report for the muon system [22-2] and in more detail in the forthcoming Technical Coordination TDR [22-1], appropriate passages and holes must be provided to allow access to the inner chambers for maintenance, repair, etc. In particular, these must allow access to the on-detector muon-trigger electronics for maintenance and repair operations.

22.6 Cooling

Issues of cooling will be addressed in the forthcoming Technical Coordination TDR [22-1], so only a very brief discussion is given here on matters relating to cooling of the LVL1 trigger electronics.

The equipment in the USA15 counting room will be mounted in racks with standard cooling systems. These will be fitted with temperature monitors and an interlock system that will cut off the electrical power supply in case of over heating.

The on-detector electronics associated with the RPC detectors will rely on natural air-cooling. As discussed in Chapter 12, active cooling is required for the on-detector electronics associated with the TGC detectors (higher density of electronics than in the RPC case). Temperature monitoring will prevent over-temperature situations causing damage to equipment; in case of a cooling failure (or any other over-temperature condition), the electrical power supply would be cut off.

22.7 References

22-1 *ATLAS Technical Coordination Technical Design Report*, to be submitted to the LHCC.

22-2 *ATLAS Muon Spectrometer Technical Design Report*, CERN / LHCC / 97–22, May 1997.

22-3 CERN / TIS safety instruction, IS 23, *criteria and standard test methods for the selection of electrical cables, wires and insulated parts with respect to fire safety and radiation resistance.*

23 Project organization, management, schedule and costs

23.1 Participating institutes

The following list contains the institutes which are participating in the LVL1 trigger project, and the names of the physicists and senior engineers from each institute working on the project.

Birmingham
P. Bright-Thomas, A. Connors, J. Garvey, S. Hillier, R. Staley, P. Watkins, A. Watson.

CERN
I. Brawn, N. Ellis, P. Farthouat, P. Gallno, C. Havet, R. McLaren, G. Schuler, C. Schwick.

Haifa
N. Lupu, S. Robins, S. Tarem.

Heidelberg
C. Geweniger, P. Hanke, E-E. Kluge, J. Krause, K. Meier, U. Pfeiffer, A. Putzer, K. Schmitt, C. Schumacher, K. Tittel, M. Wunsch.

KEK
H. Iwasaki, T. Ohska, O. Sasaki, S. Tanaka, K. Yamauchi.

Kobe
K. Kawagoe, H. Kurashige, M. Nozaki, H. Takeda.

Kyoto
H. Sakamoto.

Lecce
P. Creti.

London QMW
E. Eisenhandler, M. Landon, J.M. Pentney.

Mainz
B. Bauss, K. Georgi, K. Jakobs, G. Quast, U. Schäfer.

Naples
A. Aloisio, F. Cavenini.

Rome I
A. Di Mattia, S. Falciano, L. Luminari, A. Nisati, E. Petrolo, S. Veneziano.

Rome II
R. Cardarelli, A. Di Ciaccio, R. Santonico.

RAL
J. Edwards, C.N.P. Gee, A.R. Gillman, R. Hatley, V.J. Perera, D.P.C. Sankey, T.P. Shah.

Shinshu
T. Takeshita.

Stockholm U
C. Bohm, M. Engström, S. Hellman, S. Silverstein.

Tel Aviv
S. Boettcher.

Tokyo ICEPP
Y. Hasegawa, K. Homma, M. Imori, T. Kawamoto, T. Kobayashi.

Tokyo MU
C. Fukunaga, R. Hamatsu.

Weizmann Institute of Science
E. Duchovni, E. Gross, D. Lellouch, L. Levinson, G. Mikenberg, K. Nagai.

23.2 Responsibilities and work organization

The LVL1 trigger system is divided into a number of subsystems between which interaction is limited. These are the calorimeter trigger, the muon trigger, and the central trigger (i.e. CTP and TTC systems). The muon trigger is further subdivided into the parts associated with the RPC detectors and the TGC detectors, and the part that forms the interface to the CTP. The institutes and funding agencies responsible for each of these parts are listed in the ATLAS Memorandum of Understanding [23-1]. A slightly more detailed list of responsibilities, also indicating the sharing of work within the calorimeter trigger and the muon trigger, is given in Table 23-1.

The work in each of the main areas (calorimeter trigger, muon trigger, central trigger) is organized largely independently. Issues of interfaces and where a common approach is needed are addressed in general LVL1 trigger working-group meetings and in *ad hoc* discussions between the groups concerned.

23.3 Management organization

The overall Trigger/DAQ project of ATLAS, covering the LVL1 and LVL2 triggers, the data acquisition and event filter, and the detector-control system, is organized in accordance with the normal ATLAS rules [23-2][23-3]. A Steering Group (see Table 23-2), with members representing different areas of the project, is the executive body responsible for managing the project as a whole. An Institutes Board, with one voting representative per institute, is responsible for deciding policy and for matters related to resources. Technical and scientific matters are discussed in working-group meetings for different areas of the project and in more specialized *ad hoc* working meetings. Responsibility for leading the project is currently shared between two coordinators, N. Ellis responsible for the trigger, and L. Mapelli responsible for the data acquisition and event filter. Steering Group meetings are chaired by one or other of the coordinators. The Institutes Board is presently organized by co-chair-people, M. Abolins and J.R. Hansen, who are *ex-officio* members of the Steering Group.

Table 23-1 Sharing of responsibilities for LVL1 trigger project

Item	Funding agencies	Institutes
Calorimeter trigger		
Front-end preprocessor	Germany (BMBF)	Heidelberg
Cluster processor (e/γ, h/τ)	UK	Birmingham, London QMW, RAL
Jet/energy-sum processor	Germany (BMBF), Sweden	Mainz, Stockholm
Processor readout system	UK	Birmingham, London QMW, RAL
Muon trigger		
RPC-based part	Italy	Naples, Lecce, Rome I, Rome II
TGC-based part	Japan, Israel	KEK, Kobe, Kyoto, Shinshu, Tokyo ICEPP, Tokyo MU; Haifa, Tel Aviv, Weizmann
Interface to CTP	CERN	CERN
Central trigger		
Central trigger processor	CERN	CERN
TTC system	CERN	CERN

Table 23-2 Present composition (June 1998) of the Trigger/DAQ Steering Group.

Activity	Members(s)
DAQ Coordinator	L. Mapelli
Trigger Coordinator	N. Ellis
LVL1 calorimeter trigger	E. Eisenhandler
LVL1 muon trigger	E. Petrolo
LVL2 trigger	S. Falciano, P. Le Du, F. Wickens
DAQ	R. Jones, G. Mornacchi
Event Filter	F. Etienne
Detector Control System	H. Burckhart
Trigger Performance	T. Hansl-Kozanecki
FE electronics	Ph. Farthouat

As discussed in the preceding section, the LVL1 trigger project is divided into a number of sub-projects that interact only weakly. These are the calorimeter trigger, the muon trigger and the central trigger (CTP and TTC systems). Both of the calorimeter and the muon trigger are explicitly represented on the Steering Group; the central trigger is represented by the Trigger Coordinator.

The organization of the LVL1 calorimeter trigger sub-project is as follows. Two- or three-day meetings are held every few months, with participation from all of the collaborating institutes, at which progress is reviewed and plans are discussed. These generally take place either in one of the collaborating institutes or at CERN. A management board, with one representative per

institute, addresses issues of resources within the sub-project and is also the decision-making body for the LVL1 calorimeter trigger sub-project. Decisions are subject to approval by the Steering Group, for very major decisions the Institutes Board, and ultimately the ATLAS Collaboration Board.

For the muon trigger a similar scheme to that described above for the calorimeter trigger has been adopted for the overall sub-project management. Meetings are held every few months at which progress is reviewed and plans are discussed. These usually take place at CERN prior to each ATLAS week (four meetings per year). Decisions are subject to approval by the Steering Group, for very major decisions the Institutes Board, and ultimately the ATLAS Collaboration Board.

Responsibility for the central trigger lies with a single institute (CERN). Progress is reviewed and plans are discussed in regular meetings. As for the other parts of the system, significant decisions are subject to approval by the Steering Group, for very major decisions the Institutes Board, and ultimately the ATLAS Collaboration Board.

Monitoring of the LVL1 project and its sub-projects takes place at a number of levels, firstly within each sub-project, then within the T/DAQ system, and finally ATLAS-wide. Monitoring is achieved through technical discussions in meetings that review plans and progress, by the work of the sub-project leaders and the overall Trigger Coordinator, and via formal reviews. Several levels of review are planned — internal reviews organized by the T/DAQ community and ATLAS reviews organized by the Technical Coordination group (e.g. Production Readiness Reviews). Regular progress checks will be made with the milestones defined in the detailed schedules, following the procedures defined by the ATLAS Technical Coordination group. Major milestones that can be used by the LHCC to monitor the progress of the project are proposed.

23.4 Schedule and milestones

In this section we list milestones for each of the major components of the level-1 trigger. Detailed schedules are given in their respective chapters:

- calorimeter trigger — Section 8.9;
- barrel muon trigger — Section 11.5;
- endcap muon trigger — Section 12.14.1;
- muon trigger CTP interface — Section 13.8;
- CTP — Section 15.7;
- TTC — Section 16.9.

23.4.1 Calorimeter trigger

Table 23-3 Milestones for the level-1 calorimeter trigger.

Milestone	Date
Calorimeter trigger system specifications completed	1998 Q4
Final Design Review for PPrASIC	2000 Q4
Final Design Review for PPrMCM	2000 Q4
Final Design Review for CPASIC	2001 Q2
Final Design Review for CPMCM	2001 Q2
Final Design Review for CP and JEP ROD module	2001 Q2
Final Design Review for PPM	2001 Q4
Final Design Review for JEM	2001 Q4
Final Design Review for CPM	2002 Q3
Preprocessor system tests completed	2003 Q4
Cluster Processor system tests completed	2004 Q2
Jet / Energy-sum Processor system tests completed	2004 Q2
Preprocessor system integrated with calorimeters	2004 Q3
Calorimeter trigger system commissioning completed for cosmic ray run	2004 Q4

23.4.2 Muon trigger

Table 23-4 Milestones for the level-1 muon trigger.

Milestone	Date
Muon endcap trigger readout protocol fixed	1999 Q4
CTP interface preliminary design review	2000 Q1
Muon barrel trigger coincidence matrix final design review	2000 Q1
Muon barrel trigger pad logic final design review	2000 Q3
Muon endcap trigger first full system test	2000 Q3
Muon barrel trigger readout driver final design review	2001 Q1
Muon barrel trigger sector logic final design review	2001 Q1
Muon barrel trigger optical link final design review	2001 Q1
Muon endcap trigger full system test before mass production	2001 Q3
CTP interface final design review and PRR	2002 Q2

23.4.3 CTP and TTC

Table 23-5 Milestones for the level-1 CTP and TTC.

Milestone	Date
CTP preliminary design review	1999 Q1
TTC crate preliminary design review	1999 Q4
TTCvi preliminary design review	1999 Q4
TTC crate final design review and PRR	2000 Q2
TTCrx preliminary design review	2000 Q2
TTCvi final design review and PRR	2000 Q3
TTCrx final design review	2000 Q4
CTP final design review and PRR	2001 Q1

23.5 Cost and resources

The ATLAS Collaboration has the necessary resources to build the level-1 trigger, as described in the Memorandum of Understanding [23-1]. The cost of the level-1 trigger is estimated to be as follows, in units of millions of ATLAS Swiss Francs (MASF), as defined in [23-4].

23.5.1 Calorimeter trigger

Table 23-6 Cost of the level-1 calorimeter trigger.

Item	Cost (MASF)
Preprocessor	2.417
Cluster Processor	3.334
Jet / Energy-sum Processor	1.230
Readout system (for CP and JEP)	0.287
Joint items (cables, computing, etc.)	0.394
Total	**7.662**

23.5.2 Muon trigger

Table 23-7 Cost of the level-1 muon trigger.

Item	Cost (MASF)
Barrel RPC trigger	2.605
Endcap TGC trigger	3.340
CTP interface	0.674
Total	**6.619**

23.5.3 CTP and TTC

Table 23-8 Cost of the level-1 CTP and TTC.

Item	Cost (MASF)
Central Trigger Processor	0.533
Timing, Trigger and Control system	1.308
Total	**1.841**

23.6 References

23-1 *Memorandum of understanding for collaboration in the construction of the ATLAS detector,* RRB–D98–44, April 1998.

23-2 *ATLAS Organization,* ATLAS note GEN–NO–09, 1994.

23-3 *ATLAS System Organization,* ATLAS note GEN–NO–15, 1996.

23-4 *ATLAS Cost Planning,* Version 7.0, January 1998.

A Appendix: Product breakdown structure

10.1 LVL1 Trigger

10.1.2.1 LVL1 Calorimeter Trigger Processor

10.1.2.1.1 Front-End Preprocessor
10.1.2.1.2 Cluster Processor
10.1.2.1.3 Jet/Energy-sum Processor
10.1.2.1.3.1 Combined Items
10.1.2.1.3.2 Jet Trigger Modules
10.1.2.1.3.3 Energy-sum Trigger Modules
10.1.2.1.4 Readout System
10.1.2.1.5 Joint Items

10.1.2.2 LVL1 Muon Trigger Logic

10.1.2.2.1 LVL1 Muon Barrel Trigger Processor
10.1.2.2.1.1 Low Pt Trigger
10.1.2.2.1.1.1 Local Logic Board
10.1.2.2.1.1.1.1 Coincidence Matrix Chip
10.1.2.2.1.1.2 Pad Logic Board
10.1.2.2.1.2 High Pt Trigger
10.1.2.2.1.2.1 Local Logic Board
10.1.2.2.1.2.1.1 Coincidence Matrix Chip
10.1.2.2.1.2.2 Pad Logic Board
10.1.2.2.1.3 Sector Logic
10.1.2.2.1.4 Link from Sector Logic to Muon Interface to CTP
10.1.2.2.1.5 Common Items
10.1.2.2.2 LVL1 Muon Endcap Trigger Processor
10.1.2.2.2.1 Patch Panel
10.1.2.2.2.1.1 Patch Panel Chip
10.1.2.2.2.1.2 Patch Panel PCB for doublet wire
10.1.2.2.2.1.3 Patch Panel PCB for doublet strip
10.1.2.2.2.1.4 Patch Panel PCB for triplet wire
10.1.2.2.2.1.5 Patch Panel PCB for triplet strip
10.1.2.2.2.1.6 TTC receiver and distributor
10.1.2.2.2.1.7 DCS port
10.1.2.2.2.2 Slave Board
10.1.2.2.2.2.1 Coincidence Matrix Chip
10.1.2.2.2.2.2 Doublet Slave Board PCB for wire
10.1.2.2.2.2.3 Doublet Slave Board PCB for strip
10.1.2.2.2.2.4 Triplet Slave Board PCB for wire
10.1.2.2.2.2.5 Triplet Slave Board PCB for strip
10.1.2.2.2.3 High Pt Board
10.1.2.2.2.3.1 Coincidence Matrix Chip
10.1.2.2.2.3.2 High Pt Board PCB for wire
10.1.2.2.2.3.3 High Pt Board PCB for strip
10.1.2.2.2.3.4 High Pt crate
10.1.2.2.2.4 Link from High Pt Board to Sector Logic
10.1.2.2.2.5 Sector Logic
10.1.2.2.2.5.1 Phase adjust
10.1.2.2.2.5.2 R-Phi Coincidence Matrix
10.1.2.2.2.5.3 Track Pre-selector

10.1.2.2.2.5.4 Track Selector
10.1.2.2.2.5.5 Track Encoder
10.1.2.2.2.6 Link from Sector Logic to Muon CTP Interface
10.1.2.2.2.7 Readout link from Slave Board to Star Switch
10.1.2.2.2.8 Star Switch
10.1.2.2.2.9 Readout link from Star Switch to ROD crate
10.1.2.2.2.10 ROD crate
10.1.2.2.2.10.1 ROD
10.1.2.2.2.10.2 Local DAQ Master
10.1.2.2.2.11 Common Items
10.1.2.2.3 LVL1 Muon Interface to Central Trigger Processor
10.1.2.2.3.1 Octant Board [16]
10.1.2.2.3.2 ROI - ROD Board
10.1.2.2.3.3 Interface to CTP
10.1.2.2.3.4 Crate [9U, VME]
10.1.2.2.3.4.1 Readout Backplane

10.1.2.3 LVL1 Central Trigger Logic

10.1.2.3.1 Central Trigger Processor
10.1.2.3.2 Deadtime Control
10.1.2.3.3 Readout Driver
10.1.2.3.4 LVL2 Interface
10.1.2.3.5 Monitoring Tools
10.1.2.3.6 Crate Processor
10.1.2.3.7 Crate
10.1.2.3.7.1 Readout Backplane
10.1.2.3.7.2 ROD Busy Fan In / Out

10.1.2.4 Timing, Trigger and Control Distribution

10.1.2.4.1 TTC Crate
10.1.2.4.2 TTCvi
10.1.2.4.3 TTCrx

B Appendix: Definitions, acronyms, abbreviations

ADC
Analogue-to-digital converter.

AMUX
Analogue multiplexer.

ASD
Amplifier-shaper-discriminator circuit.

ASD-IC
Amplifier-shaper-discriminator integrated circuit.

ASIC
Application-specific integrated circuit. A custom-made integrated circuit.

ATLFAST
ATLAS software package for fast particle-level simulation.

ATRIG
ATLAS software package for trigger simulation.

Bakelite
RPC construction material: phenolic resin.

Barrel
The central-rapidity region of either the muon spectrometer, the electromagnetic calorimeter or the hadronic calorimeter.

Baud
Bits per second. Used for serial transmission; this rate is the total including all framing and protocol bits, as opposed to the net rate for transmitting actual data.

BC
Proton–proton bunch crossing in the LHC. The bunch spacing is 24.95 ns.

BCID
See Bunch-Crossing Identification.

BC-MUX; BC-multiplexing
Bunch-crossing multiplexing. Used in the Level-1 Calorimeter Trigger to double the number of trigger towers per serial link by using the fact that adjacent bunch-crossings cannot both carry valid trigger-tower data after BCID has been carried out.

BCR
See Bunch Counter Reset.

Bd
See Baud.

BER
Bit error rate.

BGA
Ball grid array.

Boundary scan (JTAG)
See JTAG.

BTL

Backplane transceiver logic. Devices intended for operation in the Futurebus+ signalling environment. They operate typically from +5 V power supplies, with signal lines terminated to +2.1 V. Receivers have a precise threshold (+1.55 V) for maximum noise immunity, and generated noise is minimized by incorporating slew-rate control.

Bunch Counter Reset (BCR)

Signal broadcast by the TTC system once per LHC orbit (88.924 s)to control the phase of local bunch counters.

Bunch-crossing identification (BCID)

The assignment of detector data to a specific bunch crossing.

Bunch_crossing_ID (ROD_BCID)

See ROD_BCID.

Calorimeter cell

The smallest unit of calorimeter information to be read out and digitized.

CAN

Control area network. A field bus for controlling and monitoring, used in the Detector Control System.

CCM

Clock Control Module of the Level-1 Calorimeter Trigger.

Central Trigger Processor (CTP)

The part of the Level-1 Trigger System which combines results from the Level-1 Calorimeter and Muon Triggers to make the global (yes / no) Level-1 Trigger decision for each bunch crossing.

CF

Cluster-finding.

Clock

The 40.08 MHz clock, linked to the LHC machine, used to synchronize the pipelined LVL1 processing system.

Cluster Processor

The part of the Level-1 Calorimeter Trigger that carries out the electron / photon and hadron / tau triggers.

CMM

Cluster Merger Module of the Level-1 Calorimeter Trigger.

CMOS

Complementary metal-oxide semiconductor. Switching logic with very low quiescent power dissipation. The logic levels are approximately equal to the supply voltage rails, covering a wide range down to +3 V. Compatibility with TTL is easily achievable.

CMS

Compact muon solenoid. The other ‘general purpose’ detector at the LHC.

CM

Coincidence matrix of the level-1 muon trigger.

CP

Cluster Processor of the Level-1 Calorimeter Trigger.

CPASIC
Cluster Processor ASIC of the Level-1 Calorimeter Trigger.

CPLD
Complex programmable logic device.

CPM
Cluster Processor Module of the Level-1 Calorimeter Trigger.

CPMCM
Cluster Processor multi-chip module of the Level-1 Calorimeter Trigger.

CPU
Central processing unit.

Counting room
See USA15.

CSC
Cathode strip chamber.

CTP
See Central Trigger Processor.

DAC
Digital-to-analogue converter.

DAQ
See Data Acquisition System.

Data Acquisition System (DAQ)
System responsible for the assembly and permanent storage of events accepted by all three levels of the trigger system (Level-1, Level-2, Event Filter) and data generated within the trigger systems.

DCS
See Detector Control System.

Derandomizer
The memory in which data corresponding to a Level-1 Accept are stored before being read out. The name is due to the fact that the data are available one event after the other whatever the time delays between successive Level-1 Accept signals was.

Detector Control System (DCS)
The system which monitors and controls physical parameters of the sub-systems of the experiment, such as gas pressure, flow-rate, high voltage settings, low-voltage power supplies, temperatures, leakage currents, etc.

DICE
ATLAS software simulation

DMA
Direct memory access: a mode of fast data transfer.

Doublet
Part of the muon spectrometer consisting of two layers of thin gap chambers.

DPMJET
Monte Carlo program used to generate simulated proton–proton interactions.

DSB

Doublet Slave Board of the muon endcap trigger.

DSL

Detector-specific logic. For each sector of the RPC and TGC systems provides an input to the muon-trigger / CTP interface.

ECL

Emitter-coupled logic. Refers to non-saturating, high-speed logic with a differential swing of ±800 mV (–1.7 to –0.9 V), operating from a power-supply voltage of –5.2 V.

ECR

See Event Counter Reset.

EF

See Event Filter.

Electromagnetic isolation (e.m. isolation) region

A continuous region in the e.m. calorimeters, surrounding a potential electron / photon or hadron / tau cluster, whose summed transverse energy is used as a veto if it exceeds some (usually low) threshold. This is to help reject jet background.

Electron/photon cluster (e.m. cluster)

The areas in η–ϕ which are summed in the e.m. calorimeters in order to compare their transverse energy with electron / photon trigger thresholds. The elements used in the summing are e.m. Trigger Towers.

Electron/photon de-cluster/RoI region

The areas in η–ϕ which are summed in order to find local maxima in transverse energy, necessary for de-clustering electron / photon triggers and for electron / photon Regions-of-Interest. The elements used in the summing are e.m. Trigger Towers.

Electron/photon window

The areas in η–ϕ which are used for the electron / photon trigger algorithm, combining Electron / photon Clusters with isolation requirements in both the electromagnetic and hadronic calorimeters.

EMI

Electromagnetic interference.

Endcap

The high-η part of the muon spectrometer and calorimeter systems. The endcap covers the two ends of the cylindrical ‘barrel’ central region of ATLAS.

E.M.

Electromagnetic. (Can refer to either the calorimeters — e.m. vs. hadronic — or to the e.m. trigger, aimed at detecting electrons and photons.)

EPROM

Erasable programmable read-only memory.

Event

The data resulting from a particular bunch crossing. At high luminosity, this could contain data from several physics processes.

Event Counter Reset (ECR)

Signal broadcast by the TTC system to reset local event counters.

Event Filter (formerly level-3 trigger, or LVL3)
The third level of event selection, responsible for reducing the trigger rate and hence the data rate to a value acceptable for permanent storage, roughly 100 Mbyte/s. A processor system which receives events from the Event Builder, these events having been selected by L2_accept or L2_request signals. The Event Filter carries out further processing and analysis and, if accepted, sends the event to data storage for offline analysis.

Event_ID (EVID)
A number which identifies an event uniquely within a run.

FADC
Flash analogue-to-digital converter.

FCAL
Liquid-argon forward calorimeter.

FDR
Final design review.

FE
See Front End.

FE_BCID
A 12-bit number corresponding to the bunch number in the LHC machine, to identify the bunch crossing. It is generated locally in the RODs and reset by BCR. It is used to cross-check against ROD_BCID when identifying event fragments to be read out.

FE_L1ID
A number, of ≥ 4 bits, corresponding to the event number. It is generated locally in the RODs by counting Level-1 Accept signals. It is used to cross-check against ROD_L1ID when identifying event fragments to be read out.

FEM
Front-end Module. Used in the Level-1 Calorimeter Trigger demonstrator programme.

FIFO
First-in first-out. A type of buffer memory.

FIR
Finite-impulse response. A type of digital filter.

FLUKA
Monte Carlo program used to simulate e.m. and hadronic particle showers in the ATLAS detector. Used for radiation calculations.

Forward
The forward trigger refers to the inner (low-radius) part of the TGC trigger system.

FPGA
Field-programmable gate array.

Front End
Shorthand for Front-End Electronics.

Front-end electronics
The detector sub-systems which generate and send trigger data to the Level-1 Trigger System and event data to their RODs for transmission to the data acquisition system.

FWHM

Full-width at half-maximum. A measure of the width of a peak.

GE

Gigabit Ethernet.

GEANT

A general Monte Carlo simulation package for describing detector geometry and tracking particles through detector material. Used to simulate the response of the ATLAS detector.

GTL

Gunning transceiver logic. Operates typically from power supply voltages of +5 V or +3.3 V, with low-amplitude voltage swings of 0.8 V (+0.4 V to +1.2 V) for reduced power dissipation. Signal lines are terminated to +1.2 V, with a receiver comparator reference voltage of +0.8 V. Variants are available with driver slew-rate control to reduce e.m. interference and associated crosstalk.

H8

Test beam at CERN.

Hadronic isolation region

A continuous region in the hadronic calorimeters, behind a potential electron / photon cluster or surrounding a hadron / tau cluster, whose summed transverse energy is used as a veto if it exceeds some (usually low) threshold. This is to help reject jet background.

Hadron/tau cluster

The areas in η–ϕ which are summed in order to compare their transverse energy with hadron / tau trigger thresholds. The elements used in the summing are Trigger Towers.

Hadron/tau de-cluster/RoI region

The areas in η–ϕ which are summed in order to find local maxima in transverse energy, necessary for de-clustering hadron / tau triggers and hadron / tau Regions-of-Interest. The elements used in the summing are Trigger Towers summed over the combined depth of the electromagnetic and hadronic calorimeters.

Hadron/tau window

The areas in η–ϕ which are used for the hadron / tau trigger algorithm, combining hadron / tau Clusters with isolation requirements in both the electromagnetic and hadronic calorimeters.

HEC

Liquid-argon hadronic endcap calorimeter.

High-p_T board

Board of the muon trigger system which implements the high-p_T-muon trigger.

HV

High voltage

IC

Integrated circuit.

I^{2}C

Inter-IC bus. Developed by Philips.

IEEE

Institute of Electrical and Electronics Engineers (USA).

IDR

Interim design review.

ISAJET

Monte Carlo program used to generate simulated proton–proton interactions for various physics processes, based on perturbative QCD plus phenomenological fragmentation models.

Isolation

See Electromagnetic (e.m.) Isolation and Hadronic Isolation.

JEM

Jet and Energy-sum Module of the Level-1 Calorimeter Trigger.

Jet de-cluster/RoI region

The areas in η–ϕ which are summed in order to find local maxima in transverse energy, necessary for de-clustering jet triggers and for jet Regions-of-Interest. The elements used in the summing are Jet Elements.

Jet element

The smallest elements in η–ϕ used to form transverse-energy sums for the jet trigger. The elements are summed over the combined depth of the electromagnetic and hadronic calorimeters. The size of these elements determines the step size for sliding the Jet Windows and Jet De-cluster / RoI Regions.

Jet/Energy-sum Processor

The part of the Level-1 Calorimeter Trigger that carries out jet, missing-E_T and scalar total-E_T triggers.

Jet window

The areas in η–ϕ which are summed in order to compare their transverse energy with jet trigger thresholds. The elements used in the summing are Jet Elements.

JMM

Jet Merger Module of the Level-1 Calorimeter Trigger.

JPM, JPMD

Jet Processor Module. Used in the Level-1 Calorimeter Trigger demonstrator programme.

JTAG (boundary scan)

A technique for loading data into or reading data out from a chip or a module using a single serial line that connects all relevant registers sequentially. The pin count is thus minimized. The standard for this is IEEE 1149.1; JTAG stands for Joint Technology Assessment Group.

L1A

See Level-1 Accept.

L1ID

See FE_L1ID.

L1MT

Level-1 muon trigger (system).

LAr

Liquid argon — refers to the ATLAS calorimeters, all of which except the Tile Calorimeter use liquid argon as a sampling medium.

LCANN
Local CAN node.

LDB
Local data-acquisition block of the muon endcap only trigger.

Level-1 Accept
A signal generated by CTP when an event has met the level-1 trigger criteria, i.e. is a level-1 trigger. It is distributed by the TTC system.

Level-1 calorimeter trigger
The part of the Level-1 Trigger System whose calculations are based on information from the ATLAS calorimeters. Trigger objects are e.m. showers, single hadrons (τ), jets, missing E_T, and total E_T.

Level-1 muon trigger
The part of the Level-1 Trigger System whose calculations are based on information from the ATLAS muon detectors. Trigger objects are high-p_T muons.

Level-1 trigger system (LVL1)
The first level of event selection, responsible for reducing the event rate from the bunch-crossing rate of 40 MHz to no more than 75 kHz averaged over short time periods (e.g. 10 ms), using a fast hardware processor. For accepted events, it issues Level-1 Accept to the front-end electronics and RoI_message to the Level-2 Trigger. The system consists of the Level-1 Calorimeter Trigger, the Level-1 Muon Trigger, and the Central Trigger Processor.

Level-2 trigger system (LVL2)
The second level of event selection, responsible for reducing the trigger rate from about 75 kHz(upgradable to 100 kHz)to a rate acceptable to the Event Filter, 1–5 kHz. It requests and receives RoI data from the ROBs and, after analysing it, sends an L2_accept or L2_reject signal for the event to the ROBs.

Level-3 trigger system
See Event Filter.

LHC
Large hadron collider.

LMB
Local control monitor board of the muon endcap trigger.

LS-Link
Local slave link. A cable link between slave boards in the muon endcap trigger logic via which data are read out.

LSB
Least-significant bit.

LSI
Large-scale integration. Refers to integrated circuits.

LUT
Lookup table.

LV
Low voltage.

LVDS

Low-voltage differential signalling. A high-speed (~500 Mbit/s) low-power, general-purpose interface standard. It features a low voltage swing of ±400 mV (+1.0 V to +1.4 V) with power-supply voltages ranging from +5.0 V down to +2.7 V.

LVL1

See Level-1 Trigger System.

LVL1_A (LVL1_accept)

See Level-1 Accept.

LVL2

See Level-2 Trigger System.

LVL3

See Event Filter.

LVDS

Low-voltage differential signal.

LVPS

Low-voltage power supply.

L1A

See Level-1 Accept.

LynxOS

A real-time UNIX-like operating system from Lynx Real-Time Systems, Inc.

MATCH

A radiation tolerant GaAs transceiver chip used in the barrel muon trigger.

MCU

Micro controller unit.

MCM

Multi-chip module. Distinguished from hybrids by the use of bare integrated-circuit dies.

MDT

Monitored drift tube.

METRAL™

A family of connectors.

MGF

Morris Garages F-type. Four cylinder, sixteen valve, available in a variety of colours.

MIBAK

Backplane that connects modules of the muon-trigger/CTP interface.

MICTP

Module of the muon-trigger/CTP interface that drives data to the CTP.

MIOCT

Module of the muon-trigger/CTP interface that processes data for an area of the muon spectrometer equal to one octant in the azimuthal direction and half the detector in η.

MIROD

Module of the muon-trigger/CTP interface that supplies data to LVL2 (for RoI building) and the ROBs.

MSB
Most-significant bit.

MUCTPI
Muon-trigger–CTP interface.

NIM
A modular system of fast logic, used for trigger systems in particle-physics experiments in simpler times. Still in use, most commonly in test-beam triggering.

NRZ
Non-return to zero. A serial bit coding in which consecutive 1s do not need separators that go to zero.

Octant
A collection of TGCs, symmetric in ϕ, comprising one eighth of a muon TGC trigger plane.

OO
Object oriented. A methodology for writing software.

ORBIT
A signal transmitted by the LHC to the TTC at a fixed point in the LHC cycle. The ORBIT signal is the broadcast to the TTC partitions.

OS9
A real-time operating system, used mainly on single-board VME processors.

PBS
Product breakdown structure

PCB
Printed-circuit board.

PCI
Peripheral component interconnect. An industry-standard bus system used mainly in personal computers.

PDR
Preliminary design review.

PECL
Positive emitter-coupled logic. Refers to ECL circuitry operated with its power-supply voltage offset by +5.2 V, giving logic levels of +3.5 V and +4.3 V.

Pivot plane
Plane of chambers in the RPC or TGC system that defines the RoI position. Equivalent to ‘reference plane’.

PLL
Phase-locked loop.

PMT
Photomultiplier tube.

PP
Preprocessor of the Level-1 Calorimeter Trigger.

PPG
Computer-controlled pulse generator.

PPM

Preprocessor Module of the Level-1 Calorimeter Trigger.

PPS

Prepulse. A signal that can be issued by the CTP a defined duration before an L1A. It is distributed to the front-end electronics by the TTC system and can be used to fire test pulses.

Preprocessor

The part of the Level-1 Calorimeter Trigger that digitizes the calorimeter signals, does bunch-crossing identification, and uses a lookup table to do pedestal subtraction, final E_{T} calibration, and apply a noise threshold.

PRR

Production readiness review.

PPrASIC

Preprocessor ASIC of the Level-1 Calorimeter Trigger.

PPrMCM

Preprocessor multi-chip module of the Level-1 Calorimeter Trigger.

PS-Pack

Patch-panel and slave-board package of the endcap muon trigger.

PYTHIA

Monte Carlo program used to generate simulated proton–proton interactions for various physics processes.

QA

See Quality Assurance.

QAP

Quality assurance plan. See Quality Assurance.

Quality Assurance (QA)

It's paperwork, but not as we know it. See PDR, IDR, FDR, URD, TP, TDR, PBS, WBS, PRR, and every other TLA you can think of.

QCD

Quantum chromodynamics.

R&D

Research and development.

RAL

Rutherford Appleton Laboratory.

Readout Buffer (ROB)

A standard module which receives data from the RODs via standard Readout Links, passes on request a subset of the data to the level-2 trigger, and buffers the data until a Level-2 Trigger decision has been reached whereupon, for accepted events, it transmits the data to the Event Filter.

Readout Driver (ROD)

The last element in the readout chain that is still considered part of the front-end electronics. This module collects one or more data streams from detector elements and merges them into a single stream which is fed via a standard Readout Link into a ROB.

Readout link
The ATLAS-standard data-transmission link between a ROD and a ROB.

Receiver station (formerly waveform-monitoring station)
Units into which analogue signals from the liquid-argon calorimeters are received and split. One output goes to the Level-1 Calorimeter Trigger front-end electronics, and the other is available for waveform monitoring by the calorimeter group.

Reference plane
Plane of chambers in the RPC or TGC system that defines the RoI position. Equivalent to 'pivot plane'.

Region-of-Interest (RoI)
A geographical region of the experiment, limited in η and ϕ, identified by the Level-1 Trigger System as containing candidates for Level-2 Trigger objects requiring further computation. Their data will be further analysed by the Level-2 Trigger System to decide if the event is to be processed further. In the case of B-physics triggers at low luminosity, some RoIs may be defined internally within the Level-2 Trigger.

RemASIC
Readout merger ASIC of the Level-1 Calorimeter Trigger.

Resistive Plate Chamber (RPC)
Muon detector used for triggering in the barrel region of ATLAS.

RMS (r.m.s.)
Root mean square.

ROB
See Readout Buffer.

ROC
Readout controller of the level-1 calorimeter trigger. This controls the transfers of data from the CPM to the ROD.

ROD
See Readout Driver.

ROD_BCID
A 12-bit number corresponding to the bunch number in the LHC machine, to identify the bunch crossing. It is provided by the TTC system to tag event fragments. This number will be reset with each LHC orbit (i.e. every 88.924 s).

ROD_BUSY
A signal to indicate that the ROD is busy, used to inhibit the level-1 trigger.

ROD_L1ID
A 24-bit number corresponding to the event number defined by counting Level-1 Accept trigger signals. It is provided by the TTC system to tag event fragments.

RoI
See Region-of-Interest.

RoI builder
A unit, inside the Level-1 Trigger System, that collects and formats level-1 RoI information for use by the Level-2 Trigger.

RoI_message
A signal sent from the Level-1 Trigger to the Level-2 Trigger which contains the coordinates of the Region-of-Interest.

ROM
Read-only memory.

RPC
See Resistive Plate Chamber.

RXM
Receiver module in the Receiver Stations.

SCT
Semi-conductor tracking detector.

Single-hadron/tau …
See Hadron / tau …

SL
Sector logic of the level-1 muon trigger.

SMM
Sum Merger Module of the Level-1 Calorimeter Trigger.

SRAM
Static random-access memory.

Star Switch
Network (e.g. Ethernet) device connecting a number of nodes in a star topology, allowing multiple simultaneous data transfers between different pairs of nodes at full speed.

SUSY
Super symmetry.

Tau/hadron …
See Hadron / tau …

TCM
Timing Control Module. Used in the Level-1 Calorimeter Trigger Demonstrator Programme.

TLA
Three-letter acronym.

TDR
Technical Design Report. What you are reading!

Thin Gap Chamber (TGC)
Muon detector system used for triggering in the endcap region of ATLAS.

Tile calorimeter (TileCal)
Hadronic barrel calorimeter, using scintillating tiles as active medium.

Timing, Trigger and Control (TTC)
The standard system which provides and distributes trigger signals (e.g. Level-1 Accept), timing signals (e.g. system clock timestamps), and control signals to the various sub-systems of the experiment.

TOF

Time-of-flight.

TP

Technical Proposal.

Trigger

A decision made by a trigger system (LVL1 or LVL2) that a particular event is of potential interest and should be retained at least until the next stage of selection. At LVL1 and LVL2, this decision is based on a subset of data from the event.

Trigger chamber

A generic term for RPCs and TGCs.

Trigger menus

The set of trigger conditions in use at any particular time. They specify a list of items, each with threshold(s) and multiplicity, and the logic to be applied to them. The level-1 trigger menu is implemented in the CTP, and is a set of logical combinations of results from the Level-1 Calorimeter and Muon Triggers.

Trigger tower

The smallest element of calorimeter information used in the Level-1 Calorimeter Trigger. It has dimensions of approximately 0.1×0.1 in η–ϕ and is summed over the full depth of either the electromagnetic or hadronic calorimeter concerned. Note that the number of Calorimeter Cells that have to be summed to produce a trigger tower ranges from 3 to 60.

Triplet

Three layers of muon thin gap chambers.

TRT

Transition-radiation tracking detector.

TSB

Triplet slave board of the muon endcap only trigger.

TTC

See Timing, Trigger and Control.

TTCrx

TTC receiver chip.

TTCvi

TTC VME interface module.

TTL

Transistor–transistor logic. Refers to a family of saturating logic operated from a +5 V supply, with logic levels typically 0–0.3 V and 3.0–5 V for 0 and 1, respectively.

TXM

Transmitter Module. Used in the Level-1 Calorimeter Trigger demonstrator programme.

U

Unit of length used to measure rack height, equal to 1.75 inches.

URD

User Requirements Document. Based (at least loosely) on software standard PSS-05 of the European Space Agency.

URL
Universal Resource Locator (address used by WWW).

USA15
The underground electronics cavern where the level-1 calorimeter trigger, with the exception of the muon barrel and endcap trigger logic, is located.

VCXD
Voltage-controlled oscillator.

VDEC
VLSI Design and Education Centre, University of Tokyo. A CAD centre for Japanese universities.

VHDL
VHSIC (Very high speed integrated circuit) hardware description language. A language for specifying the designs of electronic systems.

VIC
VME inter-crate connection. Used for both a standardized VME inter-crate bus, and for proprietary (Creative Electronic Systems SA) modules that implement it.

VLSI
Very Large Scale Integration. Refers to integrated circuits.

VME; VMEbus
Versa-Module Euro. A crate backplane bus system.

Waveform monitor station
See Receiver Station.

WBS
Work breakdown structure.

Wire-group
Group of wires that form a single readout unit in the TGC detector system.

WWW
World-Wide Web.

ZEBRA
A dynamic memory-management package for FORTRAN programs. Used in the ATLAS offline software.

Zero suppression
Compression of data by removal of values equal to zero.

ZIF
Zero insertion force (socket).

C Appendix: List of figures

D Appendix: List of tables

This document has been prepared with Release 5.5 of the Adobe FrameMaker® Technical Publishing System using the Technical Design Report template prepared by Mario Ruggier of the Information and Programming Techniques Group, ECP Division, CERN, according to requirements from the ATLAS collaboration.

To facilitate multiple author editing and electronic distribution of documents, only widely available fonts have been used. The principal ones are:

Running text:	Palatino 10.5 point on 13 point line spacing
Chapter headings:	Helvetica Bold 18 point
2nd, 3rd and 4th level headings:	Helvetica Bold 14, 12 and 10 point respectively
Figure and table captions:	Helvetica 9 point