# Bunch-Crossing Identification for the ATLAS Level-1 Calorimeter Trigger

**by**

**Ian Paul Brawn.**

A thesis submitted to the Faculty of Science
of The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY.

School of Physics and Space Research,
The University of Birmingham,
June 1996.

# Synopsis


As part of the design studies for the ATLAS level-1 calorimeter trigger, the process of Bunch-Crossing Identification (BCID) has been studied. Various digital BCID algorithms have been investigated using software simulation. The results of these investigations are presented here. They show that BCID can be performed with good efficiency for pulses corresponding to energy deposits down to a few GeV (~4 GeV).

A hardware demonstrator system has been constructed which performs BCID on 36 channels of trigger data. The BCID algorithms are implemented in Field-Programmable Gate-Arrays. This has allowed a number of BCID algorithms to be tested in hardware at a CERN test beam. The design of the BCID demonstrator module is presented here, and the off-line analysis of the data recorded at the test beam is described. The results show BCID being performed with a level of efficiency which is acceptable for the level-1 trigger system.

All of the data analysis described in this thesis is my own work, except where explicitly stated otherwise. The hardware BCID module which is described in chapter 6 was designed entirely by myself.

# Acknowledgements

There are many people I would like to thank....

For their guidance, friendship and supervision I would like to thank my two supervisors: John Garvey and Viraj Perera.

For their guidance, friendship and help in taming the beast described in chapter 6, I would like to thank Bob Hatley and Jason Leake.

For their guidance, friendship and miscellaneous help, I would like to thank the following people (in alphabetical order): Christian Bohm, Richard Carney, Brian Claxton, Adam Connors, John Dowdell, John Dowell, James Edwards, Nick Ellis, Eric Eisenhandler, Norman Gee, Steven Hillier, Murrough Landon, Giles Lewis, Mike Pentney, Steve Quinton, David Rees, Tara Shah, Richard Staley, Alan Watson.

For their lack of guidance, friendship and the occasional beer I would like to thank the following people (in roughly chronological order): Paul 'Spike' Young, Karl 'Stone Face' Hughes, J. Paul 'User' Davies, Paul 'Dr Reef' Newman, Mark '007' Pearce, Manish 'Six Pints' Odedra, David 'Boy Racer' Allen, Eithne 'Partytime' Kane, Paul 'Spaceboy' Hatswell, Mark 'Scary Grin' Smith, Carl 'Action' Mannering, James 'Genghis' Hunt, Huw 'Monkey' Morris, Keith 'Vaudeville' Norman, David 'Blue Rinse' Evans, Steve 'Beers and Bikes' Oglesby, Mark 'Baggy' Venebles, Tim 'Macman' Nicholls, James 'Behaving Badly' Bloomer, Vicki 'Strangely Female For A Physicist' Hudgson.

For their love and financial support I would like to thank my family.

For their music I would like to thank Nick Cave, Eddie Cochran and Leonard Cohen.

That's all. There are more names I could have added above and I mean no slight towards anybody I have left off. To prevent this section extending indefinitely however, I am forced to draw a line somewhere, so I'm drawing it here:

---

Ian Brawn, May, 1996.

"babadalgharaghtakamminarronnkonnbronntonnerronntuonnthunntro
-varrhounawnskawntoohoohoordenenthurnuk!"

– James Joyce, *Finnegans Wake.*

# Contents

# List of Figures

# List of Tables

# 1. Physics Overview

## 1.1. Introduction

The study of particle physics is the study of matter at the most fundamental level; it is the study of the elementary particles of nature and the interactions which they undergo. The current theoretical framework within which this study is conducted is known as the Standard Model (SM) [1, 2]. The adoption of this model has brought a degree of simplicity and consistency to the field of particle physics which contrasts strongly with the situation faced by physicists 30 years ago. At that time experiments were revealing a large ‘zoo’ of different baryons and mesons, all of which were assumed to be elementary particles. Today these particles are seen as being composed from a smaller set of more fundamental particles – the quarks, which together with the leptons form the elementary constituents of matter.

According to the Standard Model, quarks and leptons interact via the exchange of gauge bosons – particles which mediate the fundamental forces of nature. Using quantum field theories to describe these interactions, theorists have been able to predict with great precision many of today’s experimental results. So far, no experimental data have been observed which are inconsistent with the Standard Model. There are many things which the Standard Model does not explain however, and for this reason it is not regarded as a complete theory. The presence of three generations of quarks and leptons, for example, is unexplained by the Standard Model. There are also many free parameters, such as the particle masses, which cannot be predicted by theory and must be measured empirically. Most importantly, there exists a whole section of theory which has never been observed or tested directly. This is the Higgs sector.

Looking beyond the Standard Model, attempts have been made to produce a more comprehensive theory which possess a smaller number of free parameters. Such efforts have led to ‘Grand Unified Theories’ and supersymmetric models. All of these extensions to the Standard Model predict the existence of presently undiscovered particles.

The ATLAS experiment [3], due to start operation at CERN in 2005, has been proposed as a general purpose proton-proton colliding experiment, able to exploit the full range of physics opportunities provided by the Large Hadron Collider (LHC) [4]. Its design and construction are therefore guided by physics goals such as the search for Higgs bosons and supersymmetric particles – areas where our present knowledge, as embodied in the Standard Model, is incomplete. Presented below is a brief survey of the Standard Model and of possible extensions to it.

## 1.2. The Standard Model

The Standard Model consists of a collection of quantum field theories which describe the interactions of a set of elementary particles. These particles can be divided into two types: the elementary fermions, which are the basic constituents of matter, and the gauge bosons, which are the mediators of the fundamental forces of nature. All of these particles are considered to be point-like and structureless.

### 1.2.1. The Elementary Fermions: Leptons and Quarks

Any particle which posseses half-integer spin ($1/2\ \hbar$, $3/2\ \hbar$...) and therefore obeys Fermi-Dirac statistics is labelled a fermion. There are twelve elementary fermions, discounting antiparticles, all of which posses spins of $1/2\ \hbar$. These twelve particles can be grouped into two distinct families – quarks and leptons. Quarks are distinguished from leptons in that they undergo interactions via the strong nuclear force. All of the elementary fermions undergo interactions via the weak nuclear and electromagnetic forces. The elementary fermions are:

$$\text{quarks}: \begin{pmatrix} \mathrm{u} \\ \mathrm{d} \end{pmatrix} \begin{pmatrix} \mathrm{c} \\ \mathrm{s} \end{pmatrix} \begin{pmatrix} \mathrm{t} \\ \mathrm{b} \end{pmatrix} \qquad \text{leptons}: \begin{pmatrix} \nu_{\mathrm{e}} \\ \mathrm{e} \end{pmatrix} \begin{pmatrix} \nu_{\mu} \\ \mu \end{pmatrix} \begin{pmatrix} \nu_{\tau} \\ \tau \end{pmatrix}.$$

As is indicated above, the leptons and quarks can both be grouped into doublets of particles. With regard to the leptons, the lower particle of each doublet carries an electromagnetic charge of -1, whilst the upper particle, i.e. the corresponding neutrino, is neutral. With regard to the quarks, the upper particle of each doublet carries a charge of +2/3, whilst the lower particle carries a charge of -1/3.

Between the different generations of quark and lepton doublets, the coupling strengths of the particles to the strong, weak and electromagnetic forces remain the same. The mass of the particles however, increases with successive generations (the possible exception to this being the neutrino, which may have zero mass). Only quarks and leptons from the lightest generation (u, d, e, $\nu_e$) are found abundantly in nature. The fermions of higher mass are all unstable and decay into these lighter fermions via the weak interaction.

### 1.2.2. The Gauge Bosons

A boson is a particle which possesses an integer spin and obeys Bose-Einstein statistics. The elementary bosons of the Standard Model – the gauge bosons – all possess a spin of $1\hbar$.

All interactions in the Standard Model are described as an exchange of gauge bosons: the electromagnetic force is mediated by the massless photon; the weak nuclear force is mediated by three massive bosons, the $W^+$, $W^-$ and $Z^0$ particles; the strong nuclear force is mediated by an octet of massless gluons, which carry colour charge. The fourth fundamental force, gravity, is not included in the Standard Model. This is due to problems in its formulation as a quantum field theory. Fortunately, being the weakest of the known forces, gravity has no measurable effects on a subatomic scale and accurate calculations can be made using the Standard Model without its inclusion.

### 1.2.3. Gauge Field Theories

Within the Standard Model, interactions between particles are described using gauge field theories. The first such theory, quantum electrodynamics (QED) [5], was formulated in the 1940s. It describes electromagnetic interactions. The great success of QED has led to the development of similar theories describing the strong and weak interactions: quantum chromodynamics (QCD) and the unified electroweak theory respectively.

In gauge field theories, the fundamental fermions are described by complex fields. These fields are used in the construction of the Lagrangian – a mathematical term which describes the energy-momentum dynamic of a system. Being complex, the fermionic fields possess a phase, or gauge, which does not correspond to a physically observable property. It is the basic postulate of gauge field theories that the Lagrangian should be invariant under local transformations of this gauge. In other words, the physics of the system should be unchanged by such gauge rotations.

The fermionic fields themselves are not invariant under arbitrary shifts in gauge. In order for the Lagrangian as a whole to be invariant, new terms must be added to compensate for the way in which the fermionic fields transform. These new terms correspond to gauge fields. In fact, they describe the gauge bosons detailed in the previous section. Hence, by imposing the requirement that the Lagrangian be symmetric under gauge transformations, terms which describe the SM bosons are generated automatically.

Physical predictions are made from quantum field theories using a perturbative approach. This approach has the problem that divergent integrals occur in the calculations, leading to non-physical results (singularities). A crucial property of the gauge field theories used in the Standard Model is that they are renormalisable, allowing this problem to be overcome. Consider, as example, the photo-electric effect as described by QED. Figure 1.1(a) shows the first-order Feynman diagram for this process. All integrals contained in this diagram are well behaved and from it the amplitude of the process can be calculated. The result however, will only be correct to the first order. If a greater level of precision is needed, all higher order processes which can contribute towards the same amplitude must be considered. Contributions from such processes are known as radiative corrections. Figure 1.1(b) shows an example of a third-order radiative correction to the photo-electric effect, where the electron has emitted and re-absorbed a virtual photon is a closed loop.


Figure 1.1. The photo-electric effect: Feynman diagrams showing a) first order and b) third order contributions.

In calculating the amplitude of the process shown in figure 1.1(b), an integration is required over the full range of energy and momentum that can be carried by the virtual photon. This integral is divergent and thus a non-physical answer of infinity is reached. As closed-loop integrals appear in all gauge-field-theory calculations which are performed to greater than leading-order accuracy, it would appear that the theory cannot be used to make accurate physical predictions. It is possible to avoid consideration of these closed-loop interactions however, by a 'renormalisation' of the parameters associated with the Feynman diagram.

Consider for example, the charge of an electron. Any QED process calculated to higher than leading-order accuracy must consider not only the bare electron itself, but all of the possible virtual particles which the electron can emit and re-absorb in a closed-loop process. Figure 1.2 shows some examples. These processes contribute divergent radiative corrections to the charge of the bare electron. When measuring the electron charge the answer yielded must include all such radiative corrections, but it is empirically found to be finite. It follows therefore that the bare electron charge itself cannot be a fixed parameter, but must be divergent in such a way as to cancel out the divergences introduced by the loop graphs. If the bare electron charge is therefore replaced in the Lagrangian by the measured electron charge, all of the radiative corrections are absorbed into this experimentally observed parameter and the divergences disappear from the calculation. This is the process of renormalisation. Similar renormalisation schemes apply to all gauge theories in the Standard Model. It is only because of the property of renormalisability that physical predictions can be made using the SM gauge field theories.

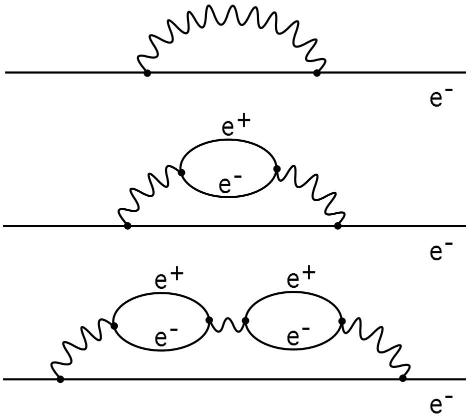


Figure 1.2. Radiative corrections to the bare electron term.

The example of electron charge given above has one further implication: the charge measured experimentally for the electron is not the bare charge – it contains contributions from the virtual particles which surround the electron. If this is so, then the value of electron charge measured should not be constant, but should vary with the separation of the probe particle from bare charge itself (or conversely, upon the momentum-transfer of the interaction). This is in fact the case. To replicate this effect within QED, the value of charge is set as a fixed parameter, but the coupling constant which describes the coupling strength of photons to charged particles is allowed to 'run', increasing as the momentum-transfer of an interaction ($Q$) increases.

### 1.2.4. Quantum Chromodynamics

Following the great success of QED in describing electromagnetic interactions, quantum chromodynamics (QCD) [6] was formulated to describe the strong interaction of quarks. To explain the observed phenomena that quarks are never found in isolation but are always confined within hadrons, it is proposed in QCD that each quark carries a 'colour' charge – either red (R), green (G) or blue(B). As hadrons have no net colour charge, bound states of quarks can only be constructed from 'colourless' combinations of these charges. For example, RGB for the three-quark baryon states, and $R\overline{R}$, $B\overline{B}$, $G\overline{G}$ for the quark-antiquark meson states. This model gives to the quarks an extra three degrees of freedom: the colour charge. The existence of these extra degrees of freedom has been verified experimentally using $e^+e^-$ experiments, where the cross-section for $e^+e^- \rightarrow q\overline{q}$ is three times larger than would be expected if colour charge did not exist.

The three colour charges of the quark form an $SU(3)$ symmetry group. It is the basic postulate of QCD that the Lagrangian describing the coloured quark fields should be invariant under local $SU(3)$ gauge transformations. For this to be accomplished it is necessary to add to the Lagrangian an octet of new gauge bosons, named gluons, which are massless and carry colour charge in a colour-anticolour combination (for example, the $\mathrm{R\overline{B}}$ gluon). The fact that the gluons carry colour charge themselves (cf. photons, which do not carry electromagnetic charge) allows them to interact with each other. One consequence of this is that the effective colour charge involved in a strong interaction increases as the momentum transfer of the interaction ($Q$) decreases. This leads to the phenomena known as asymptotic freedom, where the strength of attractive force felt between two quarks increases with their separation. The quarks are therefore confined in bound states, as required.

When the process of renormalisation is applied to QCD, it is necessary to create a 'running' coupling constant for the strong force, to reproduce the behaviour of asymptotic freedom. For low $Q$ the coupling constant approaches one, and perturbative calculations can no longer be carried out, as high-order processes contribute significantly towards the amplitude of the overall interaction. For high $Q$ interactions however, the coupling constant is lower and sensible predictions can be made using perturbation theory.

### 1.2.5. Electroweak Theory

In the late 1960s Weinberg and Salam successfully formulated the weak interaction as a gauge theory and unified it with QED to produce the electroweak theory [1]. Due to the parity violating nature of the weak interaction, it is required that electroweak theory treats left-handed and right-handed particles differently. The left-handed fermions are grouped into weak isospin doublets. The particles in these doublets have weak isospin ($T$) of 1/2, and third components ($T_3$) of +1/2. The right-handed fermions are grouped into singlets, which possess zero isospin. The fermion isospin eigenstates are thus:

$$\text{doublets (left-handed)}: \begin{pmatrix}\nu_\mathrm{e}\\ \mathrm{e}\end{pmatrix}_L \begin{pmatrix}\nu_\mu\\ \mu\end{pmatrix}_L \begin{pmatrix}\nu_\tau\\ \tau\end{pmatrix}_L \begin{pmatrix}\mathrm{u}\\ \mathrm{d'}\end{pmatrix}_L \begin{pmatrix}\mathrm{c}\\ \mathrm{s'}\end{pmatrix}_L \begin{pmatrix}\mathrm{t}\\ \mathrm{b'}\end{pmatrix}_L$$

$$\text{singlets (right-handed)}: (\mathrm{e})_R\ (\mu)_R\ (\tau)_R\ (\mathrm{u})_R\ (\mathrm{d})_R\ (\mathrm{c})_R\ (\mathrm{s})_R\ (\mathrm{t})_R\ (\mathrm{b})_R.$$

Note that right-handed neutrinos are not found in nature.

All of the fermions undergo gauge transformations according to the $U(1)_Y$ group, which acts upon weak hypercharge, $Y$. The hypercharge of a fermion state is defined by the equation:

$$Q = t_3 + \frac{Y}{2}$$

where $Q$ is electromagnetic charge. In addition, the left-handed fermion states can transform according to the $SU(2)_L$ group of gauge transformations, which act upon weak isospin. The 'step-up' and 'step-down' operators constructed from the generators of this group rotate particles between the two isospin states within each doublet. The quark eigenstates d', s' and b' which appear in the doublets do not correspond to the physical (i.e., mass) eigenstates d, s, b, but consist of linear supersitions of these states, as determined by the Kobayashi-Maskawa matrix[7]. The implication of this is that $SU(2)_L$ rotations allow a cross-coupling of the physical quark generations and such quantum numbers as strangeness and charm are not conserved in interactions described by such rotations.

For the fermionic Lagrangian to be invariant under combined $SU(2)_L \times U(1)_Y$ gauge transformations, four new fields need to be introduced. Three of these, $W^1$, $W^2$, $W^3$, couple to weak isospin, whilst one, $B^0$, couples only to the weak hypercharge. Linear combinations of these fields can be constructed which match the properties of the observed electromagnetic and weak-interaction boson fields. Fields corresponding to the $W^+$ and $W^-$ bosons can be constructed from orthogonal linear combinations of the $W^1$ and $W^2$ fields, whilst fields corresponding to the $Z^0$ boson ($Z$) and the photon($A$) can be constructed from orthogonal linear combinations of the $B^0$ and $W^3$ fields:

$$\begin{aligned} Z &= W^3 \cos\theta_W - B^0 \sin\theta_W \\ A &= W^3 \sin\theta_W + B^0 \cos\theta_W \end{aligned}$$

where $\theta_w$ is the weak mixing angle.

Unfortunately there is one major problem with the theory presented so far: the particles described are all massless. If mass terms are added to the Lagrangian 'by hand', symmetry under $SU(2)_L \times U(1)_Y$ gauge rotations is lost and the theory becomes unrenormalisable. As such a theory cannot be used to calculate physically observable quantities, it is of no use. To bypass this problem, Weinberg and Salam proposed the use of the Higgs mechanism [8] to introduce mass terms into the Lagrangian.

### 1.2.6. The Higgs Mechanism

The Higgs Mechanism generates mass terms for the particles in the electroweak Lagrangian by a process of spontaneous symmetry breaking. The motivation for using this method of mass generation is that the Lagrangian remains renormalisable. In the simplest model of the Higgs mechanism, four new scalar fields are added to the Lagrangian, arranged in an isospin doublet with hypercharge 1:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \qquad \text{where} \qquad \begin{aligned} \phi^+ &\equiv (\phi_1 + i\phi_2)^{-\frac{1}{2}} \\ \phi^0 &\equiv (\phi_3 + i\phi_4)^{-\frac{1}{2}} \end{aligned}$$

The properties of the Higgs doublet, $\phi$, are defined such that the Lagrangian now includes a self-interaction term for $\phi$ and, crucially, this term has a non-zero ground state. In fact, there exist a continuous set of degenerate ground states, $|\phi_0|$, which differ from one another in the phase of their weak isospin.

Any Lagrangian which contains quantum fields cannot be solved exactly but must be calculated using perturbation theory. To apply perturbation theory to those terms in the Lagrangian containing $\phi$, it is necessary to pick one definite ground state from the set $|\phi_0|$ and parameterise $\phi$ in terms of fluctuations about this vacuum state. It is the act of defining $\phi_0$ which is referred to as "spontaneous symmetry breaking". Whilst the Lagrangian is invariant under $SU(2)_L \times U(1)_Y$ transformations, the vacuum state $\phi_0$ is not – it has a well-defined isospin gauge.

If $\phi_0$ is defined to have a particular gauge known as the unitarity gauge, it is found to possess zero electromagnetic charge. This means that whilst $SU(2)_L \times U(1)_Y$ symmetry is broken by $\phi_0$, the residual $U(1)$ symmetry of QED is retained. Defining $\phi_0$ to possess the unitarity gauge has other desirable consequences too. If the Higgs doublet in the Lagrangian is expressed in terms of a fluctuation about $\phi_0$, terms are generated which correspond to mass terms for the $W^+$, $W^-$ and $Z^0$ bosons, whilst the photon remains massless due to the residual $U(1)_{em}$ symmetry. Of the scalar fields which comprise the Higgs doublet, three disappear. The fourth scalar field itself acquires a mass term and is found to describe a massive, spin-zero boson. The Higgs mechanism therefore predicts a new particle: the Higgs boson. This physical picture arises from the choice of the unitarity gauge for the vacuum state $\phi_0$. If some other gauge is chosen the resultant Lagrangian is totally equivalent to that generated using the unitarity gauge, but does not contain terms which manifestly relate to physically observable bosons.

The above procedure generates mass terms only for the gauge bosons. The next step is to generate mass terms for the fermions. In order to do this, it is necessary to add to the electroweak Lagrangian terms which describe the interaction of the fermions with the Higgs doublet, $\phi$. When the symmetry of $\phi$ is broken, these terms resolve into two components, one corresponding to the mass term of the relevant fermion, the other describing the coupling of that fermion to the Higgs boson. The fermions are found to couple to the Higgs boson with a strength proportional to their mass.

The addition of the Higgs doublet to the electroweak Lagrangian and the subsequent prediction of the massive Higgs boson completes the specification of the Standard Model. The Higgs sector however, is the least well-determined sector of this model. The only mass relationship directly predicted by the Higgs mechanism is the ratio between the W and Z masses ($M_W$ and $M_Z$ respectively):

$$M_W = M_Z \cos\theta_W$$

which is found to agree closely with experiment.

The masses of the fermions appear in Higgs theory only as free parameters, as does the mass of the Higgs boson, $m_H$. It is possible to place some constraints on the value of $m_H$ however.... An upper limit arises as $m_H$ is related to the self-interaction strength of the original Higgs doublet. If meaningful perturbation expansions are to be obtained from the theory, $m_H$ must be less than ~1 TeV. A lower limit for $m_H$ can be derived by considering the stability of the electroweak vacuum. For values of $m_H$ below a certain level, the Higgs field develops a minimum lower than the electroweak ground state, and the electroweak model becomes unstable. To avoid this, $m_H$ must be greater than ~120 GeV[9].

The Higgs mass is also related to the empirically observed masses of the top quark ($M_{top}$) and W boson ($M_W$): $M_W$ receives radiative corrections which are dominated by top-quark and Higgs-boson loops. It can be seen from figure 1.3 however that $M_{top}$ and $M_W$ have not yet been measured with a sufficient accuracy to constrain $m_H$[10]. Empirically, a lower limit of ~60 GeV has been placed on the mass of the Higgs boson by LEP. If LEP2 makes no discoveries in the Higgs sector, this limit will increase to 90 GeV [11].

Figure 1.3. The calculated dependence of $M_W$ on $M_{top}$ in the Standard Model using the LEP value of $M_Z = 91.1895 \pm 0.0044$ GeV, shown for Higgs masses of 50, 250, 5000 and 1000 GeV. The width of the band for each Higgs mass does not include the uncertainty on $\alpha(M_Z)$. The data point is at $M_{top} = 174 \pm 17$ GeV, $M_W = 80.23 \pm 0.18$ GeV, where the value for $M_{top}$ is from the CDF mass fit, while the value for $M_W$ is from direct $M_W$ measurements alone.

## 1.3. Beyond the Standard Model.

### 1.3.1. Grand Unified Theories

Within a single theoretical framework, the electroweak theory describes both the electromagnetic and weak interactions. The extent to which it actually 'unifies' these interactions is debatable however. The $SU(2)\times U(1)$ symmetry group which the electroweak theory employs is a product of two separate gauge transformations, $SU(2)$ and $SU(1)$, each of which has its own coupling strength ($g'$ and $g$ respectively). The ratio of these two coupling strengths is not predicted by electroweak theory and must be measured empirically.

To obtain a theory which provides a relationship between $g'$ and $g$, the $SU(2)$ and $U(1)$ symmetry groups must be embedded in a larger, unified gauge group, $G$. A requirement can be made that this group also includes the $SU(3)$ symmetry group of colour. The relationship between the groups is expressed symbolically as

$$G \supset SU(3)\times SU(2)\times U(1).$$

By requiring a Lagrangian to be invariant under gauge transformations of the group $G$, a 'Grand Unified Theory' (GUT) can be constructed which describes all particle interactions. In such a theory, the two coupling constants of electroweak theory, $g$ and $g'$, and the strong coupling constant, $g_s$, are related to a universal coupling constant, $g_c$. The precise relationship is specified by the definition of the group $G$.

The coupling strengths $g'$, $g$ and $g_s$ are all 'running constants': their values vary with the momentum transfer ($Q$) of an interaction. Figure 1.4 shows the behaviour of the three coupling constants, as predicted by the electroweak theory and quantum chromodynamics. It can be seen that at some large energy, $M_X$ (~$10^{15}$ GeV), all three interactions could possess the same coupling strength. In GUTs, interactions above this point are described as transformations according to the unified symmetry group, $G$. Below the unification energy, $G$ is spontaneously broken into $SU(3) \times SU(2)_L \times U(1)_Y$.

There are many possible choices for the unified group $G$. In 1974 Georgi and Glashow showed that the smallest possible candidate for $G$ is an $SU(5)$ symmetry group [12]. The invariant Lagrangian associated with this group has 24 gauge bosons. After spontaneous symmetry breaking these resolve into the 12 gauge bosons of the Standard Model plus an additional 12 'leptoquarks'. The leptoquarks have masses of the order of $M_X$, carry colour charge and couple leptons to quarks. They thus give rise to interactions which violate lepton and baryon numbers, such as the decay of the proton. Figure 1.5 shows a hypothetical mechanism for the decay of a proton via a leptoquark. The Georgi-Glashow GUT predicts a lifetime of ~$10^{30}$ years for the proton, which compares badly to a measured lower limit for the proton lifetime of $10^{32}$ years.

This provides the first indication that the minimal $SU(5)$ GUT is not an entirely satisfactory model.

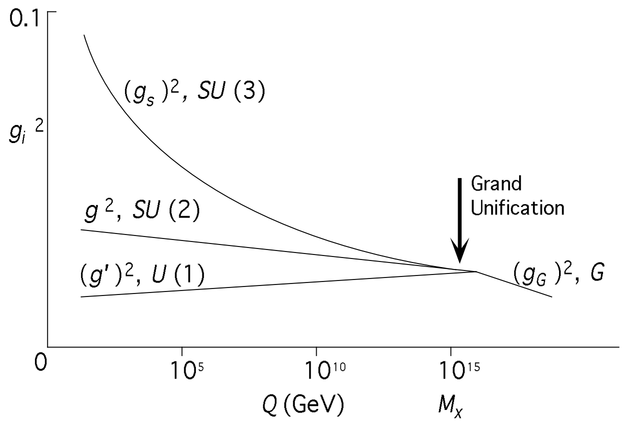


Figure 1.4. The variation of the coupling constants with $Q$, showing the speculative grand unification of strong and electroweak interactions at $Q \sim M_x$.

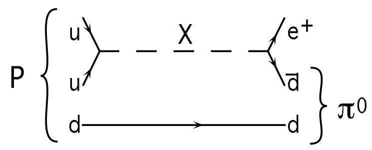


Figure 1.5. The decay of a proton via a leptoquark, X.

Symmetry breaking in the Georgi-Glashow GUT is a two-stage process. In the first stage, the $SU(5)$ symmetry is broken down via the Higgs mechanism to $SU(3) \times SU(2)_L \times U(1)_Y$. It is in this process that the leptoquarks acquire mass. Twelve Higgs bosons are generated by this process, all of which have mass $\sim M_X$. These bosons couple to the leptoquarks but not to the fermions. The second stage of symmetry breaking parallels that already discussed for the Standard Model: the $SU(3) \times SU(2)_L \times U(1)_Y$ symmetry is spontaneously broken to leave a residual $U(1)_{em}$ symmetry and the fermions and SM gauge bosons acquire mass. In this instance however, five Higgs bosons are predicted by the theory, as opposed to the single Higgs boson predicted by the Standard Model.

The Lagrangian associated with the Georgi-Glashow GUT thus contains twelve 'heavy' Higgs of mass $\sim 10^{15}$ GeV and five 'light' Higgs of mass ~100 GeV. The lighter Higgs particles inevitable receive radiative corrections to their mass from the heavier Higgs particles. As the mass of the heavier Higgs particles is so much greater than the mass of the lighter particles, these radiative corrections result in the masses of the W, Z and light Higgs bosons all becoming unstable and acquiring values in the order of $M_X$. This is know as the 'Hierarchy problem'. To avoid unstable W, Z and light Higgs masses it is necessary to cancel radiative corrections with a precision of 1 part in $10^{13}$. Such precise cancellations are not possible in the minimal Georgi-Glashow GUT, as the fine tuning of high order radiative corrections detunes the lower order corrections. This amounts to an insuperable problem which is common to many GUTs. Only supersymmetry theories appear to offer a solution.

### 1.3.2. Supersymmetry

There are many supersymmetry (SUSY) theories, all of which propose a symmetry between the fundamental fermions and the fundamental bosons. New supersymmetric partners are postulated for all of the elementary particles: leptons and quarks are partnered by slepton and squark bosons, the gauge bosons are partnered by the gaugino fermions, and the Higgs bosons are partnered by the fermionic sHiggs particles [13].

With regard to the Hierarchy problem, the radiative corrections applied to the masses of the light Higgs particles in SUSY theories must include contributions from the coupling of the light Higgs particles to the heavy Higgs bosons *and* the heavy sHiggs fermions. As the radiative corrections to the bare mass terms due to virtual boson loops and virtual fermion loops are of opposite signs, these two sets of contributions cancel out exactly. The masses of the light Higgs particles and SM gauge bosons therefore remain stable.

The symmetry between bosons and fermions proposed by SUSY theories must be broken at some stage, to explain why the new particles predicted by the theories have not yet been observed. If SUSY theories are to solve the Hierarchy problem, then the energy scale at which supersymmetry is broken cannot be much greater than that of the electroweak Higgs sector. The new SUSY particles, named sparticles, must therefore have masses no larger than ~1 TeV. With regard to the five light Higgs bosons predicted by SUSY theories, the lightest of these is constrained to have a mass of less than ~120 GeV [9]. It is therefore possible that it will be discovered at LEP2.

The exact phenomenology of the sparticles varies between SUSY theories, but in most models they can only be produced in pairs. This leads to the idea of a quantum number $R$, which is conserved multiplicatively, like parity. Particles are assigned an $R$ value of +1 and sparticles are assigned an $R$ value -1. The conservation of $R$ implies that at least one sparticle, the lightest, must be stable. If this particle exists it will pass through detectors unobserved, but its presence in the decay products of any interaction should be evident due to a dramatic imbalance in the momentum of the observed particles. To date, no evidence has been found to support the existence of SUSY particles.

## 1.4. Physics at the Large Hadron Collider

The Large Hadron Collider (LHC) [4] is the proposed particle accelerator to be built at CERN in the existing LEP tunnel. It will collide two proton beams with a centre of mass energy of 14 TeV, allowing the study of particle interactions at an energy scale never before observed. Activities which will be conducted at the LHC include the search for leptoquarks, the search for the lightest stable sparticle and the search for evidence of quark compositeness. Of primary interest though, is the investigation of spontaneous symmetry breaking at the electroweak scale.

Possible manifestations of spontaneous symmetry breaking include the single Higgs boson predicted by the Standard Model and the quintet of Higgs bosons predicted by minimal supersymmetric extensions to the Standard Model. The search strategy for the SM Higgs boson (H) is well known. In order to cover the full possible mass range 80 GeV < $m_H$ < ~1 TeV, detectors at the LHC will need a sensitivity to a variety of final states:

$H \rightarrow \gamma\gamma$, for 90 GeV < $m_H$ < 150 GeV;

$H \rightarrow ZZ^* \rightarrow 4l^\pm$, for 130< $m_H$ <180 GeV;

$H \rightarrow ZZ \rightarrow 4l^\pm$, $2l^\pm 2\nu$, for 180< $m_H$ <800 GeV;

$H \rightarrow WW, ZZ \rightarrow l^\pm\nu$ 2jets, $2l^\pm 2\nu$, $4l^\pm$ for $m_H$ < ~1 TeV;

where $l$ indicates an electron or muon. In addition, the search for the five Higgs bosons predicted by minimal supersymmetric theories ($H^\pm$, h, H and A) will require a sensitivity to such processes as:

$A \rightarrow \tau^+\tau^- \rightarrow e\mu + \nu$'s
$\quad \rightarrow l^\pm$ + hadrons + $\nu$'s;

$H^\pm \rightarrow \tau^\pm\nu$
$\quad \rightarrow$ 2 jets.

The channels through which the decay of the Higgs bosons can be observed are limited to those visible above the dominant hadronic background at the LHC. As the expected cross-sections for such processes are small for both the low ( < 180 GeV) and the very high mass range, it is proposed to run the LHC at a maximum luminosity of $10^{34}$ cm$^{-2}$s$^{-1}$.

Further models of symmetry breaking in the electroweak sector exist, all of which predict the existence of hitherto undiscovered particles, such as W' and Z' gauge bosons, with masses in the range of 1 TeV. By allowing a study of this energy scale, the LHC will inevitably provide insight into the mechanism of symmetry breaking, whether through the discovery of new particles or the revelation of their absence.

For the first few years of operation, the LHC will run at less than peak luminosity: $10^{33}cm^{-2}s^{-1}$. During this period, studies of top-quark and bottom-quark physics will be of primary interest. In addition, the LHC will able to collide beams of heavy ions, such as lead, with a centre of mass energy in excess of 1250 TeV. This will advance the study of heavy ion physics and aid the search for a quark-gluon plasma.

Three collaborations have so far presented Technical Proposals for experiments at the LHC: ATLAS [14], CMS [15] and ALICE [16]. Both ATLAS and CMS will be general-purpose proton-proton experiments, whilst ALICE will be a heavy ion experiment. One further experiment has been proposed, for the purpose of studying B-physics: LHC-B [17]. This thesis concerns the design of the ATLAS experiment.

# 2. The ATLAS Detector

## 2.1. Introduction

It is the intention of the ATLAS collaboration [14] to provide CERN with a general-purpose proton-proton experiment, able to exploit the full range of physics opportunities presented by the LHC. The ATLAS detector has therefore been designed for optimum performance over a broad range of physics tasks. The tasks of primary interest are those mentioned in section 1.4, namely the Higgs search, the search for particles predicted by extensions to the Standard Model, the search for evidence of quark compositeness, and the study of top-quark and bottom-quark physics. It is expected that by optimising the sensitivity of the detector to the broad range of physics signatures required to study these processes, the potential of the detector to discover any new, unexpected physics will also be optimised.

Figure 2.1 shows a diagram of the proposed ATLAS detector. Located at the centre of the experiment is the inner tracking detector, which consists of an inner volume of high-resolution pixel and strip detectors, surrounded by continuous straw-tube tracking detectors. The inner detector is contained within a super-conducting solenoid which provides a 2 T magnetic field over this region of the detector. Outside the superconducting solenoid lie the ATLAS calorimeters. Electromagnetic calorimetry is performed in three calorimeters (a barrel and two endcaps), all of which use liquid-argon (LAr) technology. Hadronic calorimetry is performed in the two LAr endcaps and in the scintillator tile calorimeter, which surrounds the LAr calorimeters for the full length of the detector.

A second system of magnets lies outside the tile calorimeter. This system consists of three superconducting air-core toroids: a barrel section and two end-caps. Each toroid is built from eight flat coils, positioned around the calorimeter with an eight-fold symmetry. The air-core toroid system provides the magnetic field for the muon spectrometer, which is built from muon chambers placed throughout the magnetic volume. It is the muon spectrometer which defines the overall dimensions of the ATLAS detector: 11 m in radius with an overall length of 42 m.

In sections 2.3 to 2.6, the various subsystems of the ATLAS detector are described. These descriptions are preceded by a note on the coordinate systems used in this thesis.

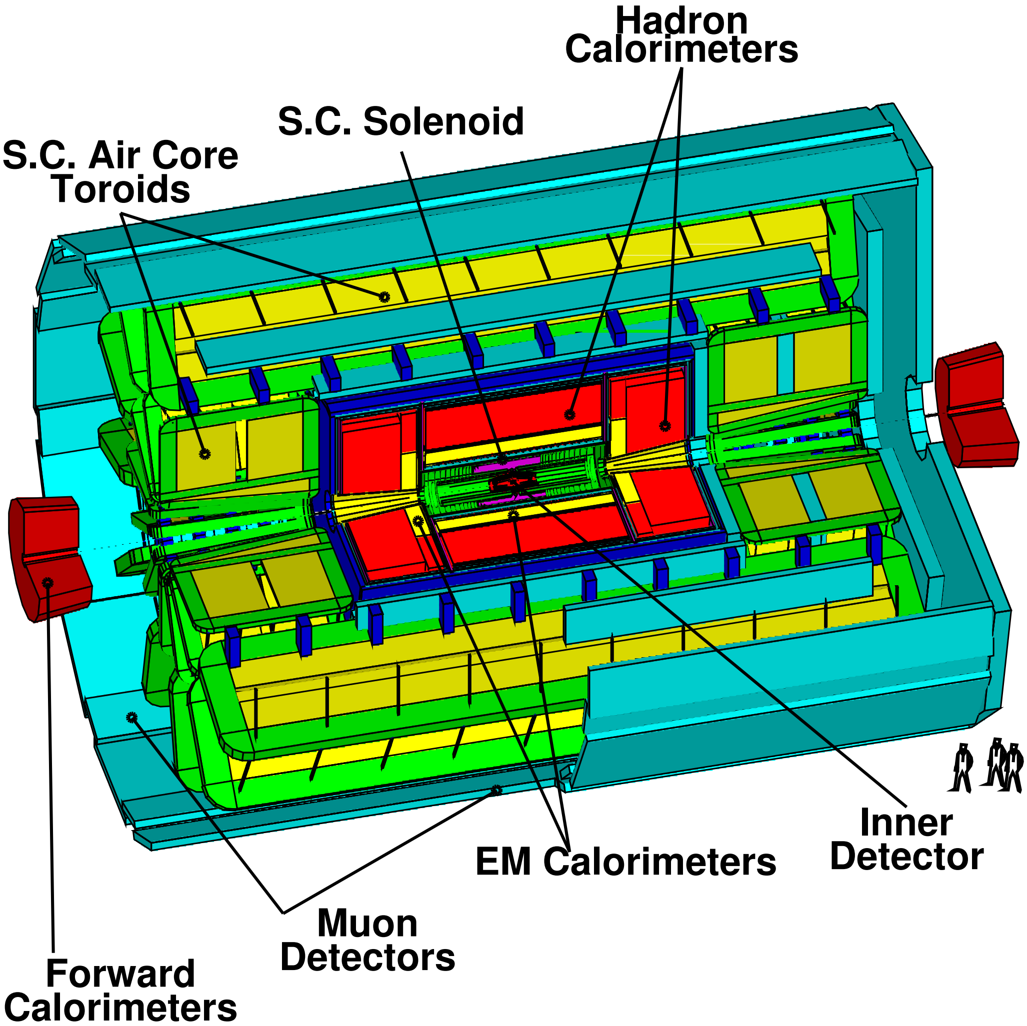


Figure 2.1. The ATLAS Detector.

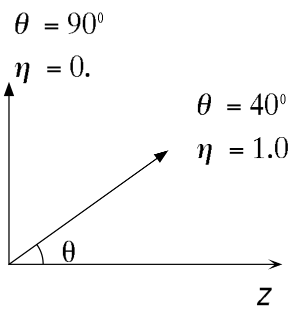


Figure 2.2. The relationship of pseudo-rapidity ($\eta$) to the polar angle ($\theta$) and the $z$ axis.

## 2.2. Coordinate Systems

Three coordinate systems are used here to describe the ATLAS detector:

- Cartesian coordinates ($x, y, z$), where $z$ lies along the beam axis, $x$ lies perpendicular to this in the horizontal direction, and $y$ points vertically upward;
- cylindrical coordinates ($R, \phi, z$), where $z$ lies along the beam axis, $\phi$ is the azimuth angle about this axis, and $R$ is the perpendicular distance from the $z$-axis;
- pseudo-rapidity–azimuth coordinates ($r$, $\eta$, $\phi$), where $r$ is the distance from the origin (defined as the interaction point), $\phi$ is the azimuth angle about the beam axis, and $\eta$ is pseudo-rapidity, defined by the equation

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right).$$

The polar angle $\theta$ is defined in figure 2.2. Pseudo-rapidity is the high-energy approximation of the Lorenz-invariant term *rapidity*.

## 2.3. The Inner Detector

At the highest LHC luminosity, the basic functions of the inner detector are:

- the efficient reconstruction of particle tracks within the range $|\eta| < 2$,
- efficient momentum measurements for leptons over the range $|\eta| < 2$,
- the tagging of b-jets,
- the measurement of the $z$ coordinate of the primary vertex of interactions.

At lower luminosities, the inner detector performs the additional task of reconstructing secondary vertices from b and $\tau$ decays.

The inner detector is contained within a cylinder of radius 1.15 m and length 6.80 m. It is divided into three sections: a barrel and two forward regions. In the barrel section, layers of detectors run parallel to the beam pipe, forming cylinders around the interaction area. In the forward regions, layers of detectors are arranged in discs perpendicular to the beam pipe. The transition from barrel to disc geometry occurs at $|\eta| \sim 1$. Figure 2.3 shows a quarter-section of the inner detector.

The inner detector employs a number of different subdetector systems. Table 2.1 lists the basic design parameters for each of these systems. The semiconductor pixel detectors provide the greatest spatial resolution of all the subdetectors and they are positioned closest to the beam pipe. For larger-area precision tracking, Si–based and

Ga-As–based semiconductor strip detectors are used. Each layer of strip detector actually consists of two or three sets of strips, running at small angles to each other. In the barrel, each strip measures the $R$ and $\phi$ coordinates of a hit directly, whilst the $z$ coordinate is obtained by correlating the results between sets of strips. In the forward regions, it is the $z$ and $\phi$ coordinates which are measured directly. A lower precision in the $z$ and $R$ coordinates than in $\phi$ is acceptable, as $z$ and $R$ lie in the non-bending plane of the magnet. The number of strip-detector layers in both the barrel and forward regions is such that each particle track in the range $|\eta| < 2.5$ must pass through at least four strip detectors.

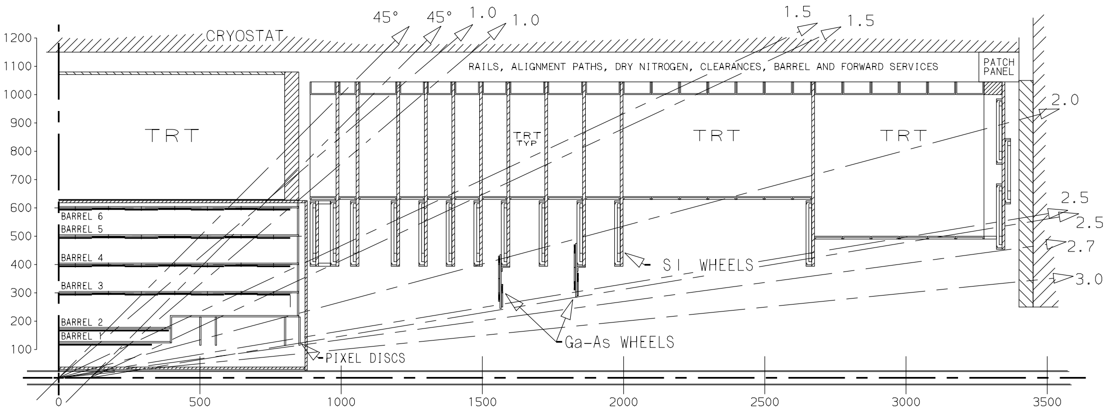


Figure 2.3. A quarter section of the ATLAS Inner Detector. The dimensions on the horizontal and vertical axes are given in millimetres.

One further precision layer is included in the design of the ATLAS detector, which is intended to be present for only the initial low-luminosity running period of the LHC. This layer is intended to increase the vertex-finding capability of the detector in this period, and hence is know as the ‘vertexing’ layer. It is foreseen as being constructed of either semiconductor pixel or strip technology and situated as close to the beam pipe as radiation levels and track-finding considerations will allow (~5 cm). For the higher luminosity running period of the LHC, the vertexing layer will be removed.

The number of pixel and strip layers that can be employed by the inner detector is limited due to their need for local electronics. As well as increasing the power dissipation and overall cost of the detector, these electronics increase significantly the thickness of the inner detector in radiation lengths. For this reason a continuous straw-tube tracker (TRT) is used at the outer radii of the inner detector.

The straw-tube detectors which compose the TRT are basically drift tubes. In the barrel section they run parallel to the beam pipe and thus measure the $R$ and $\phi$ coordinates of a hit. In the forward regions the straws are oriented axially, measuring $z$

and $\phi$ directly. The $R$ coordinate can be measured indirectly by using the first and last straws hit by a particle to determine the trajectory of the track. The spatial resolution of this tracker is lower than that of the precision pixel and strip layers, but this is acceptable at larger radii as the track density is lower. Each track in the range $|\eta| > 2.5$ must cross at least 36 straws, allowing tracks to be followed in a continuous way through the outer volume of the inner detector. This has obvious advantages for pattern recognition. In addition, the layers of straw tubes are interleaved with radiators which produce X-ray emissions from highly relativistic particles. These X-rays are detected by the straw tubes enhancing the electron identification capability of the inner detector.

| **System** | **Position** | **Resolution σ (μm)** | **Channels ($10^6$)** | $\eta$ **coverage** |
|---|---|---|---|---|
| Pixels | (1 low-lum. barrel layer)<br>2 barrel layers,<br>4 forward discs / side | ($R\phi = 14$, $z = 87$)<br>$R\phi = 14$, $z = 87$<br>$R = 87$ | (12)<br>140 | (± 2.5)<br>2.5 |
| Si Strips (barrel) | (1 low-lum. barrel layer)<br>4 barrel layers, | ($R\phi = 10$, $z = 20$)<br>$R\phi = 15$, $z = 770$ | (0.03)<br>2.9 | (± 2.1)<br>1.4, |
| Si Strips (forward discs) | 8 forward discs and<br>1 wheel on each side | $R\phi = 35$ | 1.5 | 1.4 – 2.5 |
| Ga-As strips | 2 forward discs and<br>1 wheel on each side | $R\phi = 10$,<br>$R = 1200$ | 0.8 | 2.0 – 2.5 |
| Straw tube tracker (TRT) | Axial barrel straws,<br>Radial forward straws,<br>36 straws per track | $R\phi = 170$<br>per straw | 0.1<br>0.32 | ± 2.5 |

Table 2.1. Parameters of the inner detector. The values for the two low-luminosity vertexing layer options are given in parentheses.

## 2.4. Calorimetry.

The purpose of the ATLAS calorimeter is to reconstruct the energy of photons, electrons and jets and provide a measurement of missing transverse energy (missing-$E_T$). Figure 2.4 shows a diagram of the proposed calorimetry system. It is contained within a cylinder of outer radius 4.25 m and total length 12.20 m. The inner barrel and endcap sections of the calorimeter all use liquid argon (LAr) technology and are housed inside three cryostats. These cryostats are surrounded by a scintillator tile calorimeter. Electromagnetic calorimetry is performed within the LAr sections of the calorimeter, whilst hadronic calorimetry is performed mainly within the tile calorimeter, with some hadronic calorimetry being performed in the LAr endcaps.

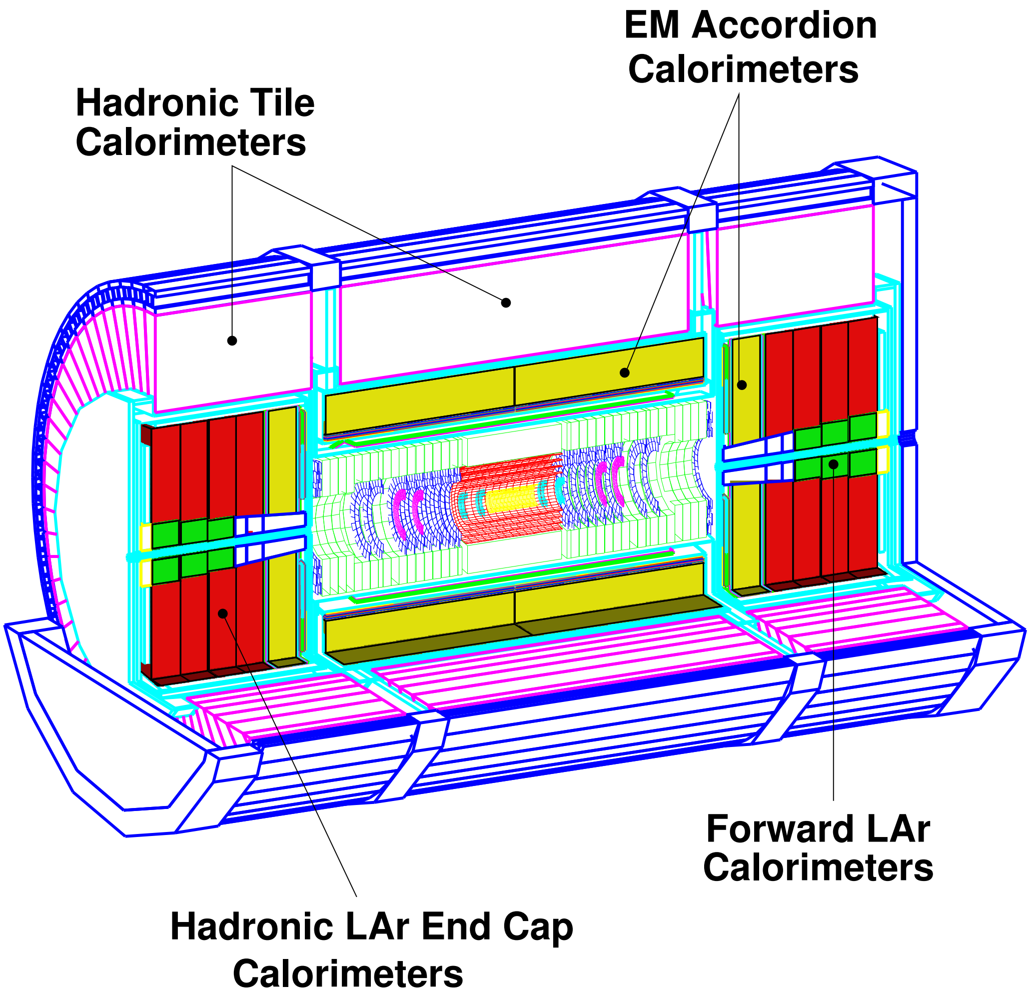


Figure 2.4. ATLAS Calorimetry.

### 2.4.1. The Electromagnetic Calorimeter

The barrel cryostat which surrounds the inner detector cavity contains the barrel electromagnetic calorimeter and the solenoidal coil for the 2 T inner magnet. This solenoid lies before the calorimeter and is integrated with the inner wall of the cryostat to minimise the thickness of material in front of the calorimeter. The electromagnetic barrel calorimeter uses lead absorber plates in a liquid argon ionisation medium. There are 1024 absorber plates, which are of an accordion geometry and are arranged radially in $\phi$. Copper readout electrodes are centred between each pair of lead plates. The barrel electromagnetic calorimeter extends to $| \eta | = 1.4$, with electromagnetic calorimetry at high pseudo-rapidity ranges being provided by the endcap calorimeters. The electromagnetic endcap calorimeters also use lead absorber plates with an accordion geometry. Each electromagnetic endcap calorimeter consists of two coaxial wheels of plates, which radiate out from the beam axis like spokes. There are 768 plates in the outer wheel, with half this number in the inner wheel.

Many of the important processes at the LHC, such as the decay of a Higgs particle into two photons, place stringent requirements on the performance of the electromagnetic calorimeter. In the search for rare processes, such as the decay $H \rightarrow \gamma\gamma$, it is essential to maintain excellent acceptance. A large rapidity coverage is therefore required. Simulation studies have shown that the coverage provided by the electromagnetic calorimeter, of $| \eta | < 3.2$, allows high-efficiency reconstruction of essentially all the interesting high-transverse-momentum (high-$P_T$) physics channels.

In order to be sensitive to all of the processes of interest at the LHC, good electron and photon identification and reconstruction are required over a large range of energies. For the efficient tagging of b-jets via the search for leptonic b decays, it is necessary to reconstruct electrons with transverse energies down to 2 GeV. At the other end of the scale, sensitivity to processes such as $Z' \rightarrow ee$ and $W' \rightarrow e\nu$ requires electrons with energies of up to 5 TeV to be reconstructed. A 16-bit dynamic range is required to adequately reconstruct photons and electrons within these limits. An energy resolution of $\Delta E/E \leq 10\ \%/\sqrt{E} + 1\%$ is thought to provide adequate sensitivity to all of the physics processes of interest at the LHC. At design luminosity and $|\eta| = 0$, the total noise contribution to the resolution of an electromagnetic shower has an R.M.S. of 400 MeV.

Optimising the segmentation of the calorimeter involves balancing performance issues such as the efficiency of electron and photon identification, spatial resolution and the level of pile-up, against cost and technical restraints such as the routing of the signals from inside cryostats. The dominant background to electrons and photons are jets. These produce showers in the electromagnetic calorimeter which are longer and broader than those of isolated electrons and photons. The longitudinal and transverse segmentation of the calorimeter gives a measure of the shape of a shower and is therefore an essential tool in jet rejection. The segmentation of the ATLAS electromagnetic calorimeter is $\Delta\eta \times \Delta\phi \approx 0.025 \times 0.025$ in the second sampling and $\Delta\eta \times \Delta\phi \approx 0.025 \times 0.05$ in third sampling, for $|\eta| < 2.4$. At larger pseudo-rapidity a coarser segmentation of $\Delta\eta \times \Delta\varphi \approx 0.05 \times 0.05$ is used. This segmentation, when combined with data from the hadronic calorimeter, gives a jet rejection factor of ~2000, compared to a target figure of $10^4$. The remaining background is mainly due to high-$P_T$ isolated $\pi^0$s from jet fragmentation. To reduce this, a fine-grained, position-sensitive device is required to identify the two merged photons produced by a $\pi^0$ decay. This is achieved by segmenting the first sampling of the barrel electromagnetic calorimeter into narrow strips of $\Delta\eta \times \Delta\varphi \approx 0.003 \times 0.1$. This first sampling is effectively an integrated pre-shower detector.

The electromagnetic calorimeter is preceded in the barrel by a pre-sampler detector layer, located immediately behind the cryostat inner wall. The pre-sampler has two functions: to correct for the energy which a particle loses in the material in front of the electromagnetic calorimeter, and to assist in measuring the direction of electromagnetic showers.

### 2.4.2. The Hadronic Calorimeter

The hadronic scintillator-tile calorimeter surrounds the cryostats which house the LAr sections of the calorimeter. The tile calorimeter itself is contained in an outer support cylinder which also acts as the main flux return for the 2 T inner solenoid. It is divided into three sections: a barrel and two extended barrels. These sections are composed of plastic scintillator plates embedded in an iron absorber. The tiles are arranged in planes perpendicular to the beam axis, with fibre optic cables being used to transmit the scintillation light to photo-multiplier tubes.

The scintillator-tile calorimeter covers the pseudo-rapidity range $|\eta| < 1.6$. At higher pseudo-rapidity, where a higher radiation resistance is required, hadronic calorimetry is based on LAr technology. Hadronic calorimeters constructed from flat plates of copper in a liquid argon ionisation medium sit in the endcap cryostats behind the electromagnetic calorimeter sections, extending the coverage of the hadronic calorimeter to $|\eta| = 3.2$.

The major tasks of the hadronic calorimeter are to identify jets and measure their energy and direction, to contribute towards the measurement of the total missing-$E_T$, and to enhance the particle identification capability of the electromagnetic calorimeter by measuring such quantities as isolation and leakage. The reconstruction of both jet energy and missing-$E_T$ are limited by intrinsic effects. Uncertainty is introduced into the measurement of jet energy via such effects as jet fragmentation and the loss of energy into non-interacting particles (neutrinos and muons), whilst pile-up noise contributes a significant uncertainty to the measurement of missing-$E_T$. It is therefore acceptable for the hadronic calorimeter to possess a lower energy resolution than the electromagnetic calorimeter. It is believed that a resolution of $\Delta E/E = 50\%/\sqrt{E} + 3\%$ for the range $|\eta| < 2.4$, and $\Delta E/E = 100\%/\sqrt{E} + 10\%$ for $|\eta| < 2.4$, provides an adequate sensitivity to all of the physics processes of interest.

The search for processes such as $H \rightarrow WW \rightarrow l\nu$ + 2jets requires the reconstruction of the decay of a high-$P_T$ W to jet pairs. This can only be performed with a well segmented hadronic calorimeter. A segmentation of $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ is used for the range $|\eta| < 2.4$, with a coarser segmentation of $\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$ being used for the range $|\eta| > 2.4$. The required pseudo-rapidity coverage of the hadronic calorimeter is determined by the measurement of missing-$E_T$. This measurement obviously requires a good acceptance, but at $|\eta| < 4.5$, the accuracy of the measurement becomes limited by the energy loss due to prompt neutrinos and muons. The hadronic endcap calorimeters extend to $|\eta| < 3.2$, with coverage at higher rapidities being provided by the integrated forward calorimeter.

### 2.4.3. The Integrated Forward Calorimeter

The integrated forward calorimeter is housed in the same cryostats as the endcap calorimeters. It employs LAr technology and provides both hadronic and electromagnetic calorimetry for the extreme forward regions of the experiment ($3.1 < |\eta| < 4.9$). In this pseudo-rapidity range there is an extreme particle and energy flux. Normally in such conditions, LAr calorimeters do not work properly due to the accumulation of positive ions in the LAr gap which distort the electric field and degrade signals from the calorimeter. To avoid this problem, the forward calorimeter uses a tube and rod electrode structure with only a very narrow LAr gap between the electrode and the absorber – 250 μm compared to the more normal gap of 1.94 mm used in the LAr barrel calorimeter. See figure 2.5.

Figure 2.5. The front face of the forward Calorimeter.

The front module of the forward calorimeter performs electromagnetic calorimetry. Copper electrodes are embedded in a copper matrix which acts as an absorber. The outer two modules of the forward calorimeter perform hadronic calorimetry and are constructed of stainless steel electrodes embedded in a tungsten-alloy absorber matrix. Tungsten has an interaction length which is about half that of iron or copper and thus limits the transverse spreading of hadronic showers. This gives the forward hadronic calorimeter an angular resolution comparable to that of the other calorimeters which are further from the interaction point ($\Delta\eta \times \Delta\phi \approx 0.15 \times 0.15$).

## 2.5. The Muon Spectrometer

The muon spectrometer surrounds the ATLAS calorimeter. An air-core toroid system is used to generate a magnetic field of strong bending power over the large volume of the spectrometer. Muon chambers are positioned throughout the magnetic volume, providing measurements of muon track trajectories. The performance requirements and design parameters of the muon spectrometer are derived from physics targets such as the ability to see the decays $H \rightarrow ZZ^* \rightarrow 4l$ or $Z' \rightarrow \mu\mu$ above a large background. The major components of this background are low-energy neutrons, photons, muons and hadrons originating from the interaction of primary hadrons with various machine elements of the detector. In order for the muon spectrometer to provide adequate sensitivity to the physics processes of interest, a good momentum resolution over the pseudo-rapidity range $|\eta| < 3$ is required.

The magnet system used by the muon spectrometer consists of three air-core toroids: a long barrel toroid and two endcaps. The endcaps are inserted into the barrel toroid at either end. Each toroid is composed of eight flat, super-conducting coils. The coils which compose each endcap toroid are housed in a common cryostat, whilst the barrel toroid employs a separate cryostat for each coil. The complete system produces a magnetic field of approximately 4 T over the range $|\eta| < 3$. The light, open nature of the air-core toroid system minimises the contribution of multiple scattering effects to the measured muon momentum.

The position of the muon chambers in the spectrometer is calculated to balance such things as the efficient use of the bending power of the magnets and the pseudo-rapidity coverage of the spectrometer, against practical considerations such as the ease of access to the chambers after installation and the size of chambers that can be easily produced and transported, etc.. The positioning of chambers varies between the different sections of the spectrometer. There are three distinct sections: the barrel section, which covers the pseudo-rapidity range $|\eta| < 1.05$, the transition section ($1.05 < |\eta| < 1.4$) and the endcap section($1.3 < |\eta|$). In the barrel region, the muon

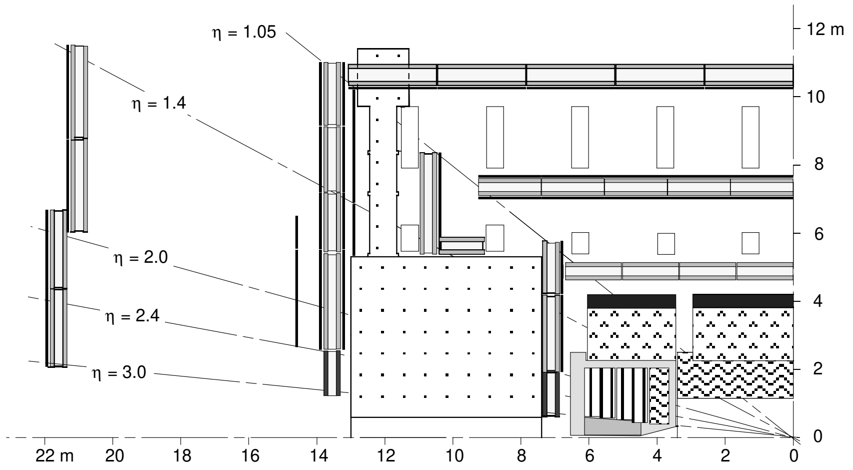


Figure 2.6. Quarter section of the ATLAS detector showing the positioning of the muon chambers. Light shades of grey indicate the precision chamber support structures and dark shades of grey the precision chamber multilayers. Black lines show the trigger chambers.

chambers are arranged in cylindrical shells concentric with the beam axis. In the transition and endcap regions, most chambers are mounted vertically. Figure 2.6 shows the positioning of the muon chambers. The outermost layer of chambers in the endcap region (at $z = \pm$ 22 m) is mounted on the ATLAS cavern wall.

The muon spectrometer employs two types of muon chambers: precision chambers and trigger chambers. The precision chambers perform precise tracking and momentum measurements but, for the most part, they have readout times which are much greater than the bunch-crossing period of the LHC. The trigger chambers provide faster, less precise tracking and momentum measurements for the purposes of triggering. They also allow the muon tracks identified by the precision chambers to be assigned to a bunch crossing.

The precision chambers are positioned at all of the detector stations throughout the spectrometer. Each precision muon chamber consists of two multilayers of detectors, which are mounted to either side of a support structure and separated by 15–35 cm. The basic measurement of each precision chamber is a muon track vector. The multilayers are composed of 3 or 4 planes of individual detectors of a type dependent upon the particular precision chamber. There are two types of precision muon chambers: monitored drift tubes (MDTs) and cathode strip chambers (CSCs). With the MDT chambers, the multilayers are constructed from layers of closely-packed individual drift-tubes. The drift time of these detectors is ~600 ns – much longer than the 25 ns bunch-crossing period. With the CSCs, each detector layer consists of a multiwire proportional chamber. MDTs are employed throughout most of the muon spectrometer, with CSCs being used only in the inner and middle detector stations of the endcap region, for rapidities > | ~2 |. The high track density in this area requires a granularity and level of radiation hardness greater than that which the MDT chambers can provide. It is the increased cost of the CSCs compared to the MDTs that prohibit their use throughout the spectrometer.

The muon trigger chambers form part of a stand-alone trigger system for the muon spectrometer. There are two types of trigger chamber. Resistive plate chambers are employed in the barrel region at the inner and mid-plane detector stations. They consist of gaseous detectors, with parallel electrode plates which are segmented into strips to provide position information. In the transition and endcap regions, thin gap chambers are employed, at the inner and mid-plane detector stations only. The thin gap chambers are multiwire chambers operating in a saturated mode. Their improved rate capacity over the resistive plate chambers makes them more suitable for this region of the spectrometer. It is their increased cost which prohibits their use elsewhere.

## 2.6. The Trigger

The conditions present at the LHC, namely a high bunch-crossing frequency and a very large event rate, present a considerable challenge to the trigger systems of LHC experiments. ATLAS uses a three-level trigger, with successive levels refining the trigger decision of previous levels. Full event building is performed by the final trigger level, and only events which pass this final trigger are passed on to permanent storage. Figure 2.6 shows a block diagram of the ATLAS trigger.

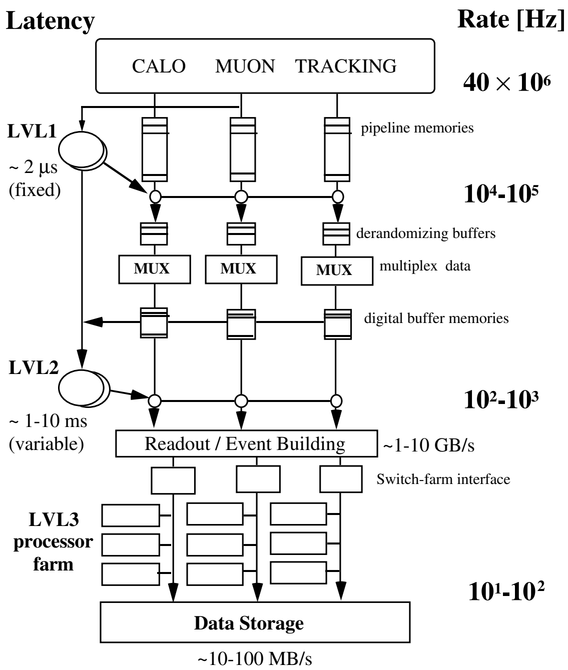


Figure 2.7. The ATLAS Trigger System.

### 2.6.1. The Level-1 Trigger

The level-1 trigger receives events at the full LHC interaction rate of ~$10^9$ Hz. To reduce the event rate to a magnitude that can be accepted by the level-2 trigger, it must trigger at a rate of not more than 100 kHz, a feat which requires an unprecedented rejection power with only one event in $10^4$ being accepted. This must be achieved whilst providing a high trigger efficiency for events which are of interest to the physics programme at ATLAS. The provision of an unambiguous bunch-crossing identification for each event of interest is also a function of the level-1 trigger.

The proposed signatures by which the level-1 trigger recognises an event of interest are shown in table 2.2. All of these triggers are built from four basic building blocks: electron/photon triggers, muon triggers, jet triggers and missing-$E_T$ triggers. It is important for the ATLAS trigger to be flexible enough to react to any unexpected physics or unforeseen situations in the LHC environment and whilst the algorithms used by the level-1 trigger are hard-wired, some flexibility is maintained by making them programmable at the level of parameters. The thresholds quoted in table 2.2 should be taken only as examples.

| **Process** | **Level-1 Trigger** |
|---|---|
| Higgs $\rightarrow \gamma\gamma$<br>(80 GeV < $m_H$ < 130 GeV) | 2 γ's with $P_T$ > 20 GeV |
| Higgs $\rightarrow$ 4l<br>(120 GeV < $m_H$ < 800 GeV) | e-e, e-μ, or μ-μ with $P_T$ > 20 GeV for both leptons |
| Higgs $\rightarrow$ 2l + 2ν<br>(very large $m_H$) | e-e, e-μ, or μ-μ with $P_T$ > 20 GeV for both leptons |
| Top $\rightarrow$ 3 jets + lepton | e or μ with $P_T$ > 40 GeV |
| Top $\rightarrow$ 2 leptons | e-e, e-μ, or μ-μ with $P_T$ > 20 GeV for both leptons |
| W–Z pairs | e-e, e-μ, or μ-μ with $P_T$ > 20 GeV for both leptons |
| SUSY $\rightarrow$ jets + missing-$E_T$ | ≥ 3 jets with $P_T$ > 200 GeV + missing-$E_T$ > 200 GeV |
| SUSY cascade to leptons | e-e, e-μ, or μ-μ with $P_T$ > 20 GeV for both leptons |
| Z', W' $\rightarrow$ leptons | e or μ with $P_T$ > 40 GeV |
| Z', W' $\rightarrow$ jets | 2 jets with $P_T$ > 200 GeV |

Table 2.2. Some benchmark physics processes and their level-1 trigger criteria.

Whilst it is impossible to judge the effectiveness of these signatures in detecting unexpected new physics, it can be noted that they are efficient for most processes which are visible above the large (mainly hadronic) background of the LHC. The calculated acceptance rates for each of the level-1 triggers are shown in table 2.3. The total level-1 trigger rate of around 60 kHz is compatible with the level-2 maximum intake rate of 100 kHz.

In terms of architecture, the level-1 trigger consists of a number of trigger processors, each associated with a different sub-detector, and a global processor to combine results and form the overall trigger decision. All of the level-1 processors are pipelined, with a large number of events being processed concurrently. Only data from the calorimeter and muon detectors are examined by the level-1 trigger. Information from the inner detector is not necessary when making this initial cut on events.

| **Trigger** | **Rate** |
|---|---|
| ≥ 1 isolated em cluster with $P_T > 40$ GeV<br>(isolation not required for clusters with $P_T > 65$ GeV) | 31 kHz |
| ≥ 2 isolated em clusters, each with $P_T > 20$ GeV | 16 kHz |
| ≥ 1 μ with $P_T > 20$ GeV | 8 kHz |
| ≥ 2 μ's, each with $P_T > 20$ GeV | 67 Hz |
| ≥ 2 jets, each with $P_T > 200$ GeV | 5 kHz |
| **Total:** | **60 kHz** |

Table 2.3. Level-1 trigger rates at $L = 10^{34}$ cm$^{-2}$s$^{-1}$.

The level-1 calorimeter trigger consists of a digital processor. It processes information from the LAr electromagnetic calorimeters and the hadronic tile calorimeter, using less than full detector granularity. The calorimeter trigger retains events which contain high-$E_T$ electrons, photons or jets, and events which have large missing $E_T$. The level-1 calorimeter trigger is fully discussed in chapter 3. The level-1 muon trigger processes information from the muon trigger chambers, which provide digital hit information. The events accepted are those which produce high-$E_T$ muon tracks which point back to the region of interaction. Both the muon trigger and the calorimeter trigger include several programmable $E_T$ threshold levels, for use in inclusive muon triggers, inclusive electron/photon triggers and combined multi-lepton triggers. They also include additional, lower thresholds which are not used in the level-1 trigger decision but are used to define Regions Of Interest (RoIs) for the level-2 trigger (see below). The level-1 central processor correlates the muon and calorimeter trigger results and forms combinatorial triggers, transmitting the global trigger decision to the level-2 trigger and to the front-end electronics.

During the latency of the level-1 trigger, the front-end electronics store the data from all of the sub-detectors in pipeline memories. These memories may be either digital or analogue and the technology used varies between detectors. Due to the large number of channels required to store all of the data (~$10^5$ for the calorimeter, ~$10^7$ for the inner tracking detector) the total cost of the pipeline memories rises very rapidly with their length. For this reason it is desirable to keep the latency of the level-1 trigger to a minimum.

The latency of the level-1 trigger is dominated by the calorimeter trigger, which implements relatively complicated algorithms. Other factors which have to be considered are the processing time of the central processor and the propagation delay involved in transmitting information to and from the front-end electronics. A realistic estimate for the total latency of the level-1 trigger is < 2 μs. Approximately half of this is due to the time taken to transport signals to and from the detector.

### 2.6.2. The Level-2 Trigger

If an event is accepted by the level-1 trigger, the required data from that event are read out of the pipeline memories in the front-end electronics into a derandomising buffer. The data are then compressed, multiplexed and transmitted to the level-2 buffer memories, where they are stored for the duration of the level-2 trigger latency. If either the derandomising buffer or the level-2 memories become full, the level-2 trigger sends an inhibit signal to the level-1 central processor. No level-1 triggers are generated while this condition persists.

The aim of the level-2 trigger is to reduce the event rate by a factor of around 100. It is able to refine the level-1 trigger decision in several ways. The use of full calorimeter granularity allows the electron, photon, jet and missing-$E_T$ cuts to be sharpened, and the electron trigger can be refined by the use of information from the inner detector to match calorimeter hits with high-$P_T$ tracks; the muon cut can be refined with the use of information from the precision muon chambers and possibly from the inner detector. The level-2 trigger may also require events of interest to satisfy additional signatures, such as extra leptons at low $P_T$ or large $|\eta|$, or it may make a selection based on the invariant mass of secondary vertices. The full level-2 triggering algorithms are still under study. It is possible that they will be fully programmable, compared to the hard-wired algorithms employed by level-1, which are programmable only at the level of parameters.

In terms of architecture, the level-2 trigger can be divided into two conceptual parts: feature extraction, in which a candidate trigger feature is parameterised using information from a localised area of one or more detectors, and global processing, in which information from local features is combined and an overall trigger decision is made. The RoIs flagged by the level-1 trigger define the areas for the level-2 feature-extraction processors to examine, so that only data from these regions of the detector need be moved to the level-2 trigger. This reduces the data rate into the level-2 processing system by a factor of around 100 and allows for longer processing times for the feature extraction processors.

If an event is accepted by the level-2 trigger, the data from all detectors is read out to the level-3 trigger.

### 2.6.3. The Level-3 Trigger

The level-3 trigger performs full event building and makes a final event selection, before sending data to permanent storage at a rate of 10–100 events/s (10–100 Mbyte/s). A partial event read-out is made available for some types of trigger, e.g. calibration triggers or a high mass jet-pair trigger which may indicate new particles decaying to jets.

The bandwidth required for full event building is in the order of a few GBytes/s. It is not possible to achieve this bandwidth using traditional bus-based event builders and instead an architecture is required which involves a high level of parallelism. It is expected that products currently being developed for the telecommunications industry will be used in the construction of the level-3 trigger.

Once full event building has taken place, the data is routed to a farm of processors of sufficient power to allow sophisticated algorithms to be run. It is expected that these processors will run the off-line analysis algorithms. This final stage of event reduction will reduce the data rate to a level that is manageable for storage.

# 3. The Level-1 Calorimeter Trigger

## 3.1. Introduction

The formation of a level-1 trigger decision is a two-stage process. In the first stage, a calorimeter trigger processor and a muon trigger processor work independently and in parallel to implement four inclusive subtriggers, based on the following event characteristics:

- the multiplicity of electrons/photons at several different energies;
- the multiplicity of jets at several different energies;
- the missing transverse energy (missing-$E_T$);
- the multiplicity of muons at several different energies.

It can be seen that the first three of these four subtriggers require information from the ATLAS calorimeters alone, hence they are implemented by the calorimeter trigger processor. The fourth is implemented by the muon trigger processor. The data output from all four subtriggers are passed on to a central trigger processor which forms the overall level-1 trigger decision, based on combinations of the subtrigger results. Figure 3.1 shows a diagram of this system.

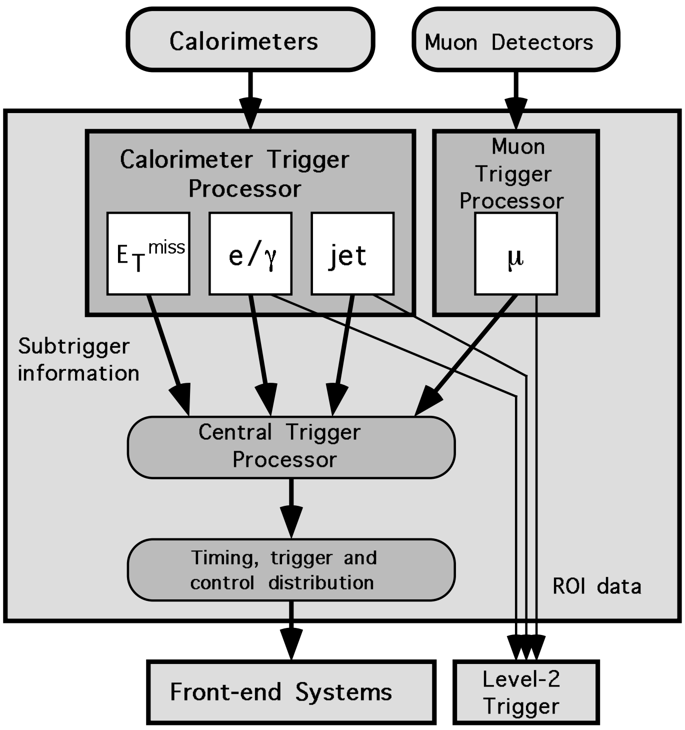


Figure 3.1. A Block diagram of the ATLAS level-1 trigger system.

The design of the level-1 calorimeter trigger is the subject of an ATLAS R & D programme [18]. Extensive simulation studies have been carried out to optimise the algorithms used by the level-1 calorimeter subtriggers, and technical demonstrator systems have been built to validate various aspects of the proposed trigger design. Investigations have also been made into bunch-crossing identification (see later). The calorimeter trigger design which has resulted from this R & D programme is described in the following sections of this chapter [19].

## 3.2. Subtrigger Algorithms.

This section describes the algorithms proposed for the *final* ATLAS level-1 calorimeter trigger. These algorithms should not be confused with those implemented by the demonstrator system described elsewhere in this thesis (chapters 5 and 7). In particular, the e/γ trigger algorithm described below differs slightly from that implemented by the demonstrator system.

### 3.2.1. The Electron/Photon Trigger

***Local Cluster Identification***

The algorithm used by the e/γ trigger is illustrated in figure 3.2. It is based upon the detection of isolated $E_T$ deposits in the electromagnetic calorimeter. A $4 \times 4$ window of trigger cells is considered by the trigger, where each cell is of granularity $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$. This window slides in both the $\eta$ and $\phi$ directions by one trigger cell, such that all possible window positions in the range $|\eta| < 2.5$ are processed. Results from a window are referenced to the trigger cell indicated.

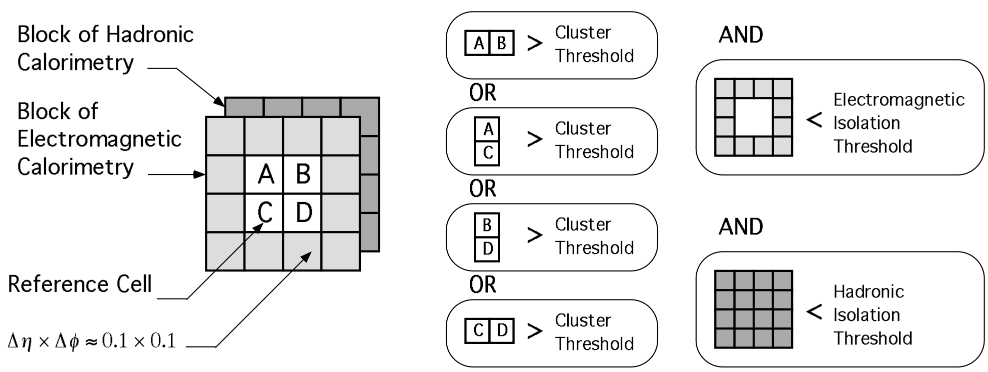


Figure 3.2. The e/γ trigger algorithm.

To produce a positive trigger result the $E_T$ deposits in a window must satisfy the following conditions. Firstly, the energy deposited in one of the four cluster windows formed by the trigger must be greater than the cluster threshold level. The four cluster windows are shown in figure 3.2: they cover all possible $2 \times 1$ combinations of electromagnetic trigger cells in the centre of the 4 × 4 window. The second condition necessary for a positive trigger result is that the energy sums in both the electromagnetic isolation window and the hadronic isolation window must be lower than their respective isolation thresholds. All of the thresholds used by the e/γ trigger are programmable.

The figures of merit for the e/γ trigger are the sharpness of the trigger threshold (i.e., the rate at which the acceptance falls as a function of $E_T$) and the trigger rate observed when the thresholds are set at a level efficient for the desired physics. The factors which affect the sharpness of the threshold are the resolution of the calorimeter, the degree of containment of the e/γ shower within the cluster window, and the level of pile-up and electronic noise present in the window. The trigger rate, which is dominated by the QCD jet background, is determined primarily by the area of the cluster window. It is therefore important to optimise this area. Ideally, it should be large enough to fully contain an e/γ shower but be small compared to the size of a jet core.

The use of a 2 × 1 area of trigger cells for the e/γ cluster window, rather than a single trigger cell, significantly improves the sharpness of the trigger threshold as it reduces containment inefficiencies for those showers which spread their energy between two cells. This is particularly important in the forward regions of the calorimeter where the physical size of a trigger cell is relatively small. The use of a window of 2 × 2 trigger cells would further reduce inefficiencies in cluster containment, but simulation shows that this has a negligible effect on the sharpness of the threshold. Conversely, doubling the area of the cluster window doubles the expected background trigger rate from jets. For these reasons the 2 × 1 area of trigger cells is the preferred window size.

Jets can produce fake e/γ triggers by one of two processes: high-$P_T$ $\pi^0$s can leave large $E_T$ deposits in the electromagnetic calorimeter via a decay to photons and, with a lower probability, high-$P_T$ charged hadrons can produce em showers via the charge exchange process. Whilst many of the events of interest at the LHC will result in isolated $E_T$ deposits, an $E_T$ deposit produced by a jet will, in general, be accompanied by a further, nearby $E_T$ deposit. It is for this reason that an isolation requirement is included in the e/γ trigger.

The effectiveness of the isolation veto is determined primarily by the tightness of the cuts which can be applied whilst retaining a high efficiency for e/γ showers. This tightness is limited by three factors: pile-up, electronic noise, and leakage of the e/γ shower into the isolation window. The sensitivity of the isolation veto to pile-up and electronic noise can be reduced significantly by requiring each trigger cell to possess an energy deposit greater than 1 GeV before including it in the total isolation sum. The sensitivity of the veto to shower leakage can also be reduced by raising the isolation threshold for high-$P_T$ e/γ candidates (a task which requires the trigger to possess more than one set of cluster and isolation thresholds).

Figure 3.3 shows the effectiveness of the isolation veto in reducing the background trigger rate. This plot is taken from a Monte Carlo study of QCD jets. The e/γ trigger algorithm used to generate this plot is not that illustrated in figure 3.2. Here only two of the four possible cluster widows are formed: $A + C$ and $C + D$. This algorithm (which is the precursor of that described above) has slightly lower efficiency for isolated electrons. It can still be used to judge the effectiveness of the isolation veto however. Figure 3.3 shows the background rate for the single electron trigger as a function of the cluster threshold, both without any isolation requirement and with an isolation threshold of 5 GeV in the em calorimeter. With a typical cluster threshold of 40 GeV, imposing this isolation threshold reduces the background trigger rate by an order of magnitude, to below 10 kHz. This compares well with the maximum level-2 acceptance rate of 100 kHz.

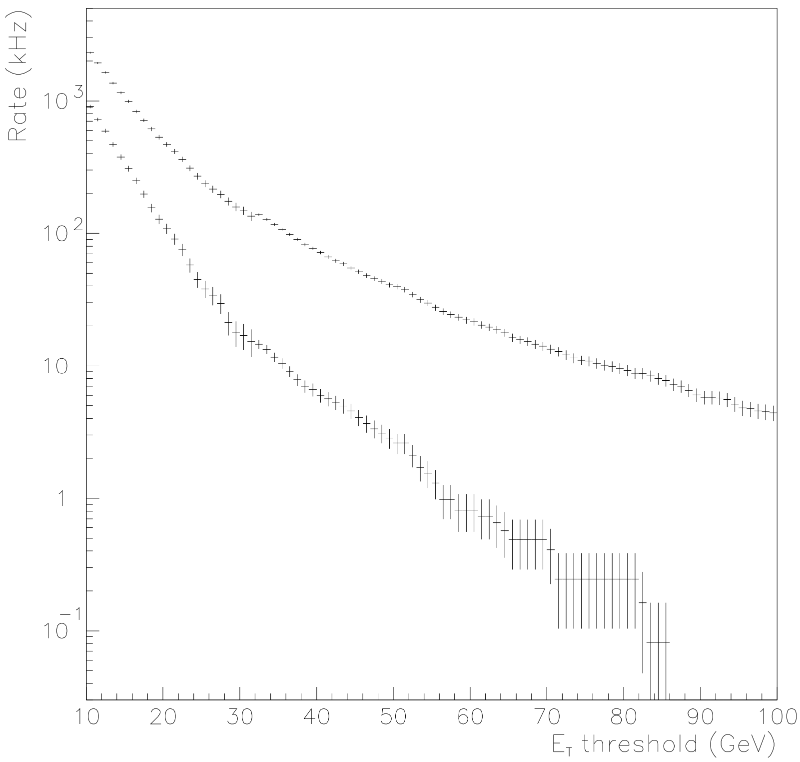


Figure 3.3. Background trigger rate for the e/γ trigger due to QCD jets, versus the cluster threshold for $|\eta| < 2.5$ and luminosity $10^{34}$ cm$^{-2}$s$^{-1}$, with (lower curve) and without (upper curve) making the isolation requirement stated in the text.

The proposed e/γ trigger has eight sets of thresholds, where each set consists of a cluster threshold, an electromagnetic isolation threshold, and a hadronic isolation threshold. Eight different e/γ trigger requirements can therefore be specified simultaneously. This allows isolation vetoes to be relaxed for high-$E_T$ deposits, different thresholds to be set for single and multiple cluster triggers, and RoIs to be defined using thresholds which are more relaxed than those used to produce actual level-1 triggers.

***De-clustering and RoI Generation***

Since the $4 \times 4$ windows used by the e/γ trigger overlap, a shower in one trigger cell may pass the trigger conditions in up to four windows, producing a cluster of trigger hits. Logic is therefore needed to 'de-cluster' these overlapping hits and avoid multiple counting. There are two distinct tasks to be performed: the number of e/γ hits produced by an event must be counted, and the location of each shower which passes the trigger conditions must be unambiguously identified. The former task is required to allow global trigger decisions based upon multiple triggers, whilst the latter is needed to identify RoIs for the level-2 trigger.

As mentioned in section 2.6.2, the level-2 trigger only reads out those sections of the detector which are flagged as RoIs by the level-1 trigger. To minimise the amount of data which the level-2 trigger must accept, an RoI should therefore be defined as precisely as possible, i.e., at a trigger-cell level. To define the position of an RoI the following algorithm can be used: in addition to being part of an $E_T$ cluster which passes trigger conditions, a candidate RoI cell must also contain a greater $E_T$ than its neighbours. To resolve the rare cases where a shower is shared equally by two cells, the algorithm can be modified such that the $E_T$ must be greater than that of the neighbours on two specified sides, and greater or equal to the $E_T$ of the neighbours on the opposite two sides. Clearly, no two neighbouring cells can both pass such a condition. Each e/γ shower which passes trigger conditions therefore produces a single RoI candidate, which defines the position of the shower for the benefit of the level-2 trigger.

A variation on the algorithm described above can also be used to perform cluster counting: if an e/γ hit is required to satisfy the conditions illustrated in figure 3.2, plus the RoI condition described above, then no multiple counting of trigger hits occurs. Such an algorithm however, requires a significant increase in the fan-out of data signals between processing units in the calorimeter trigger, and for this reason it is not the preferred solution. Instead, the following global method of de-clustering is used: results from neighbouring trigger windows are combined (using an *OR* function) to form de-clustering windows of granularity $\Delta\eta \times \Delta\phi \approx 0.4 \times 0.4$. These windows do not overlap. The algorithm illustrated in figure 3.4 is then used to count the number of de-clustering cells which produce trigger hits.

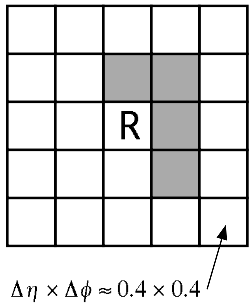


Figure 3.4. The Global cluster-counting algorithm. A hit in the window *R* is counted provided there are no hits in the shaded windows.

The use of a reduced granularity for global cluster-counting inevitably results in a loss of efficiency for multiplicity based triggers, as multiple e/γ hits which are separated by less than twice the de-clustering granularity are counted as a single hit. The use of full trigger-cell granularity however, is prohibited by the volume of data that would need to be processed. Physics simulation studies show that if de-clustering windows of $\Delta\eta \times \Delta\phi \approx 0.4 \times 0.4$ are used, the loss in efficiency for cluster pair triggers is minimal. See table 3.1.

| **Process** | **0.1 × 0.1** | **0.4 × 0.4** | **0.8 × 0.8** |
|---|---|---|---|
| $H \rightarrow \gamma\gamma$ | 0.979 ± 0.008 | 0.969 ± 0.010 | 0.928 ± 0.015 |
| $H \rightarrow e^+ e^- e^+ e^-$ | 0.987 ± 0.006 | 0.957 ± 0.010 | 0.839 ± 0.019 |
| $tt \rightarrow e^+ e^-$ jets | 0.938 ± 0.013 | 0.903 ± 0.016 | 0.824 ± 0.021 |

Table 3.1. The efficiency of the two-cluster e/γ trigger as a function of granularity for a number of key physics processes. All events contain two electrons or photons with $P_T > 20$ GeV within $|\eta| < 2.5$.

### 3.2.2. The Jet Trigger

The jet trigger compares the total $E_T$ in a window of trigger cells with a programmable threshold, producing a positive trigger decision if the $E_T$ is greater than the threshold. Studies are still being conducted to determine the optimum size of the jet window: if the window is too small it will not fully contain the jet $E_T$ and resolution will be lost, whereas if the window is too large the summation of electronic noise and pile-up over the increased area will also lead to a loss in resolution. One consequence of this latter fact is that the optimum window size for the low luminosity running period of the LHC is larger than that for the high luminosity running period. As the window size must be fixed, it should be optimised for the high luminosity running period.

The resolution of the jet trigger is its only figure of merit. The resolution alone determines the sharpness of the threshold curve and thus the trigger rate obtained when the threshold is set at a level efficient for the required physics. There is very little background to the jet trigger, due to jet production being the dominant process at the LHC.

The proposed jet trigger is based upon the total $E_T$ in a window of size $\Delta\eta \times \Delta\phi \approx 0.8 \times 0.8$ in both the electromagnetic and hadronic calorimeters. This window slides in units of half-granularity (i.e., $\Delta\eta$ or $\Delta\phi \approx 0.4$ for a window of size $\Delta\eta \times \Delta\phi \approx 0.8 \times 0.8$). This avoids any inefficiency due to jets falling between windows. Unfortunately, it also means that a de-clustering algorithm is required to avoid the multiple counting of jet hits and to define RoIs. The algorithm used is this: each jet window is divided into 2 × 2 jet cells and trigger hits generated within a window are associated arbitrarily with the bottom left cell (see figure 3.5). The jet hits are then de-clustered at jet-cell granularity using the same algorithm used to perform global cluster-counting for the e/γ trigger (see figure 3.4). Note that here, the logic used to avoid multiple counting can also be used to determine RoI coordinates, as the de-clustering granularity matches the finest granularity used by the jet trigger. The RoI candidate cells are those jet cells which contain jet-trigger hits not vetoed by the de-clustering algorithm.

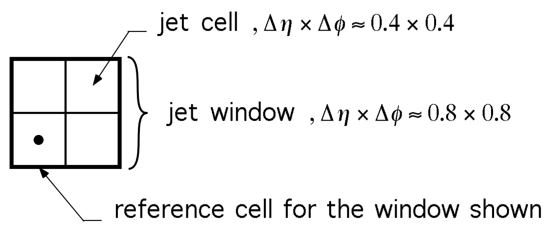


Figure 3.5. A jet window.

### 3.2.3. The Missing-Transverse-Energy Trigger

The principle of the missing-$E_T$ trigger is simple. Calorimeter $E_T$ measurements are resolved into $E_x$ and $E_y$ components and these quantities are summed over the acceptance range of the trigger. The total missing-$E_T$ measurement is given by the quadratic sum of the global $E_x$ and $E_y$ measurements, and this quantity is compared to a programmable threshold to produce a trigger decision.

The missing-$E_T$ trigger processes data from all of the ATLAS calorimeters, giving it an acceptance of $|\eta| < 5$. This large pseudo-rapidity coverage is essential to avoid miss-measurements of missing-$E_T$ due to jets falling outside the acceptance of the trigger. To reduce the sensitivity of the missing-$E_T$ trigger to noise, only trigger cells which possess an $E_T$ above a 1 GeV threshold are included in the global $E_x$ and $E_y$ sums. The largest contribution to the background of the missing-$E_T$ trigger is from the degradation of $E_T$ measurements in the transition regions between calorimeters, due to both the extra material in front of the calorimeters in these regions and the difficulty in combining measurements from different calorimeters. There is also a non-physics background due to beam-wall and beam-halo events. Only if this background can be controlled will an inclusive missing-$E_T$ trigger be employed by ATLAS. If not, the missing-$E_T$ measurement will be used as a secondary signature in conjunction with other trigger criteria.

## 3.3. Architectural Overview

The following sections describe the design proposed for the ATLAS level-1 calorimeter trigger, as presented in the ATLAS technical proposal and the accompanying back-up document [19]. The design is based upon purpose-built digital processors and extensive use is made of custom micro-electronics. The proposed processors implement a fixed set of algorithms which are programmable at the level of parameters. A pipelined architecture is employed by the trigger, with many operations being performed in parallel. The transmission of data to the processors and the pipelined processing are synchronous with the 40 MHz LHC clock (in some cases clock frequencies which are multiples of 40 MHz are used).

Figure 3.6 shows the components of the proposed calorimeter trigger system. It can be seen that signals from the calorimeter are digitised by hardware mounted on the detector and are then transferred by fibre-optic cable to a shielded control room, where the trigger processor itself is housed.

### 3.3.1. The Front-End Digitisation System

The first task performed by the level-1 calorimeter trigger is the summation of the analogue calorimeter signals to form trigger towers of the desired granularity: $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$. An ADC system is then used to digitise the resultant signals by sampling them every 25 ns. The ADC clock is derived from the central LHC clock, but it has an adjustable phase which allows the analogue calorimeter pulses to be sampled at their peak. The dynamic range of the ADC is eight bits and the sensitivity is 1 GeV/count.

As the calorimeters do not measure $E_T$ directly, but rather $E$, Look-up tables (LUTs) are employed to convert the ADC output into $E_T$ values. The signals then undergo a process of bunch-crossing identification – a subject which is covered in detail in chapter 4 and so is not explained here. As indicated in figure 3.6, all of the above processes may be performed either by a stand-alone level-1 front-end system, or by the FERMI front-end readout module [20], if it is accepted by ATLAS.

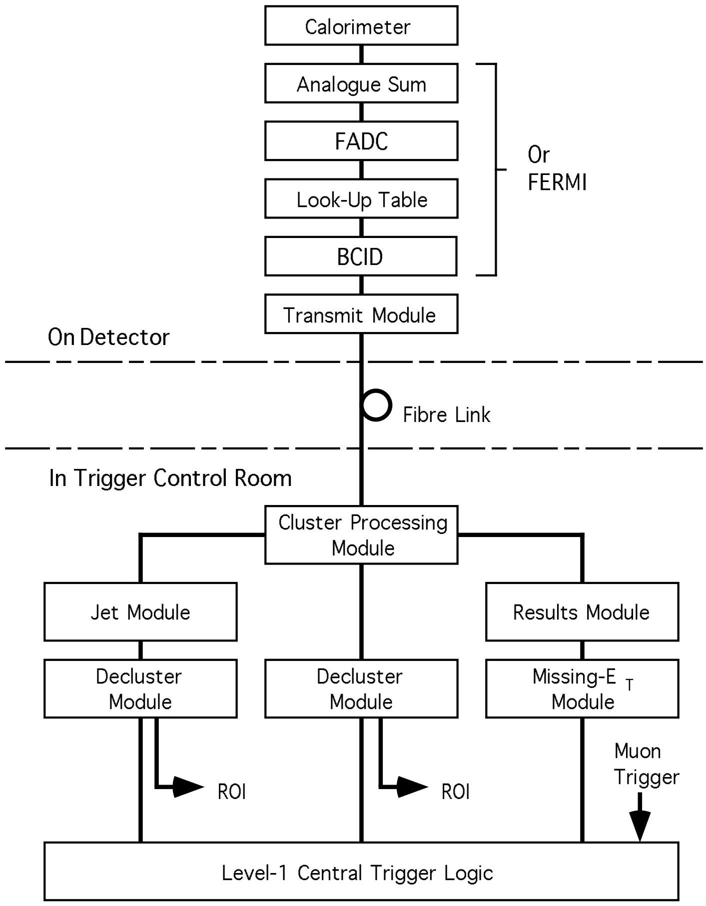


Figure 3.6. Components of the Level-1 calorimeter trigger.

### 3.3.2. The Optical Data Link

Fibre-optic cables are used to transfer trigger data off the detector. Data bytes from pairs of trigger channels are combined to form 16-bit words and these are serialised and transferred over an optical link at an effective rate of 640 Mbit/s. (The actual clocking rate is 800 Mbaud across the link due to the transfer of additional framing and protocol bits, which are internal to the link and not seen by the outside system). A total of 4000 cables are required to transfer all of the calorimeter data off the detector. In the counting room the cables are received into the trigger processor system described below.

### 3.3.3. The Trigger Processor System

Figure 3.7 shows a crate-level diagram of the level-1 calorimeter trigger processor. This is the part of the system which implements the trigger algorithms described in section 3.2. Six different types of module are used in the system:

- cluster processing modules (receive optical data, produce e/γ hit decision, perform initial processing for jet and missing-$E_T$ triggers);
- jet modules (form jet hits);
- results modules (calculate $E_x$ and $E_y$);
- missing-$E_T$ module (calculates missing-$E_T$ from local $E_x$ and $E_y$ results);
- de-clustering modules (de-cluster e/γ hits and jet hits);
- readout controller modules (interface to host computer and level-2 trigger).

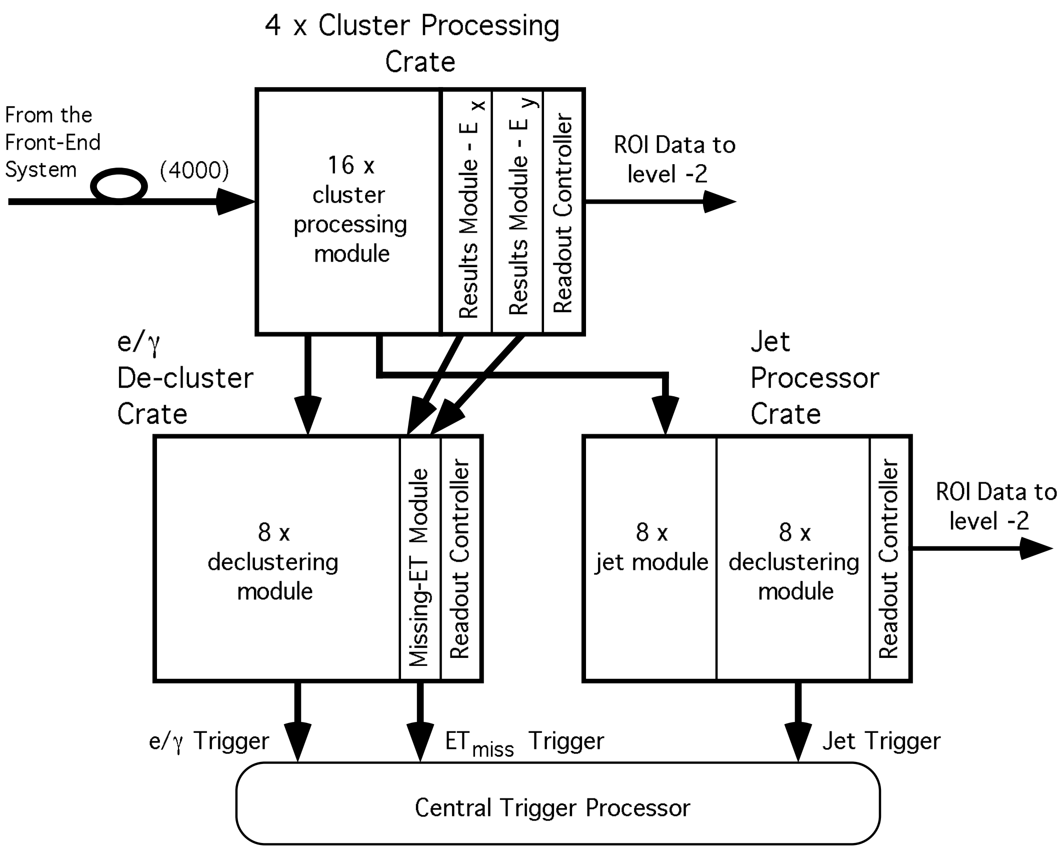


Figure 3.7: Crate-level diagram of the ATLAS level-1 calorimeter trigger processor.

## 3.4. Trigger Implementation

The sections below describe the data path through the trigger processor for each of the three calorimeter subtriggers (e/γ, jet and missing-$E_T$).

### 3.4.1. The Electron/Photon Trigger

Figure 3.8 shows a block diagram of the logic for the e/γ trigger. Optical data from the front-end digitisation system are received on to cluster processing modules, each of which receives the data for 242 trigger cells. The cluster processing modules contain multichip modules (MCMs) to perform optical-to-electrical conversion, and cluster-processing ASICs to identify e/γ hits. The hit data produced by the cluster processing modules are passed on to de-clustering modules, which perform global cluster counting. A total of fifteen cluster processing modules and eight de-clustering modules are employed by the e/γ trigger.

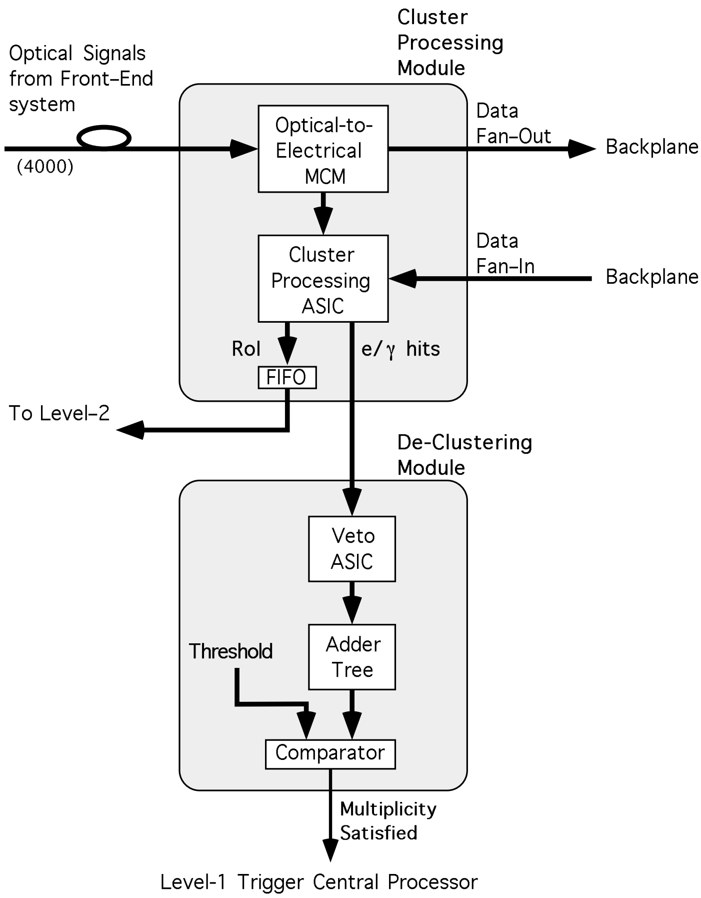


Figure 3.8: Electron/Photon Trigger Logic

### *The Optical-To-Electrical Multichip Modules*

Each cluster processing module contains four optical-to-electrical MCMs of the type shown in figure 3.9. The fibre-optic cables which carry the calorimeter data off the detector are received directly on to the MCMs, with four fibres being received per module. As shown in figure 3.9, each MCM contains four different types of chip in bare-die form: optical-to-electrical converters, serial-to-parallel converters, serialising ASICs and a diagnostic memory. The serial-to-parallel converters shown form a chip-set with the parallel-to-serial converters used at the transmitter end of the optical link. Problems of framing and protocol for the transferred data are transparent to the user.

In order to minimise the number of signal tracks required to carry the output from the MCM, the 16-bit parallel words produced by the serial-to-parallel converters are serialised on to four links each (16 links per MCM). These links operate at 160 Mbaud. The maximum baud rate which can be used to output data from the MCM is determined by the cluster ASIC, which is implemented in 0.5 μm CMOS gate-array technology and is not able to handle rates beyond 160 Mbaud. Rates much above 160 Mbaud also require an optical fan-out technique to be employed, rather than the proposed electrical backplane (see below).

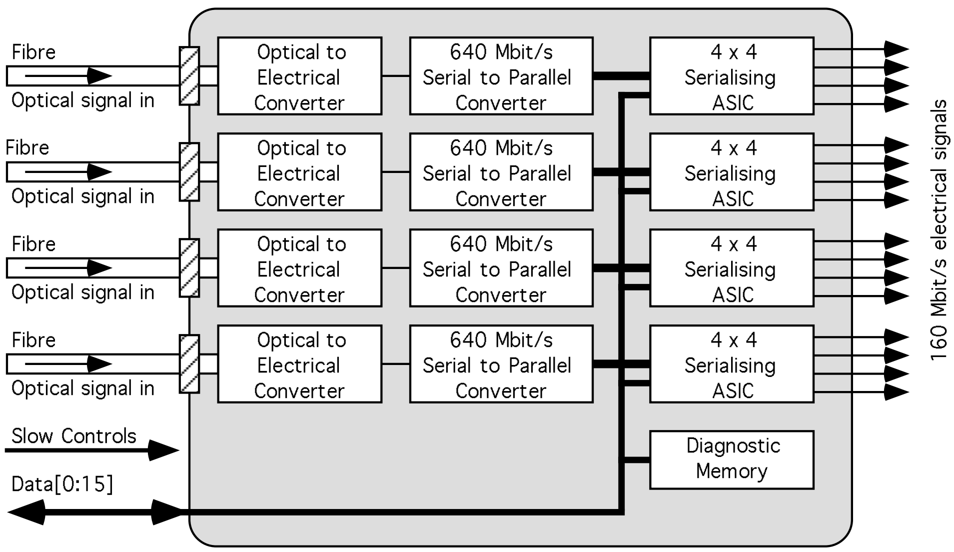


Figure 3.9. MCM block diagram.

### *The Cluster ASIC*

The cluster processing ASIC lies at the heart of the calorimeter trigger processor. It identifies e/γ trigger hits. Each ASIC computes 16 overlapping windows of the type shown in figure 3.2, requiring data from a $7 \times 7$ area of electromagnetic and hadronic trigger cells (see figure 3.10). The data for each pair of trigger cells arrive at the cluster

ASIC on four 160-Mbaud links. The data are converted into two 8-bit parallel words for processing by the core logic of the ASIC, which operates at 40 MHz.

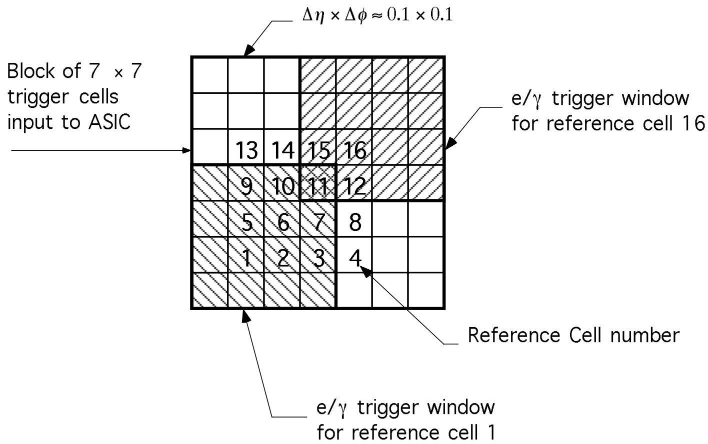


Figure 3.10. Input data to a cluster processing ASIC.

The cluster processing ASIC provides eight sets of thresholds for the e/γ trigger, where each set consists of a cluster threshold, an electromagnetic isolation threshold and a hadronic isolation threshold. There are therefore eight sets of results produced for each window location. Within the ASIC, the hit results for different window locations (but not different threshold values) are combined using an *OR* function: if any of the window locations produces a positive e/γ-hit result for a given set of thresholds, then the output of the whole ASIC for that set of thresholds is positive (i.e., an e/γ cluster has been found). This is the first step in the global cluster-counting algorithm described in section 3.2.1. The output of each cluster processing ASIC for the e/γ trigger consists of an 8-bit hit pattern, where each bit gives the e/γ hit decision for one set of thresholds over the $\Delta\eta \times \Delta\phi \approx 0.4 \times 0.4$ area covered by that ASIC.

In addition to the above, the cluster processing ASICs generate RoI candidates according to the algorithm described in section 3.2.1. Each cluster ASIC outputs a 16-bit array of RoI data, covering the 16 reference cells of the e/γ windows processed. A positive RoI result for a given reference cell indicates that an e/γ shower which passes trigger conditions is centred on that reference cell. The RoI data are stored in a FIFO on the cluster processing module. This FIFO can be read out into the level-2 trigger as required.

***Data Fan-In/Fan-Out***

The overlapping nature of the windows used by the e/γ trigger means that each trigger cell contributes towards 16 different windows. Inevitably, the data from some trigger cells must be shared out between different cluster processing ASICs, different cluster processing modules, and in the most extreme case different cluster processing crates.

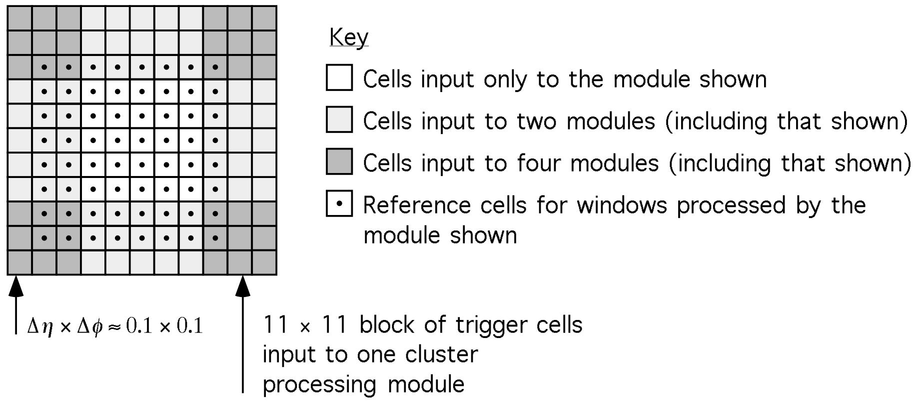


Figure 3.11. Data fan-in/fan-out for a cluster processing module.

Each cluster processing module houses four cluster ASICs, processing a total of $8 \times 8$ e/γ trigger windows. This requires data from an $11 \times 11$ area of trigger cells in both the electromagnetic and hadronic calorimeters. As shown in figure 3.11, of the 242 cells which form the input to a cluster processing module, data for 120 cells are shared between two modules and data for 72 cells are shared between four modules.

Data sharing between modules in a cluster processing crate is achieved via a custom high-speed backplane, which transfers the data as ECL signals at 160 Mbaud. The fan-out/fan-in requirement for each cluster processing module is this:

- data for 128 cells arrive at a module optically, from the front-end digitisation system;
- data for a further 114 cells arrive via the crate backplane from neighbouring cluster processing modules;
- of the data received optically, data for 78 cells are output on to the crate backplane and transferred to neighbouring modules.

Data sharing between crates is accomplished using passive optical splitters to fan-out the data whilst they are being transmitted optically from the front-end system.

***De-Clustering***

The sixteen cluster processing modules employed by the e/γ trigger output a total of 2048 bits of hit data (256 position pixels $\times$ 8 cluster thresholds). These data are transmitted via parallel cable links to eight de-clustering modules, such that each de-clustering module receives all 256 bits of hit information for one set of thresholds. The de-clustering modules implement the global de-clustering algorithm shown in figure 3.4 and count the number of non-vetoed hits. This number is then compared to a programmable multiplicity threshold and a *multiplicity-satisfied* flag is generated. These flags (one per set of cluster thresholds) constitute the final output of the e/γ trigger and they are transmitted to the central level-1 trigger processor.

### 3.4.2. The Jet Trigger

Figure 3.12 shows a diagram of the jet trigger logic. As described in section 3.2.2, the jet trigger is based upon the $E_T$ sum in a $2 \times 2$ window of jet cells. It is assumed here that the jet cells used are of dimension $\Delta\eta \times \Delta\phi \approx 0.4 \times 0.4$, and are formed in the cluster ASICs by summing over $4 \times 4$ areas of trigger cells. These sums are output from the cluster processing modules on 160 Mbaud cable links and are sent to the jet modules, which implement the bulk of the jet trigger logic.

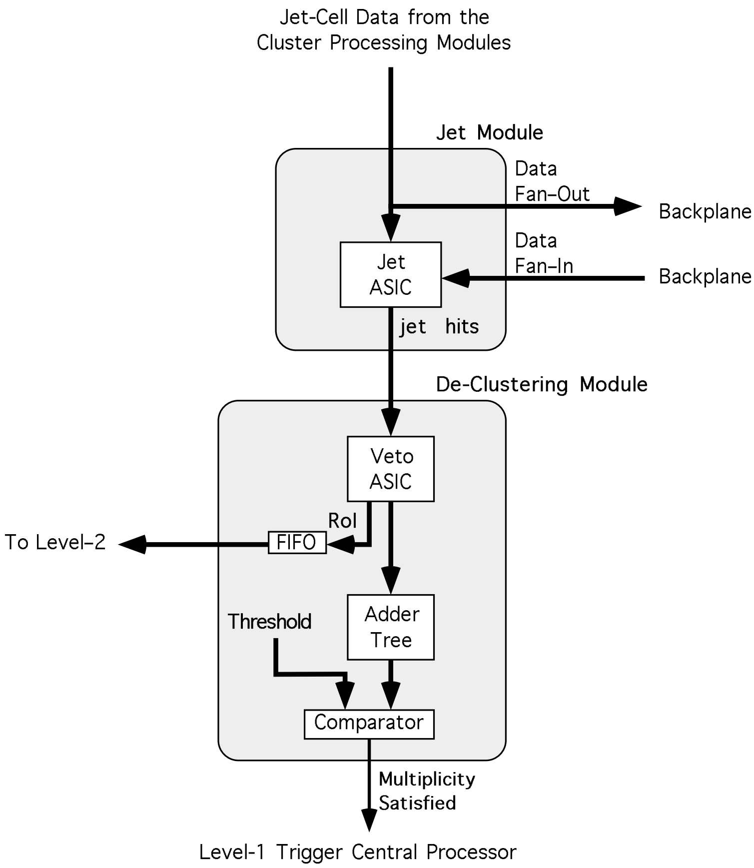


Figure 3.12. Jet Trigger Logic.

Each jet module houses four jet ASICs, each of which processes $4 \times 2$ jet windows. As each jet window is composed of $2 \times 2$ jet cells and the windows slide by one cell in both the $\eta$ and $\phi$ directions, this requires an input of $5 \times 3$ jet cells per ASIC. See figure 3.13. Extending this argument further shows that the input to each jet module is 45 jet cells ($9 \times 5$). Data for 32 of these cells arrive directly on the module via twisted pair cable from the cluster processing modules, and data for the remaining 13 arrive via the crate backplane from neighbouring jet modules, in a similar fan-out scheme to that used for the cluster processing modules. There are 265 jet cells in total, so eight jet modules are required by the trigger system.

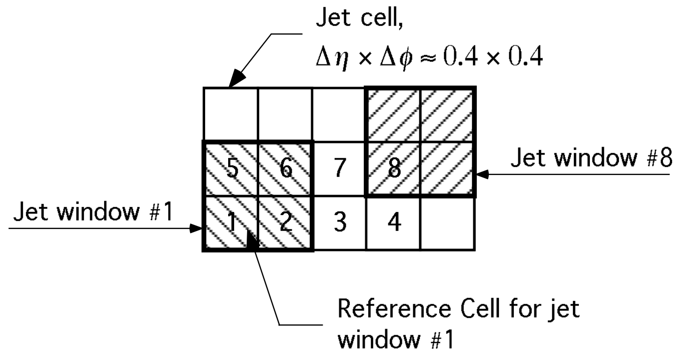


Figure 3.13. The 5 × 3 block of jet cells which form the input to a single jet ASIC.

Each jet module compares the $E_T$ in the 32 windows which it processes to eight different thresholds. The output from each module thus consists of 32 8-bit hit patterns. The results from the jet modules are passed on via 160 Mbit/s cable links to eight de-clustering modules which are identical to those used for the e/γ trigger. One block of logic which is present on all de-clustering boards but is superfluous to the e/γ trigger is the FIFO, which stores the pattern of non-vetoed hits produced by the veto ASICs. For the jet trigger this pattern constitutes the required RoI information and it is transferred from the FIFO to level-2 upon request.

### 3.4.3. The Missing-Transverse-Energy Trigger

Figure 3.14 shows the missing-$E_T$ trigger logic. The calculation of missing-$E_T$ begins on the cluster processing modules. Each of the four cluster ASICs on a module measures the $E_T$ deposited in a $\Delta\eta \times \Delta\phi \approx 0.4 \times 0.4$ area of the calorimeter. These $E_T$ measurements are resolved into their $E_x$ and $E_y$ components by LUTs, before being summed over the area of calorimeter covered by that cluster processing module. The partial $E_x$ and $E_y$ sums are then transferred via crate backplane to the results modules.

There are two results modules per cluster processing crate: one to form a local $E_x$ sum for the entire crate and one to form a local $E_y$ sum. The outputs from all of the results modules are transmitted to the missing-$E_T$ module, which implements the final stage in the missing-$E_T$ adder tree, forming the global $E_x$ and $E_y$ sums. These sums are then sent to a LUT which tests $E_T = \sqrt{(E_x^2 + E_y^2)}$ against four pre-programmed thresholds. The output of the missing-$E_t$ trigger is a four-bit hit flag which gives the trigger results for each of the four threshold values. This hit flag is passed on to the central level-1 trigger processor.

This concludes the overview of the level-1 calorimeter trigger.

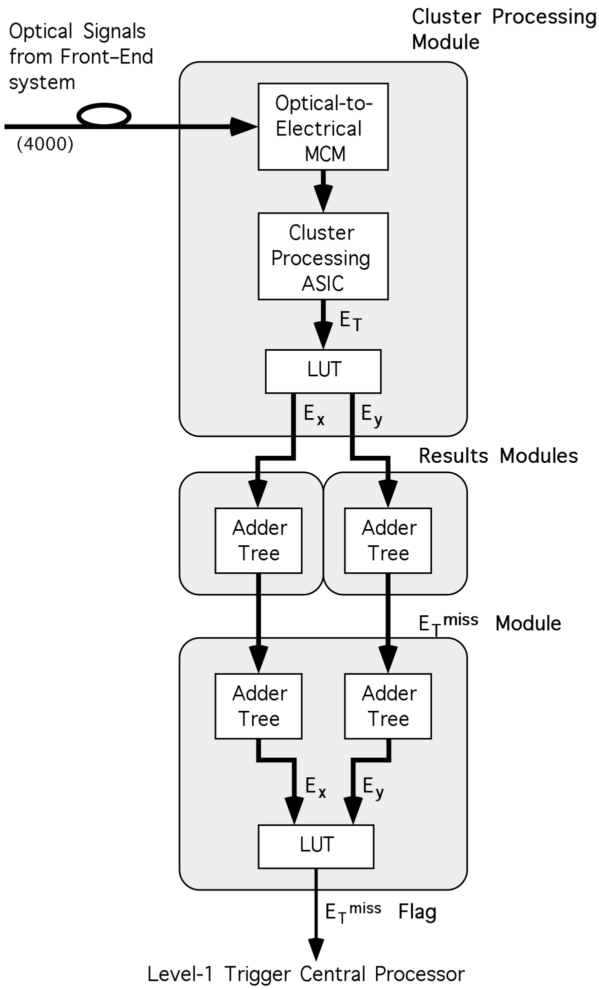


Figure 3.14. Missing-$E_T$ Trigger Logic

# 4. Bunch-Crossing Identification for the Level-1 Calorimeter Trigger

## 4.1. Introduction

One of the tasks of the level-1 trigger will be to provide unambiguous bunch-crossing identification for each event of interest. The implication of this for the individual level-1 trigger processors is that each trigger signal must be identifiable with one unique bunch crossing.

For the muon trigger, a dedicated muon detector can be built using resistive plate chambers (RPCs) which have a response time much less than 25 ns. The problem of bunch-crossing identification in this case is essentially one of synchronisation, allowing for the time of flight of the muons (which will be of the same order as the bunch-crossing period) and propagation delays within the trigger. Synchronisation can be achieved easily using programmable delays.

For the calorimeter trigger the situation is more complicated. The liquid-argon calorimeters proposed for ATLAS have the drawback that they have a charge collection time much longer than the bunch-crossing period (~400 ns). To correct for this, electronic shaping of the calorimeter signals is employed. Figure 4.1 shows an example of the resultant pulse shape relative to the LHC clock. This pulse was taken from the RD3 prototype liquid-argon accordion calorimeter [21] and has been digitised by sampling with a 25 ns period (the bunch-crossing period). The pulse has a narrow positive lobe which would span three to four bunch crossings, followed by a much shallower negative lobe which would span around 15 bunch crossings. It is the width of the positive lobe which is of the most concern.

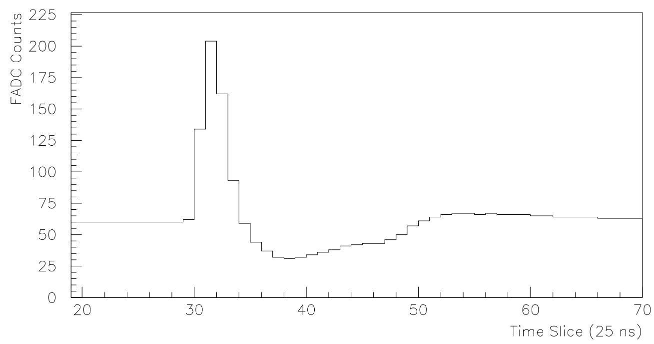


Figure 4.1. A typical pulse from the RD3 prototype Liquid-Argon Accordion Calorimeter, sampled at 40 MHz.

If signals such as the one shown in figure 4.1 were sent directly to the calorimeter trigger, several undesirable consequences would result. Most noticeably, the trigger rates of the e/γ and jet triggers would be increased as large energy deposits would cause the triggers to fire for several consecutive bunch-crossings. At the same time, the efficiency of the e/γ trigger for isolated electrons would be decreased, as genuinely-isolated em clusters would occasionally be vetoed by energy deposits from other bunch-crossings. There would also be an overall rise in the level of pile-up noise which would degrade the resolution of the jet and missing-$E_T$ triggers. The process of bunch-crossing identification can be used to improve this situation. Once each calorimeter signal has been identified with a single bunch crossing, that signal can be truncated in time and the information from that signal sent to the trigger processor for one bunch crossing only.

Bunch-crossing identification (BCID) may be performed using either digital or analogue signal processing methods. Only digital implementations of a BCID filter have been studied during the course of this project, although analogue implementations of a BCID filter have not been ruled out for the final LHC system. The respective advantages of the two approaches are listed below.

Digital Signal Processing (DSP):

- programmable systems can be used, giving maximum flexibility;
- non-linear DSP algorithms can be very effective in the presence of noise;
- ease of diagnostics.

Analogue Signal Processing (ASP):

- ASP systems can be very fast;
- there is no quantisation noise on the observed signal;
- simple (and possibly cheap) circuits may be used.

If the FERMI digital front-end system is accepted for calorimeter readout, the bunch-crossing identification of calorimeter signals will be performed using digital algorithms and the required logic will be housed in the FERMI module itself. If an analogue system is used for calorimeter readout, the level-1 trigger will need to employ both an independent ADC system and a BCID module. This BCID module could be placed before or after the ADC system, allowing either analogue or digital signal processing methods to be used.

There are three sources of noise present on the signals observed by the BCID logic:

- Electronic noise. This consists of pick-up etc. and is a function of how well the system is designed.
- Pile-up noise. This occurs when two calorimeter signals overlap. It is a function of the luminosity of the machine, the granularity of the detector, and the length of the pulses produced by the calorimeter shaping circuits. Although the calorimeter will be of a high granularity to reduce the occurrences of pile-up, the level-1 trigger system will use only a reduced granularity.
- Quantisation noise. This is a source of noise for digital signals only. It is the rounding error which occurs in the analogue-to-digital conversion process.

All of these sources of noise lead to possible errors in the BCID process by causing distortion of the calorimeter pulses from their ideal shape. One further cause of distorted pulse shapes is saturation. Energy deposits in the calorimeter which are greater than 255 GeV will saturate the level-1 ADC system, producing digital pulses with truncated peaks. At some energy greater than 255 GeV the front-end electronics for the calorimetry will also become saturated, such that it no longer possesses a linear time-response. The pulses produced in this situation will generally be of a longer duration than non-saturated pulses. Both of these effects – the truncation of large pulses and the non-linear time-response of the front-end electronics – present problems for the BCID process.

On a global level, the bunch-crossing identification of an interaction of interest can be derived from the bunch-crossing identifications of the individual signatures which constitute that event. Looking at the signatures listed in table 2.2, it can be seen that any event of interest will, in general, produce a large $E_T$ deposit in the calorimeter (either from an e/$\gamma$ cluster or a high-$E_T$ jet) which can be used to identify the event as a whole. Bunch-crossing identifications made using large calorimeter signals (from large $E_T$ deposits) will generally be more accurate than those made using small signals, due to an improved signal-to-noise ratio (the exception being saturated pulses). BCID for smaller $E_T$ deposits (< 5 GeV) is less important as it will not directly determine the bunch-crossing identification of an event as a whole. It will however, affect the performance of the triggers themselves. Poor BCID for small pulses will result in small energy deposits being missed or allocated to the wrong bunch crossing. This will degrade the sharpness of the jet, missing-$E_T$ and e/$\gamma$ triggers and weaken the effectiveness of the isolation veto for the e/$\gamma$ trigger.

The requirements for the BCID filter are therefore that it should provide bunch-crossing identification with good efficiency for high-$E_T$ pulses, with good efficiency for low-$E_T$ pulses being desirable but less important. The crucial factor in determining the efficiency of the filter, particularly for low-$E_T$ pulses, will be its immunity to noise.

Another important requirement is that the filter be of simple design, introducing a minimum additional latency to the level-1 trigger.

## 4.2. Bunch-Crossing Identification Algorithms

The purpose of the BCID filter can be defined as this: to take a pulse which spans several bunch crossing and identify it with only one of those bunch crossings. If it is digital signals which are being processed, then those signals will have been digitised by taking samples synchronously with the bunch-crossings. The requirement of the BCID filter in this case is to take a pulse which occupies a finite number of samples and transform it to an impulse, which occupies only one sample.

It should be noted that normally, although digital signals consist of a series of discrete samples, the timing information which can be extracted from the signals is not limited to the discrete set of times at which the samples were taken. For example, if dealing with a bi-polar pulse of known shape but unknown phase with respect to the sampling clock, it is possible to locate the point of zero crossing by examining the relative heights of the samples, even though in general the point of zero-crossing will not fall directly on a sample. In this way a measure of the phase between the pulse and the samples can be obtained. For the level-1 trigger however, this procedure is not required.

In order to cope with saturated pulses of the type described in the previous section, it is believed that the BCID system for the ATLAS level-1 calorimeter trigger will require two separate algorithms: one optimised for non-saturated pulses and one optimised for saturated pulses. Most pulses will be processed by the former of these two algorithms, but when a saturated pulse is identified (by its height), it will be processed instead by the saturated-pulse algorithm. This separate algorithm is required as most of the timing information associated with a calorimeter pulse (such as the peaking time and the point of zero-crossing) is destroyed or corrupted by saturation, leaving only the rising edge of the pulse from which to extract reliable information.

All of the BCID algorithms described in this thesis have been developed to process non-saturated calorimeter pulses. They identify the timing of a calorimeter pulse by looking for some characteristic feature which can be identified with a single sample. For example, the peak of the pulse. All of the algorithms studied consist of either an FIR (Finite Impulse Response) filter, a peak-finding algorithm, or a combination of the two (namely an FIR filter followed by peak finder). Below are described the theories of the FIR filter and the peak finder, followed by a description of each of the candidate BCID algorithms in turn.

### 4.2.1. The Finite Impulse Response Filter

The Finite Impulse Response (FIR) Filter is an example of a discrete linear system which is invariant in time. To express these conditions more formally:

1. A discrete system is one which converts an input data set, $x(n)$, into an output data set, $y(n)$ :

   $$x(n) \rightarrow y(n)$$

   where $n$ is an integer representing the sample number.

2. The definition of a linear system is that the equation

   $$x_1(n) + ax_2(n) \rightarrow y_1(n) + ay_2(n)$$

   should be true for any value of $a$.

3. A time-invariant system is one for which

   $$x(n + n_0) \rightarrow y(n + n_0)$$

   is true for any integer $n_0$.

The behaviour of a discrete, linear, time-invariant system is governed by a convolution equation, i.e. the output from the system is given by a convolution of the input data with the impulse response of the system. The impulse response of the system is defined to be the output produced by the system when the input consists of a unit impulse at time zero. The set of values, $h(n)$, which constitute the impulse response completely define the system. The definition of an FIR filter is a filter that has a finite set of such values.

The equation which describes the behaviour of an FIR filter can be written as

$$y(n) = \sum_{i=0}^{N-1} h(i)x(n-i)$$

| | | | |
|---|---|---|---|
| where | $h(i) = a_i$ | if | $0 \leq i \leq N-1$ |
| and | $h(i) = 0$ | elsewhere. | |

Here the impulse response of the FIR filter consists of a set of $N$ coefficients, $a_i$ . The output sample, $y(n)$, is determined by a sum of the $N$ previous input samples, weighted by the coefficients $a_i$. A hardware implementation of an FIR filter is shown in Figure 4.2. The series of latches shown at the top of the figure form a pipeline memory which contains the $N$ previous data samples. These samples are multiplied by the coefficients $a_i$ and summed using an adder tree.

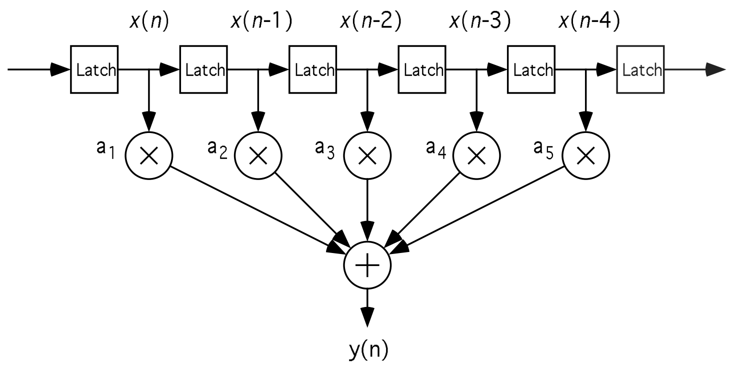


Figure 4.2. A hardware implementation of the FIR filter.

It is possible to construct Infinite Impulse Response (IIR) filters by the use of recursive techniques: feeding the output of the filter back into the input. These filters have an infinite memory and effectively perform a convolution involving all past samples. They have no practical application for bunch-crossing identification however.

### 4.2.2. The Peak Finder

This is an example of a feature extraction algorithm which is used to identify local peaks within a data set. It can be used either in isolation, as the whole of the BCID algorithm, or in conjunction with an FIR filter, in which case it is fed the filtered data. Whenever the peak finder identifies a peak in the data it generates a flag and ultimately only flagged samples are sent to the trigger processor.

The input to the peak finder is a set of consecutive data samples: $x(n), x(n-1), x(n-2)$.... For convenience and clarity it is best to denote these samples by the nomenclature $R_1, R_2, R_3$... where the sample with the largest subscript number is the most recent. With every clock cycle the samples shift forward by one position and $R_1, R_2, R_3$... take new values.

As an example, consider a peak finder which takes as its input three consecutive samples $(R_1, R_2, R_3)$. For this algorithm to recognise a peak within these three samples it requires that they fulfil the condition:

$$R_1 < R_2 > R_3$$

If this is true then there is obviously a local peak in the data sample at the position $R_2$ and a flag will be generated and attached to one of the samples. Which sample this will be depends on whether the peak finder is examining filtered or unfiltered data. The purpose of the BCID logic is to pass on to the level-1 trigger processor the information from each calorimeter pulse for one bunch crossing only: the flagged bunch crossing. For the trigger processor to make the most accurate measurement of $E_T$, it is the sample which contains the peak of the unfiltered calorimeter pulse that should be flagged. This will not necessarily coincide with the peak of the filtered pulse being examined by the peak finder, but will be offset from it by a fixed number of time slices according to the FIR filter used. This offset needs to be taken into account when determining which sample to flag.

The number of samples examined by the peak finder and the criteria by which a peak is defined can be used to control the response of the algorithm to different pulse shapes and to different levels of noise. Any noise on the signals from the calorimeter will have two effects:

- it will produce fluctuations in the pedestal level of the calorimeter signal which will cause false peaks to be identified when there is no true pulse present;
- it will cause pulses to be distorted from their ideal shape.

This latter effect becomes increasingly important as the magnitude of the calorimeter signal falls and the signal-to-noise ratio decreases. In the ideal situation, with no electronic noise, the rounding error of the digitisation process will always limit the information carried on low energy pulses and compromise the pulse shapes. Figure 4.3 shows an extreme example of this.

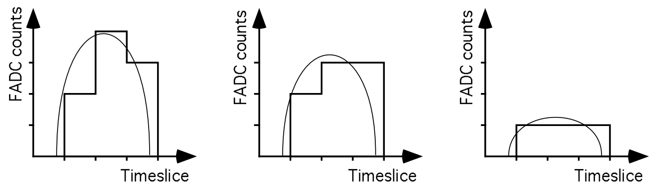


Figure 4.3. The effect of quantisation noise on low energy signals

False flags generated entirely from background noise are less serious than incorrect flags generated from misshapen pulses. If a pulse is wrongly identified a large amount of energy is incorrectly entered into one bunch crossing, whilst being left out of the correct one. For background flags only small energy discrepancies are involved and the net result is an increase in noise on the signals sent to trigger processor.

The use of a larger number of samples in the criteria used by the peak finder to define a peak makes the peak finder more selective. For example:

$$R_1 < R_2 < R_3 > R_4 > R_5$$

$$R_1 < R_2 > R_3 > R_4 > R_5.$$

The second of these could be used to search for a pulse which was of an asymmetric shape, with a sharp rising edge followed by a long falling edge.

The use of a large number of coefficients cuts down on number of false flags generated by fluctuations in the pedestal but also reduces the number of real peaks found in the low energy region, where a low signal-to-noise ratio causes signals to be increasingly distorted from their ideal shape. For a less selective peak finder one could use criteria such as

$$R_1 \leq R_2 > R_3.$$

By using 'greater or equal to' requirements in the peak definition one can identify those pulses which have peaks that cover more than one sampling. This situation is more prevalent at low energies, where quantisation noise causes the flattening of peaks. See figure 4.3. The accompanying disadvantage of using such wide criteria is that they lead to an increase in the rate at which flags are generated by pedestal fluctuations. One further alternative is to use 'either/or' conditions when defining the peak. For example:

$$R_1 \leq R_2 > R_3 \quad \text{OR} \quad R_1 < R_2 \geq R_3 .$$

The best peak-finding criteria for a given situation ultimately depend upon the shape of the pulse.

Peak-finding algorithms can be easily implemented in hardware using a set of comparators. The outputs from the comparators can be combined using the required logic functions to give the desired algorithm. Figure 4.4 shows a simple example.

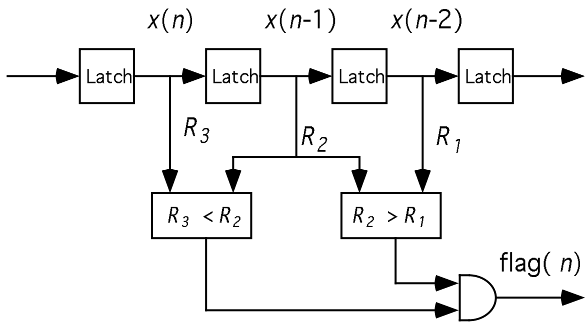


Figure 4.4. The Hardware implementation of a peak-finding algorithm with the criterion $R_1 < R_2 > R_3$.

## 4.3. The Candidate Algorithms

This section describes in turn each of the algorithms investigated during the course of this project.

### 4.3.1. The Deconvolution Filter

The aim of the deconvolution filter [22] is to deconvolute the shaped calorimeter pulse back into an impulse signal. This impulse will contain both timing and energy information. Figure 4.5 illustrates the theory behind the deconvolution filter.

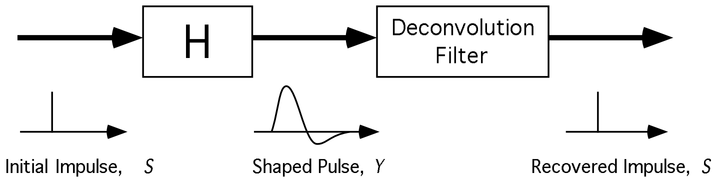


Figure 4.5. The deconvolution filter

The characteristic pulse shape seen by the BCID filter can be regarded as having been produced by some system, $H$ (labelled by its impulse response), in response to an initial input of impulse, $S$. This is a hypothetical system: in reality the pulse shape is the product of the calorimeter shapers acting on the initial calorimeter signal (which is not an impulse) and the subsequent digitisation by the ADC system.

The output from the system $H$ is given by a convolution of the input, $S$, with the impulse response of the system, $H$. As $S$ itself will always be an impulse, the output, $Y$, will simply be the impulse response multiplied by the height of $S$. For the special case where $S$ is unitary the output from the system will be the impulse response itself – a situation which can be expressed by the matrix equation:

$$Y = H.S = H$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} h_1 & 0 & 0 & \dots \\ h_2 & h_1 & 0 & \dots \\ h_3 & h_2 & h_1 & \dots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ \vdots \end{bmatrix}.$$

The elements of $H$ can therefore be written down from the inspection of the ideal pulse shape (i.e., the right-hand side of the above equation). If $H$ is known it is possible to calculate its inverse matrix, $W$. This matrix can be used to reconstruct the original impulse from the shaped output:

$$W = H^{-1}$$
$$S = H^{-1}.H.S$$
$$S = W.Y$$
$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} w_1 & 0 & 0 & \dots \\ w_2 & w_1 & 0 & \dots \\ w_2 & w_2 & w_1 & \dots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ \vdots \end{bmatrix}$$

This equation again illustrates the case where $Y$ is equal to $H$ and a unitary impulse has therefore been reconstructed.

In general the matrix $W$ will be of infinite dimensions. For certain pulse shapes though it will contain only a finite number of non-zero or non-negligible elements. In such cases the coefficients of the $W$ matrix ( $w_1$ , $w_2$ , $w_3$...) can be used as the coefficients of an FIR filter. In the ideal case (i.e., with no noise) the output from such a filter will consist of a series of impulses, each representing one calorimeter pulse and having a height proportional to the $E_T$ deposit in the calorimeter. This data can therefore be sent directly to the trigger processor.

The drawback of using the deconvolution filter is its effect on noise. The filter used to perform the deconvolution has the adverse effect of increasing the noise present on signal. For this reason it is only suitable for low noise signals.

### 4.3.2. The Stand-Alone Peak Finder

The term 'stand-alone peak finder' is used here to denote that case where BCID is performed using just a peak-finding algorithm (with no prior FIR filter). Here, timing information is extracted from the calorimeter pulse by identifying its peak. The peak however, is not necessarily the optimum feature of the pulse from which to extract timing information.

For analogue pulses, timing information should ideally be extracted from that part of the pulse where the gradient, d$V$/d$t$, is at its maximum. At this point any shift in the timing of the pulse, $\Delta t$, will produce the greatest change in signal height, $\Delta V$. Conversely, any noise on the signal, $\Delta V$, will produce the smallest change in the extracted time information, $\Delta t$ . For maximum sensitivity to small time shifts and minimum sensitivity to signal noise, timing extraction should therefore be performed by investigating the height of the pulse in this region.

For digital systems the situation is similar: small phase shifts between the original analogue pulse and the sampling clock are most accurately measured by examining the sample heights at positions where *dV/dt* is at its maximum, rather than at the peak, where *dV/dt* is zero. This is not an important consideration for the level-1 BCID logic however, as there is no need to extract timing information about the phase of a pulse relative to its samples. It is the sensitivity of the BCID algorithm to noise which is important.

Noise will produce incorrect results from a feature extraction algorithm when it alters the relative heights of the samples which constitute the feature. What matters is not the gradient of the original analogue pulse over this feature, but the variation in sample height. The sensitivity of the peak finder to noise therefore depends upon the sharpness of the peak of calorimeter pulses when observed in the time-frame of the 25 ns sampling clock.

### 4.3.3. Zero-Crossing Identifier

The zero-crossing identifier is more commonly implemented as an analogue method of extracting timing information from analogue signals. Those signals must necessarily be bi-polar. As the gradient of a pulse at the point of zero crossing is non-zero, timing measurements made from this point are more precise than those made from the peak.

The digital implementation of this algorithm uses an FIR filter to integrate the incoming calorimeter pulse, forming a peak where formerly the point of zero-crossing occurred. A peak-finding algorithm is then used to identify those peaks. The advantage of using this method of zero-crossing identification, rather than a feature extraction algorithm which searches for the feature

$$R_1 > \text{Pedestal} \quad AND \quad R_2 < \text{Pedestal}$$

is that no knowledge of the local pedestal height is required. An algorithm which compares the data samples with some preprogrammed pedestal level is susceptible to errors due to pedestal drift and pile-up.

The integration of the calorimeter pulse is performed by convolving it with a set of unity coefficients (1, 1, 1...). Only a finite set of coefficients is required as the integration need not be over all time, only over the duration of a single pulse. The minimum number of coefficients required to transform the point of zero-crossing into a peak is the number required to cover the positive region of the pulse.

The subsequent peak-finding algorithm should be optimised for the shape of the pulses produced by the FIR filter.

### 4.3.4. Constant Fraction Discriminator

The constant fraction discriminator [23] is another algorithm usually implemented as an analogue signal-processing method. It extracts timing information from the rising edge of the pulse – the optimum location for analogue signals as here the gradient of the pulse is generally at its maximum.

The constant fraction discriminator uses the basic principle of a threshold counter, identifying that time when the pulse height rises above some set threshold. The use of a fixed threshold however, is not suitable for pulses of varying height, as the time at which the threshold is crossed will depend upon the total pulse height. Instead the constant fraction discriminator sets a threshold which is a constant fraction of the total pulse height.

Figure 4.6 illustrates the theory behind the constant fraction discriminator. For clarity the pulses are shown as analogue signals. The initial input pulse from the calorimeter is sent down two channels. Down one of these it is delayed by a time $\Delta t$ and multiplied by some number, $X$. The value of $\Delta t$ is chosen such that the rising edge of the delayed pulse coincides with the peak of the undelayed pulse. The delayed pulse is then subtracted from the undelayed pulse and the result is a pulse which has a point of zero-crossing at the time when the delayed pulse is equal to $1/X$ of its total height. The point of zero-crossing is then identified by integrating the pulse and using a peak-finding algorithm, as for the zero-crossing identifier.

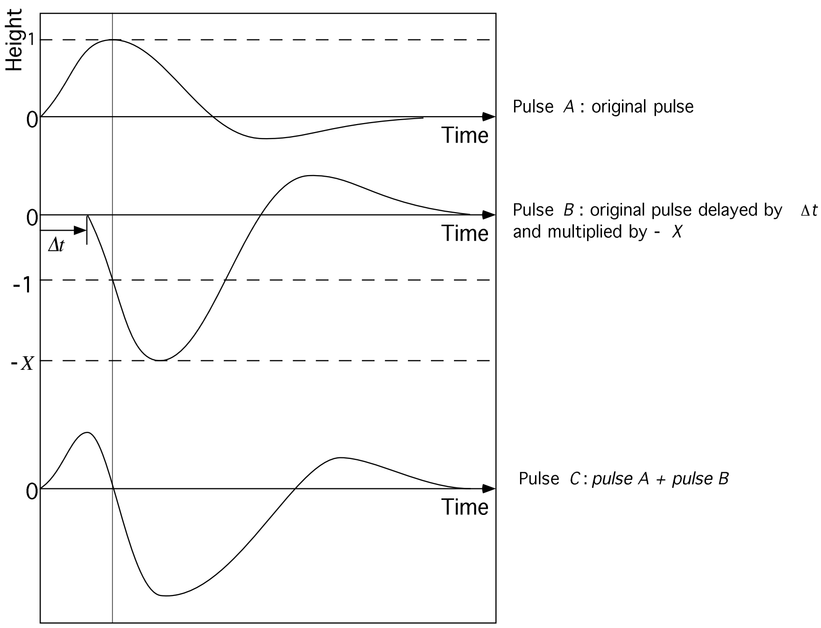


Figure 4.6. The Constant Fraction Discriminator

For a digital constant fraction discriminator, the generation of the delayed pulse and its subtraction from the undelayed pulse can be performed using an FIR filter. A second

FIR filter is then needed to integrate the resultant pulse. The two convolutions required for this process can easily be combined though, and performed by a single FIR filter. This is shown in figure 4.7. Appendix A explains the notation used in this figure.

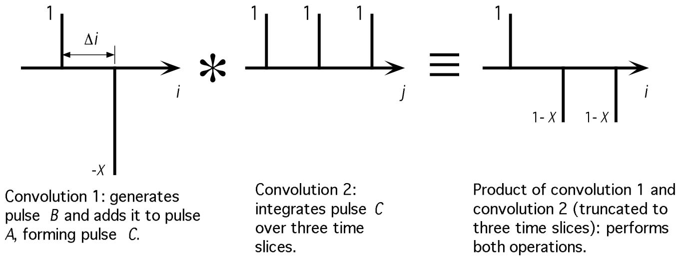


Figure 4.7. The FIR coefficients required to implement a constant fraction discriminator. The width of $\Delta i$ corresponds to the number of timeslices required to generate the delay $\Delta t$ shown in figure 4.6. It is assumed in this figure that $\Delta i = 1$. The pulses *A, B* and *C* referred to here are those shown in figure 4.6.

When dealing with digital signals rather than analogue signals there are additional factors which must be taken into consideration when determining the optimum values of $\Delta t$ and $X$. The most obvious is that $\Delta t$ is limited to a set of discrete values. When determining the optimum value of $X$, regard must be made to the fact that the height of a digital pulse is not a continuous function of time (as it is for analogue pulses). The threshold of the constant fraction discriminator can therefore be set at a level corresponding to the height of one of the samples on the rising edge of the pulse, or at a level which falls between the heights of two consecutive samples.

If the threshold is set to correspond to the height of one of the samples, then even a small amount of noise on that sample will be able to raise or lower the sample height above or below the threshold level, shifting the time at which the constant fraction discriminator fires by one whole time slice. For maximum immunity to noise the threshold should therefore be set midway between consecutive sample heights.

The peak finder required by this algorithm should be optimised for the shape of pulses produced by the FIR filter.

### 4.3.5. The Comparison of the Pulse with its Integral

This algorithm is very similar to the constant fraction discriminator. It fires when the pulse height is equal to a constant fraction of its own integral. An analogue implementation of this algorithm was proposed as the BCID system for the GEM collaboration on the SSC [24].

The digital implementation of this algorithm requires two convolutions. The first of these multiplies the original pulse by a factor $X$, and subtracts from it the integral of the original pulse. The result of this is a pulse which has a point of zero-crossing when the original pulse height is equal to *1/X* of its own integral. A second convolution is then required to integrate this pulse, before sending it to a peak finder. As with the constant fraction discriminator, these two convolutions can be combined and performed by a single FIR filter. Figure 4.8 illustrates this. The same considerations as for the constant fraction discriminator apply when optimising the value of $X$.

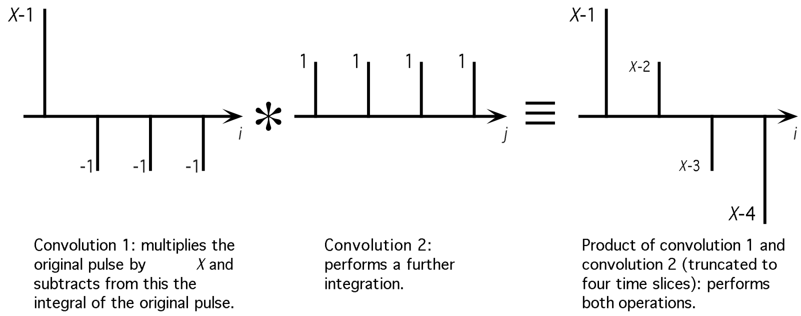


Figure 4.8. The FIR coefficients required to perform the comparison of the pulse with its integral.

### 4.3.6. Peak-Sharpening Algorithms

These algorithms are essentially the same as the stand-alone peak finder, except that they use an FIR filter to sharpen the peak of the pulse before trying to identify it. The optimum coefficients for use by the FIR filter depend upon the shape of the calorimeter pulse. The basic form is of positive coefficients in the centre of the FIR filter, to give a large positive output when the pulse is centred on the filter and negative coefficients at the extreme, to reduce the magnitude of the output when the pulse is off-centre. The net effect is to reinforce the peak of the pulse, whilst suppressing the samples on the rising and falling edge. This operation is then followed by a peak finder.

## 4.4. Conclusion

The testing of the above algorithms is described in the following chapter. It is already possible to draw one important conclusion however. All of the above algorithms possess a very similar structure: an FIR filter followed by a peak finder. The stand-alone peak finder and the deconvolution filter, whilst not explicitly consisting of this structure, can be considered as possessing a transparent FIR filter and peak finder respectively. Any hardware system consisting of a programmable FIR filter and a programmable peak finder would therefore be able to implement all of the above algorithms and more besides. This would allow the FIR coefficients and peak-finding criteria to be optimised for the shape of pulses observed at the LHC. Such a system is proposed for the level-1 calorimeter BCID system. It is foreseen as existing in parallel to a second system which will handle saturated pulses. It should be remembered that the front-end electronics for the ATLAS calorimeters have not yet been fixed and neither therefore has the expected shape of the pulses. Certainly the shape of pulses will vary between different calorimeters. The large degree of flexibility provided by this proposed BCID system is therefore essential.

# 5. BCID software simulation

## 5.1. Introduction

All of the algorithms described in section 4.3 have been simulated in FORTRAN and tested with real data from a prototype of an ATLAS calorimeter. The data were recorded at a CERN test beam using a technical demonstrator of the level-1 trigger system connected to the RD3 prototype LAr accordion calorimeter. Two sets of data were used for the algorithm tests. The first was recorded in November 1992, the second during April and May 1993 [25].

The primary purpose of the level-1 beam tests was to test the phase-1 prototype Cluster Finding Module (CFM) [26]: a module which implemented an early version level-1 e/γ trigger. It was also recognised however, that the data recorded from the calorimeter could be used to test candidate BCID algorithms.

Data were recorded with a wide variety of beam conditions, including pions and electrons at a range of energies, from 10 to 400 GeV. The pulses observed from the calorimeter were of the characteristic shape shown in figure 4.1.

## 5.2. Test Beam Set Up

### 5.2.1. April/May 1993

Figure 5.1 shows a block diagram of the system used to record the data during the April/May beam tests.

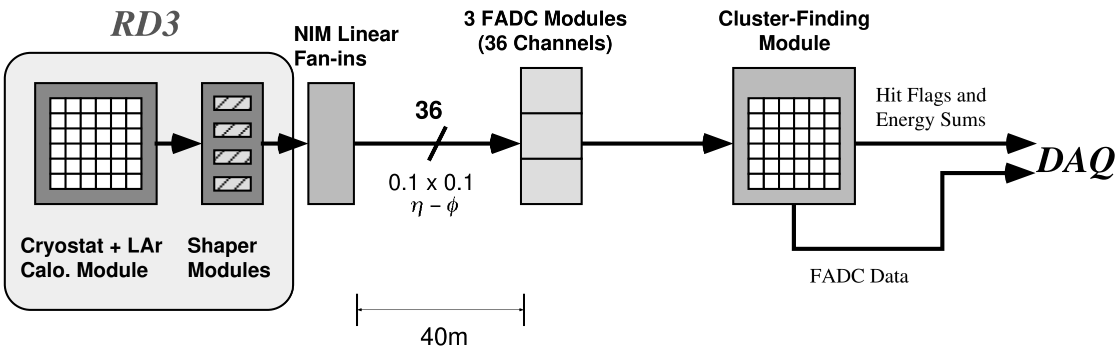


Figure 5.1. Block Diagram of the calorimeter and trigger demonstrator system used in the April/May 1993 beam tests.

The level-1 trigger cells were formed from a $4 \times 4$ area of calorimeter cells, summed over the first two layers of the electromagnetic calorimeter. This was achieved using linear fan-in modules to sum the analogue signals output from the calorimeter shaping modules. The summed signals were then transmitted down 40 m long BNC 50 Ω coaxial cables to the level-1 demonstrator system.

The analogue signals were digitised by FADCs using a free-running 40 MHz clock and a TDC was used to record the phase of the clock relative to the arrival of the test beam particle. The FADCs had a calibration of 0.9 GeV/count. Once digitised the data were passed on to the CFM, which implemented a version of the e/γ trigger algorithm. The output of this module contained the results from the trigger algorithm and the unprocessed data (as received from the FADCs). Only the unprocessed data were of interest for the BCID tests.

The data recording process was triggered by a signal from a scintillator counter placed in the beam line, upstream of the calorimeter. On receipt of this signal the system clock module produced a programmable number of clock cycles before freezing and waiting for the next event signal. These clock signals were used to drive the FADCs and the CFM, such that for each event a programmable number of samples were recorded. Programmable delays were used to adjust the position of the calorimeter pulse within the time-frame of these samples, so that all of the pulse could be observed. The data were read out from the trigger using a VME-based data-acquisition (DAQ) system [27] and recorded using an exabyte storage system.

### 5.2.2. November 1992

The system used to record the November 1992 data was a simpler version of that later used to record the April/May data. The CFM was not present for the November 1992 beam tests and the data passed straight from the FADCs to a Mac-driven data acquisition (DAQ) system. The data were recorded on floppy disks before being transferred to a hard disk on CERNVM. The FADCs used in the November 1992 test beam were not the same as those used in the later tests, but were borrowed from the H1 drift-chamber readout system [28]. They had a calibration of 1.4 GeV/count.

At the time of the November test beam, the proposed bunch-crossing period of the LHC was 15 ns and so most of the data were recorded using a 67 MHz clock. This was the original design frequency of the demonstrator system.

## 5.3. Data Playback.

The data recorded during the beam tests required sorting before they could be used to test the BCID algorithms.

With the trigger demonstrator system, the sampling process was driven by a free-running clock and the calorimeter pulses arrived at the FADC with a random phase with respect to the clock. At the LHC however, sampling will be driven by the LHC clock and will be synchronised with the bunch crossings. The phase of the ADC strobe will be adjusted so that each pulse is sampled at its maximum. To replicate this situation with the data recorded by the trigger demonstrator, the recorded TDC value was used as a measure of the phase between the original analogue signal and the sampling clock. Only events which had a phase within a defined window were used to test the BCID algorithms. This window was centred on the region which corresponded to the pulse being sampled at its maximum. The width of the window was determined by a compromise between the need for a narrow window to more accurately reflect the situation at the LHC and the need for a large number of test events. For the November 1992 data the window was 5 ns ($2/3\ \pi$) wide for the data recorded with a 15 ns clock, and 10 ns ($4/5\ \pi$) wide for the data recorded with a 25 ns clock (this was due to the smaller volume of these data). For the April/May 1993 data, all of which were recorded with a 25 ns clock, the window was 5 ns ($2/5\ \pi$ ) wide.

In addition to random noise, the data recorded in the April/May test beam showed the presence of coherent pick-up noise. This was generated by one of the demonstrator systems power supplies and it was visible on around 15% of all events as a 1.7 MHz sinusoidal oscillation of the pedestal level. Events which displayed this pick-up were removed from the test data by an algorithm which searched for oscillations of the characteristic frequency in a trigger cell which was not targeted by the beam.

To analyse the performance of the BCID algorithms as a function of pulse height, the remaining events were sorted into an energy spectrum. To obtain a large number of test events at the low energy range of the spectrum, pion data were used, together with data runs where the beam had been targeted between two trigger cells. These data produced a range of low-energy signals.

## 5.4. Results

### 5.4.1. November 1992

During the November 1992 test-beam, data were recorded using both a 25 ns sampling period and a 15 ns sampling period. As 15 ns was the proposed LHC bunch-crossing period at that time, most of the data were recorded with this sampling period. The emphasis of early BCID tests was also placed on evaluating the performance of the algorithms with these data. When the proposed bunch-crossing period was changed to 25 ns, the results of these tests became less relevant. Nonetheless, they did allow several important conclusions to be drawn, most notably with respect to the deconvolution filter.

The deconvolution filter was the first of the BCID algorithms to be studied. Its performance was tested on data recorded both with a 15 ns clock and a 25 ns clock, with the same method being used for both. Ideal pulse shapes were formed from the summation of many pulses (~1000) to increase the signal-to-noise ratio. The coefficients of the deconvolution matrix ($w_1$, $w_2$, $w_3$... ) were then extracted from this pulse shape by means of a FORTRAN program.

Unfortunately the shape of pulse observed in the test data produced a deconvolution matrix with an infinite number of coefficients (for both the 15 ns and 25 ns data). This would have been acceptable if the coefficients tended towards zero, but in this case the coefficients tended towards infinity. The deconvolution process could therefore not be approximated by an FIR filter. Deconvolution was also tried using pulses which had undergone pedestal subtraction and suppression of the negative lobe of the pulse (an operation which is proposed for the 'final' level-1 trigger system), but here too the resultant matrix was infinite.

All of the other algorithms described in section 4.3 were tested with the November data. As the stand-alone peak finder was the simplest algorithm, its performance was used as the baseline with which to compare that of the other algorithms.

For the data recorded with a 15 ns clock, the stand-alone peak finder achieved 100% efficiency for pulses of a height of ten counts (14 GeV) or greater. The efficiency here is defined as the percentage of pulses correctly identified. Below this level the peak finder still achieved an efficiency of 88% for pulses down to a height of four counts (5.6 GeV). These results were based on a data sample of around 100 events per energy level. The performance of the stand-alone peak finder showed that bunch-crossing identification could easily be performed for large pulses and that it was only for small pulses that the situation was more demanding.

The only algorithm to improve significantly on the performance of the stand-alone peak finder for the data recorded with a 15 ns clock was the constant fraction discriminator. This algorithm achieved 100% efficiency down to a pulse height of ten counts, and 98% efficiency down to a pulse height of four counts.

The results produced for the data recorded with a 25 ns clock were based on much lower statistics – around 30 events per energy level. It was clear however that all of the BCID algorithms showed a significant improvement in performance when working with the 25 ns clock as compared to the 15 ns clock. The stand-alone peak finder was the most efficient of the algorithms in this case, achieving 100% efficiency down to a height of three counts above pedestal. The other algorithms produced efficiencies which were only a few percent below this, allowing no real distinction between the performances of the algorithms because of the low-statistic data.

The BCID algorithms gave an improved performance for the data recorded with the 25 ns clock because when the analogue signal is sampled at 25 ns intervals rather than 15 ns intervals, the resultant digital signal in the sampling time-frame is significantly closer to an impulse. This is the ideal situation for the stand-alone peak finder, which explains why this algorithm had the highest efficiency of all the algorithms for the 25 ns data. If a smaller sampling period is used, as was the case for the 15 ns data, the resultant digital signal is closer to the original analogue signal. This explains why the constant fraction discriminator, the optimum analogue signal processing method, had the highest efficiency of the algorithms tested on the 15 ns data.

The limiting factor on the efficiency of the algorithms for all of the data was quantisation noise, rather than electronic noise. This was shown by the high efficiencies of the algorithms for pulses of only a few counts height, and the very low rate at which false bunch-crossing identifications were generated from the background pedestal. Fewer than 0.1% of pedestal samples produced a false bunch-crossing identification which contained a deviation from the pedestal height of two counts or greater. The R.M.S. fluctuation of the pedestal due to noise was found to be ~1 count.

A wide variety of FIR coefficients and peak-finding criteria were tested with the November 1992 data. From the results obtained for both the 25 ns data and the 15 ns data several conclusions could be drawn:

- The number of coefficients used in any FIR filter could always be reduced to a maximum of five samples, with no loss in the efficiency of the algorithm.
- The values of the coefficients used in an FIR filter could be limited to powers of two, with no loss in the efficiency of the algorithm. (The implication of this is that in a hardware implementation of the BCID algorithm the multiplication required for the FIR filter can be performed by bit-shift operations, reducing the latency of the BCID algorithm.)
- The use of complex peak-finding criteria involving many samples and logical *OR* conditions did not improve the performance of a BCID algorithm compared to the use of more simple criteria.

These conclusions were used to determine which algorithms were tested more exhaustively with the high-statistic 25 ns data, recorded at the April/May test beam. The results obtained from the November 1992 data were published as an RD27 internal note [29].

### 5.4.2. April/May 1993

The data recorded in the April/May '93 test-beam period which were used to test the BCID algorithms are shown in table 5.1. This table shows the number of events in each run before and after the TDC and pick-up-noise cuts described in section 5.3. The beam position is given in terms of trigger-cell coordinates. The area of calorimeter visible to the trigger was a $6 \times 6$ array of trigger cells, where each cell was of a granularity corresponding to $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ in the barrel. Where more than one trigger-cell coordinate is given in table 5.1, the beam was targeted between cells and data from all of the cells in which the beam was visible were used to test the algorithms. The break-down of the test events with respect to pulse height is shown in table 5.2.

| Run no. | Beam Type | Beam Position | Initial no. of events | No. of events after TDC cut | No. of events after Pick-up cut |
|---|---|---|---|---|---|
| 198 | 200 GeV $\pi$ | (4, 4) | 10, 054 | 2109 | 1694 |
| 212 | 50 GeV $\pi$ | (4, 4) | 14, 397 | 2958 | 2572 |
| 213 | 50 GeV $\pi$ | (2, 4) | 12, 879 | 2691 | 2236 |
| 262 | 30 GeV $\pi$ | (4, 4) | 11, 129 | 2313 | 1720 |
| 227 | 50 GeV e | (4, 3), (5, 3) | 7380 | 1544 | 1356 |
| 230 | 50 GeV e | (4, 3), (5, 3), (4, 4) | 4944 | 1038 | 897 |
| 231 | 50 GeV e | (5, 3), (4, 3), (5, 4) | 9264 | 1995 | 1707 |
| 232 | 50 GeV e | (5, 3), (5, 4) | 8258 | 1810 | 1558 |

Table 5.1. The April/May 1993 test data.

| Pulse height / FADC counts | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | ≥11 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Number of Pulses | 2100 | 621 | 428 | 438 | 294 | 401 | 291 | 336 | 289 | 288 | 4958 |

Table 5.2. The number of test pulses as a function of pulse height.

The algorithms tested with the April/May 1993 data are shown in tables 5.3 and 5.4. Table 5.3 shows the FIR filter coefficients used in each case and Table 5.4 lists the different peak finders used. Each algorithm was tested with all of the different peak-finding criteria. It can be seen in table 5.4 that peak-finding criterion 1 uses five samples to define the peak of a pulse. Tests with the November 1992 data had already indicated that using such a strict definition of the peak was undesirable, as it reduced the number of correct bunch-crossing identifications to a greater extent than it reduced the number of incorrect bunch-crossing identifications. Peak-finding criterion 1 was included in the April/May data tests to confirm this result with high-statistic data.

| **Algorithm** | **FIR filter Coefficients** | | | | |
|---|---|---|---|---|---|
| Stand-Alone Peak Finder | 1 | 0 | 0 | 0 | 0 |
| Zero-Crossing Identifier | 1 | 1 | 1 | 0 | 0 |
| Constant Fraction Discriminator #1 | 1 | 1 | -1 | 0 | 0 |
| Constant Fraction Discriminator #2 | 1 | 1 | -2 | 0 | 0 |
| Constant Fraction Discriminator #3 | 1 | -2 | -2 | 0 | 0 |
| Comparison of Pulse With Integral #1 | 1 | 0 | -1 | -2 | 0 |
| Comparison of Pulse With Integral #2 | 1 | 0 | -1 | 0 | 0 |
| Peak-Sharpening Algorithm #1 | -1 | 1 | 1 | -2 | 0 |
| Peak-Sharpening Algorithm #2 | -2 | 1 | 1 | 1 | -2 |
| Peak-Sharpening Algorithm #3 | -1 | 0 | 1 | 0 | -1 |
| Peak-Sharpening Algorithm #4 | -1 | 0 | 2 | 0 | -1 |
| Peak-Sharpening Algorithm #5 | -1 | 0 | 4 | 0 | -1 |

Table 5.3. The BCID algorithms tested with the April/May test data.

| **Criterion Number** | **Peak Definition** |
|---|---|
| 1 | $R_1 < R_2 < R_3 > R_4 > R_5$ |
| 2 | $R_1 < R_2 > R_3$ |
| 3 | $R_1 \leq R_2 > R_3$ |
| 4 | $R_1 < R_2 \geq R_3$ |
| 5 | $R_1 \leq R_2 \geq R_3$ |

Table 5.4. The Peak-finding criteria tested with the April/May test data.

The FADCs used to collect the April/May test beam data were of a higher sensitivity than those used in the November test beam: 0.9 GeV/count. It was hoped that the resultant reduction in quantisation noise would lead to an improvement in the efficiency of the BCID algorithms. The decrease in quantisation noise however was accompanied by an increase in electronic noise, so that this became the dominant source of noise on the observed signals. The resultant performance of the BCID algorithms was slightly worse than for the November 1992 data. The increase in electronic noise was understandable as the FADCs being used here were prototype modules, rather than the fully developed F1001 drift chamber readout FADCs used for the November 1992 beam tests.

The main results for each algorithm are detailed in the following sections.

***The Stand-Alone Peak Finder***

Figure 5.2 shows the results produced by the stand-alone peak finder for each of the five peak-finding criteria. Each graph shows three plots as a function of pulse height:

- The percentage of correct identifications. This is the percentage of pulses correctly identified.
- The percentage of incorrect identifications. This is defined by the equation:

  $$\text{Percentage of incorrect id.s} = \frac{\text{no. of incorrect id.s}}{\text{no. of pulses}} \times 100.$$

  Note that as some algorithms may miss a pulse entirely and others may produce more than one identification per pulse, the number of correct identifications plus the number of incorrect identifications do not necessarily total to 100%. The percentage of incorrect identifications may exceed 100% using this system.
- The background identification rate. This is the percentage of samples which lay off the calorimeter pulse entirely which produced false bunch-crossing identifications due to noise of the specified height above pedestal.

The efficiency of a BCID algorithm should be measured not only by the number of correct identifications made, but also by the number of incorrect identifications made. The rate at which an algorithm produces false identifications from pedestal noise should also be considered. From the graphs it can be seen that the most efficient peak finder was criterion 3 from table 5.4:

$$R_1 < R_2 \geq R_3.$$

Using this criterion, the stand-alone peak finder correctly identified 100% of pulses down to a height of five counts (5 GeV) above pedestal. Below this, it was this criterion which reached the best compromise between the number of correct identifications made and the number of incorrect identifications made. This is to be expected. It can be seen by inspection that this peak-finding criterion is more sensitive to the rising edge of a pulse than it is to the falling edge. As the pulses observed from the RD3 calorimeter had sharper rising edges than they did falling edges (see figure 5.3), it was the rising edge which was least likely to be corrupted by noise. Peak-finding criterion 3 was therefore the optimal criterion for such pulses.

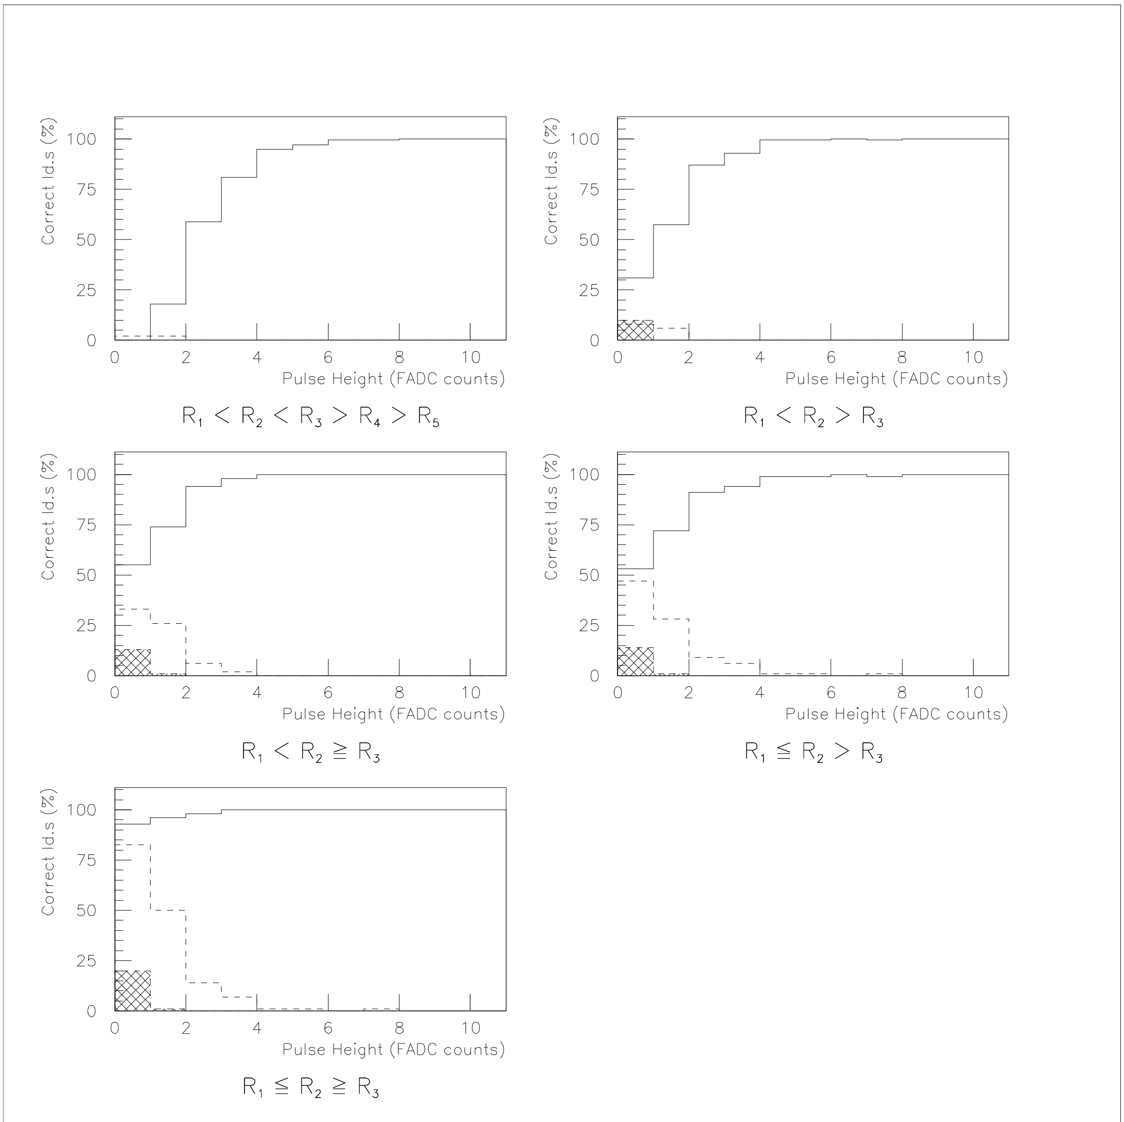


Figure 5.2. Results from the BCID peak finder as a function of energy. Each graph shows the percentage of correct identifications (solid line), the percentage of incorrect identifications (dashed line), and the background identification rate (hatched area).

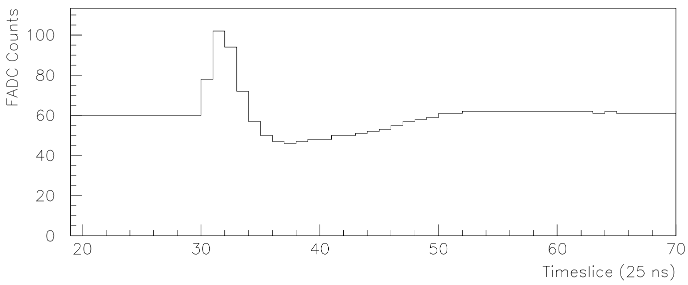


Figure 5.3. A typical pulse from the RD3 prototype LAr calorimeter.

Peak-finding criterion 5:

$$R_1 \leq R_2 \geq R_3$$

produced a greater percentage of correct identifications than criteria 3, but it also produced double the amount of incorrect identifications. This was due to its ability to produce more than one bunch-crossing identification per pulse. At the LHC, inefficient BCID for small pulses (< 5 GeV) will have the greatest effect on the energy sums formed by the jet and missing-$E_T$ triggers, and the isolation sum formed by the e/γ trigger. Doubling the rate at which spurious identifications are generated obviously reduces the accuracy of such measurements to a greater extent than does accepting a loss in efficiency of a few percent. Peak-finding criterion 5 is therefore less suitable for the task of BCID than criterion 3.

The rate at which an algorithm generates background identifications from pedestal noise is less important than the rate at which it generates incorrect identifications for actual pulses, as these will produce much smaller errors in the $E_T$ values sent to the trigger processor. A measure of the importance of these background identifications can be gained from looking at their effect on the e/γ isolation sum. When no pulse was present, the stand-alone peak finder using criterion 3 produced an R.M.S. deviation from the pedestal of +0.14 GeV (+0.14 counts) per sample. Summed over the 12 electromagnetic and 16 hadronic calorimeter cells this would produce an energy deposit of 4 GeV in the e/γ isolation window. If a threshold of 1 GeV is applied to individual triggers cells before they are included in the isolation sum however, the R.M.S. contribution to the isolation sum from background bunch-crossing identifications drops to only 0.3 GeV per bunch crossing. This compares to 0.02 GeV per bunch crossing if BCID is performed using peak-finding criterion 1, and 0.3 GeV if BCID is performed using criterion 5.

At the time of conducting these tests, the background identification rate had a more important implication for the level-1 trigger. At that time it was proposed to use an asynchronous fibre-optic link to transfer the level-1 calorimeter data off the detector. In order to reduce the bandwidth of the data prior to transmission, it was proposed to use zero-suppression. The rate at which bunch-crossing identifications were generated from background noise therefore had an important role in determining what factor of data reduction could be achieved. For the above data using peak-finding criterion 3, if a zero-suppression threshold were set at 0 GeV, 87% of the off-peak samples would be suppressed. If the threshold were set at 1 GeV, 95.5% of the off-peak samples would be suppressed. It was primarily for making measurements such as these that the background identification rate was recorded. These results became less important when a fully synchronous data-transfer system was adopted and zero-suppression became redundant.

***The Zero-Crossing Identifier***

The zero-crossing identifier produced its best results when the number of coefficients in the FIR filter was truncated to three. This corresponded to the number of samples occupied by the positive lobe of the calorimeter pulse. The reason why this produced the best results is illustrated in figure 5.4.

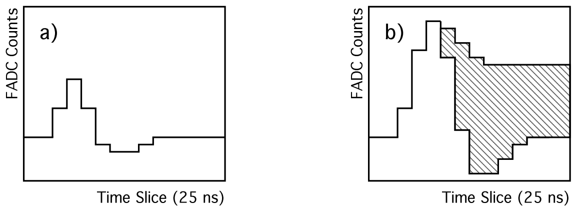


Figure 5.4. Pulse shape a) before FIR filter and b) after integration over three samples (unhatched area) and all samplings (additional hatched area).

When the integration is performed over only three samples, the point at which the integral first encounters negative values (and thus begins to fall in value) corresponds to the point at which it no longer covers the entire positive region of the pulse. The effect of this is to produce an output pulse from the FIR filter with a steeper falling edge than would otherwise be the case. For small calorimeter pulses the shallow negative region of the pulse is rounded to zero by the quantisation process. In such cases the algorithm relies entirely on this secondary effect to produce correct bunch-crossing identifications.

The results produced by the zero-crossing identifier are shown in figure 5.5. They are worse than those produced by the BCID peak finder, a fact accounted for by the shallow nature of the negative lobe of the calorimeter pulse.

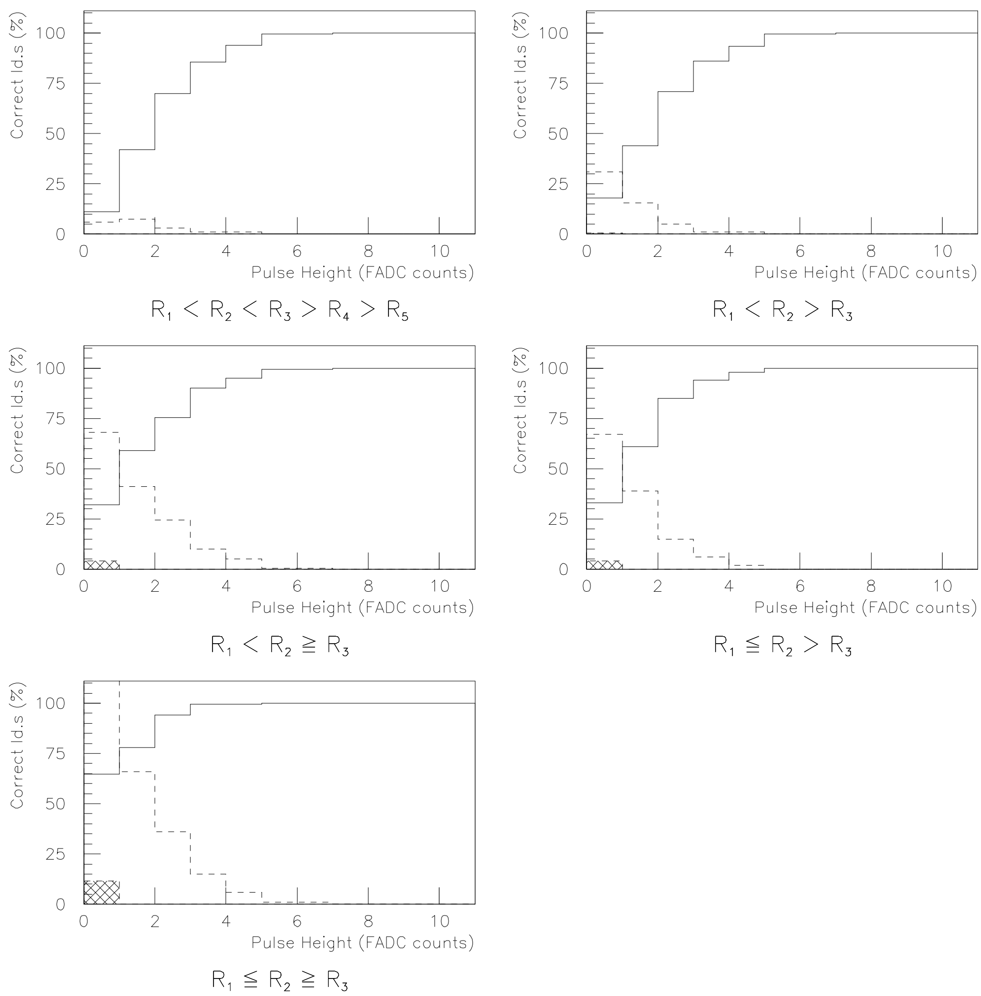


Figure 5.5. Results from the zero-crossing identifier as a function of energy. Each graph shows the percentage of correct identifications (solid line), the percentage of incorrect identifications (dashed line), and the background identification rate (hatched area).

***Constant Fraction Discriminator***

Figure 5.6 shows how the efficiency of the constant fraction discriminator varies as a function of the fraction used. These results were obtained using a value of one time slice for $\Delta t$. As the observed calorimeter pulses had only one sample on their rising edge, this was the only delay which corresponded to the algorithm functioning as a true constant fraction discriminator. The test events used to generate this plot were taken from a subset of the April/May 1993 data containing only pulses of a height of 10 counts (10 GeV) or greater. It has already been shown that BCID can be performed for such pulses with 100% efficiency using a stand-alone peak finder.

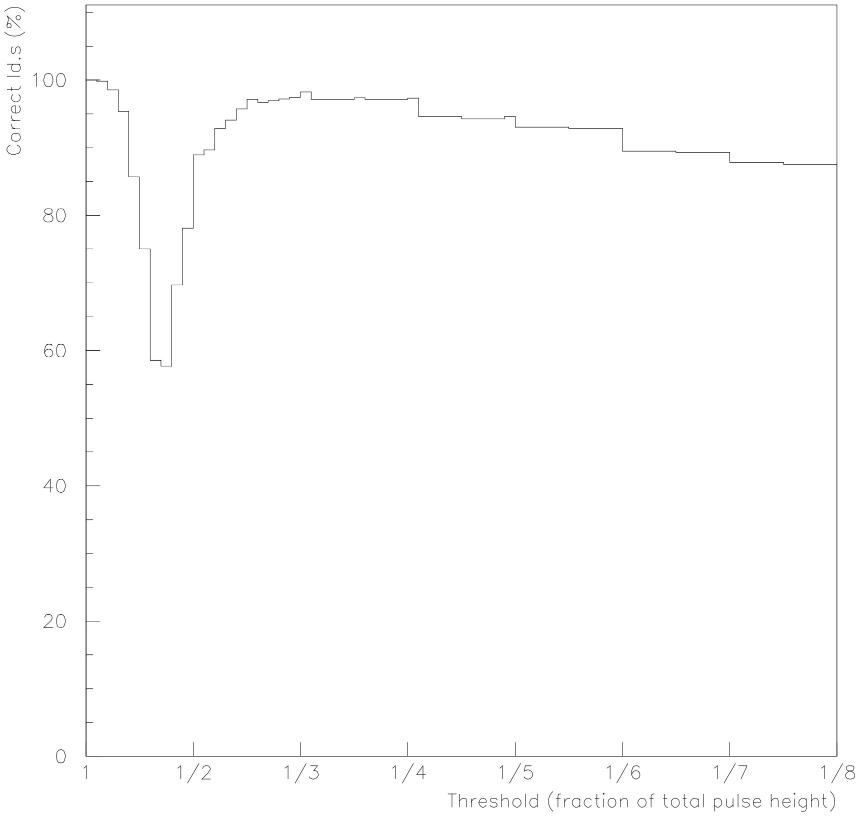


Figure 5.6. The efficiency of the constant fraction discriminator as a function of the threshold used.

From figure 5.6 it can be seen that even for pulses of 10 counts or greater the constant fraction discriminator achieved 100% efficiency only when using a threshold of one. In this situation the constant fraction discriminator was behaving exactly as a peak finder – identifying the point where the pulse reached its total height. This is illustrated in figure 5.7. The efficiency reached a minimum when fractions of around one half were used. This corresponded to the height of the sample on the rising edge of the pulse.

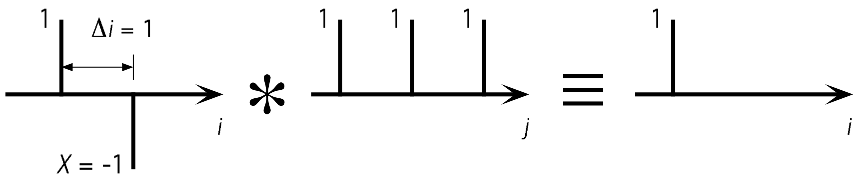


Figure 5.7. The two convolutions necessary to implement a constant fraction discriminator with the values $\Delta i = 1$ ($\Delta t = 1$ timeslice) and $X = 1$, and the single, transparent convolution to which they are equivalent.

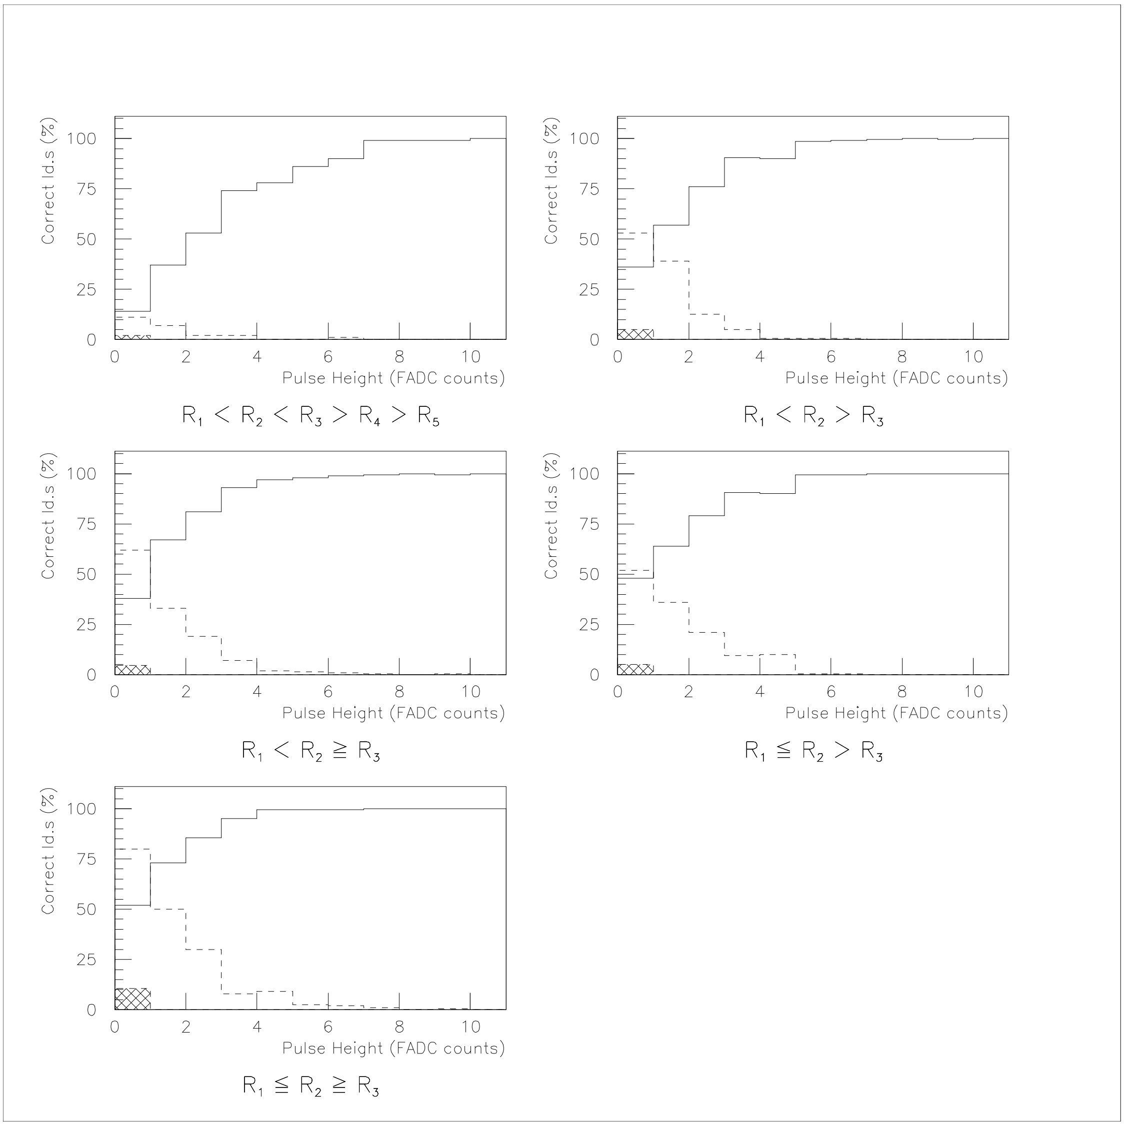


Figure 5.8. Results from the constant fraction discriminator #2. Each graph shows the percentage of correct identifications (solid line), the percentage of incorrect identifications (dashed line), and the background identification rate (hatched area).

Improved results were obtained if a delay of two time slices was used for the constant fraction discriminator. In this case the algorithm was not functioning as a true constant fraction discriminator as it was comparing the height of a point on the rising edge of the pulse, not with the peak of the pulse, but with a point on the falling edge of the pulse. The true constant fraction discriminator measures the height of the pulse at its peak (the optimum place for height measurements as $\mathrm{d}V/\mathrm{d}t$ is at its minimum), and uses this to extract timing information from the rising edge of the pulse (the optimum place for timing measurements as $\mathrm{d}V/\mathrm{d}t$ is at its maximum). The improved performance of this algorithm when using delay of two time slices though demonstrates that such considerations are less important here than they would be if dealing with analogue pulses. From table 5.3, constant fraction discriminator #2 uses a delay of two time slices and a threshold level of one third. The full results for this algorithm are shown in figure 5.8. They are worse than those of the stand-alone peak finder.

***Comparison of the Pulse with its Integral***

This algorithm produced its best results when the FIR filter used was truncated to just three coefficients (1, 0, -1). Figure 5.9 shows a comparison of the results produced by the algorithm when three coefficients were used (1, 0, -1), and when four coefficients were used (1, 0, -1, -2). The results in both cases are worse than those of the peak finder.

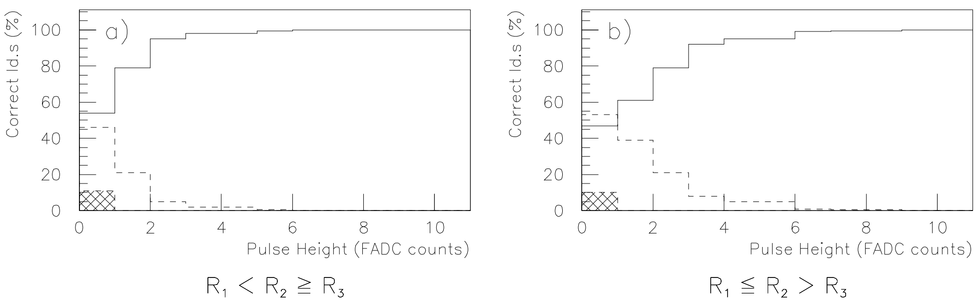


Figure 5.9. Results from comparison of the pulse with its integral, using the coefficients a) 1, 0, -1, and b) 1, 0, -1, -2. Each plot shows the percentage of correct identifications (solid line), the percentage of incorrect identifications (dashed line), and the background identification rate (hatched area).

The convolution performed by the FIR filter in this algorithm has the effect of suppressing the falling edge of the pulse and making the peak sharper. Unfortunately it also has the effect of increasing the noise on the pulse. If the number of coefficients employed by the filter is limited to three, this latter effect is reduced.

***The Peak-Sharpening Algorithms***

Peak sharpening algorithm #4 and peak-sharpening algorithm #5 were the only algorithms to improve on the performance of the stand-alone peak finder. The improvements were minimal. Table 5.5 shows the percentage of pulses identified incorrectly by each algorithm.

All of the algorithms listed in table 5.5 produce only one identification per pulse, so the number of correct identifications is directly related to the number of incorrect identifications. The rate at which background identifications were produced is similar for all three algorithms. The results produced by *peak sharpener #5* are shown in figure 5.10.

| Pulse Height | Stand-Alone Peak finder, $R_1 < R_2 \geq R_3$ | Peak Sharpener #4, $R_1 < R_2 \geq R_3$ | Peak Sharpener #4, $R_1 < R_2 \geq R_3$ |
|---|---|---|---|
| ≥ 5 GeV | 0 % | 0 % | 0 % |
| 4 GeV | 2 % | 1 % | 0.5 % |
| 3 GeV | 6 % | 6 % | 3 % |
| 2 GeV | 25 % | 18 % | 17 % |

Table 5.5. A comparison of the percentage of pulses identified incorrectly by a selection of BCID algorithms.

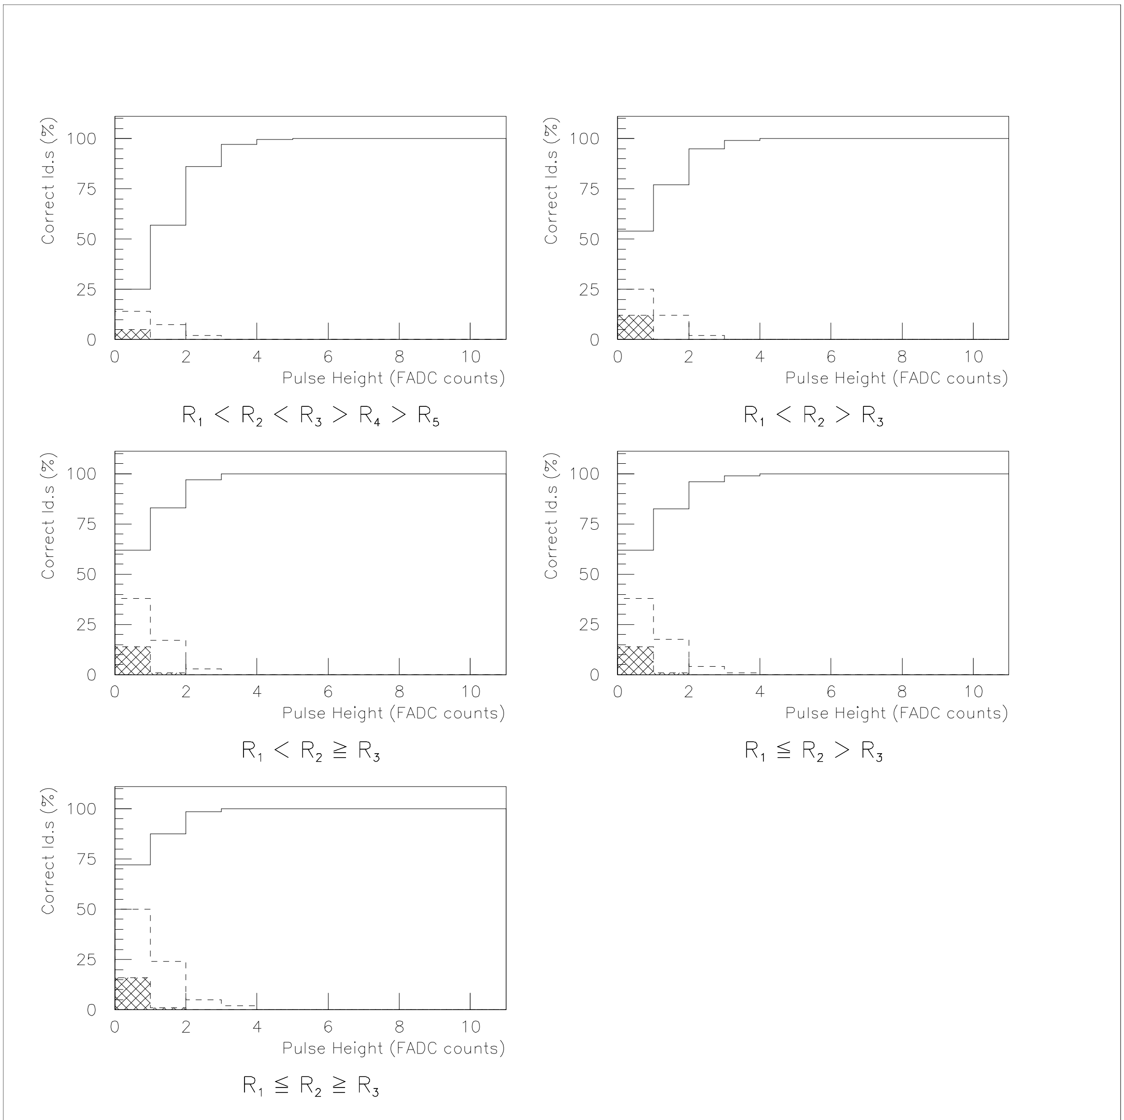


Figure 5.10. Results from peak sharpener #5. Each graph shows the percentage of correct identifications (solid line), the percentage of incorrect identifications (dashed line), and the background identification rate (hatched area).

## 5.5. Summary of results

The results produced by the software tests show that a simple peak-finding algorithm alone can be used to perform bunch-crossing with very good efficiency, down to energy deposits of a few GeV. This level of performance is enough to correctly identify any single pulse large enough to produce an e/γ trigger. More complex BCID algorithms were in general less efficient than the stand-alone peak finder. This is because those algorithms identified the timing of the pulse by searching for features which were more easily swamped by noise than the peak of the pulse. The peak was an efficient feature from which to extract timing information because of the narrow width of the pulse in the time-frame of the 25 ns sampling clock.

Two algorithms were found which could improve on the performance of the BCID peak finder by convolving the data with a simple set of coefficients. The improvements in efficiency produced were only minimal however, and it is questionable whether the improvements produced would warrant the added complexity and latency of the algorithm in the final system.

The results produced by any BCID algorithm will depend critically on the shape of the pulses used. As the shape of pulses that will be used in the final LHC systems is not fixed it is important that the range of algorithms which are studied is maintained as large as possible. For this reason it is important that any demonstrator system constructed for the purpose of bunch-crossing identification should be able to implement a variety of the above algorithms. The results from the software tests of the algorithms were used to formulate a set of rules for a general BCID demonstrator system:

- FIR filters which sum over more than five consecutive samples are not required.
- The coefficients of the FIR filter can be limited to powers of two.

# 6. The BCID Demonstrator Module

## 6.1. Introduction

As part of the design studies for the ATLAS level-1 calorimeter trigger, two BCID demonstrator systems have been constructed. The first of these consists of a single board designed to be used in conjunction with the CFM demonstrator system described in section 5.2. It performs BCID on a single channel of trigger data, implementing BCID algorithms in discrete logic. The use of Look-Up Tables (LUTs) to perform the multiplication necessary for the FIR filters, and the use of programmable comparators for the peak finder enable this board to implement a number of the BCID algorithms described in chapter 4. This board was tested in a CERN test beam in 1993 and the results are described in reference [30].

In 1995 a more comprehensive BCID demonstrator system was constructed. Whereas the single-channel BCID board was built to demonstrate the feasibility of implementing digital BCID algorithms in hardware, this second BCID system was constructed to study the effect of BCID on the CFM demonstrator in a test-beam environment. It consists of three identical boards which perform BCID on a total of 36 trigger channels – the number of channels processed by the CFM demonstrator. Field-Programmable Gate-Arrays (FPGAs) are used to implement the BCID logic. They provide a much more compact solution than discrete logic without incurring any of the great expense or development time of an ASIC. Their field-programmable nature allows different BCID algorithms to be downloaded at will, providing a desirable degree of flexibility to the demonstrator system.

This chapter describes in detail the design of the 36-channel BCID system. The boards themselves are discussed in section 6.5, whilst sections 6.2 to 6.4 describe the FPGA-based BCID logic.

## 6.2. A Hardware Implementation of the BCID Algorithms

All of the BCID algorithms described in chapter 4 consist of a single FIR filter followed by a peak finder. The hardware implementation of such a system is obvious and is shown in figure 6.1. Like the rest of the level-1 trigger, the BCID logic is implemented as a systolic array. This form of architecture requires the algorithm to be broken down into functions of sufficient simplicity that they can be completed within one clock cycle. These functions are then allocated fixed compute stations within the system and are connected via a network of data paths and delays which ensure that all of the arguments required at a compute station arrive within the same clock period.

The latency of the system is determined by the number of clock cycles required to implement the given algorithm.

In the logic shown in figure 6.1, incoming data are stored in a pipeline memory whilst BCID is performed. Data samples are tapped from this pipeline by the FIR filter, which performs a weighted sum of the samples and writes the result to a second pipeline memory. This memory contains the filtered waveform. The peak finder then examines this filtered waveform and upon finding a peak, gates out the corresponding data sample from the original pipeline memory. If no peak is found in the filtered waveform for a given timeslice, the output from the BCID logic for that timeslice is suppressed to zero. Different algorithms employ different FIR coefficients and peak-finding criteria.

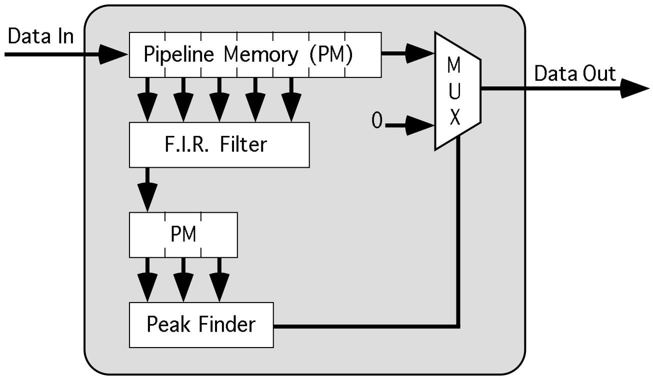


Figure 6.1. Block Diagram of the BCID Logic.

## 6.3. Field-Programmable Gate-Arrays

An FPGA is a user-programmable logic device. It is configured for a specific application by downloading configuration data to an on-chip static RAM. The BCID demonstrator system uses Xilinx XC4008 devices, capable of implementing logic functions of up to ~8000 gates in size [31].

Architecturally, an FPGA is composed of three different configurable elements: logic blocks, I/O blocks and programmable interconnections. The configurable logic blocks, which provide the functional elements from which the user-specified logic is constructed, are arranged in a matrix in the centre of the device. They contain LUTs to perform combinatorial logic functions (adders, comparators etc.) and storage elements in the form of flip-flops. The data path through a logic block is controlled by multiplexers set up during configuration. Configuration also specifies the functionality of the LUTs.

The XC4008 contains 324 identical logic blocks, arranged in an 18 × 18 array. Figure 6.2 shows the design of these logic blocks. Each logic block contains three LUTs to implement Boolean logic functions and two 1-bit registers for data storage. As an example of the computational power provided by such logic blocks, four blocks are required to implement an 8-bit adder.

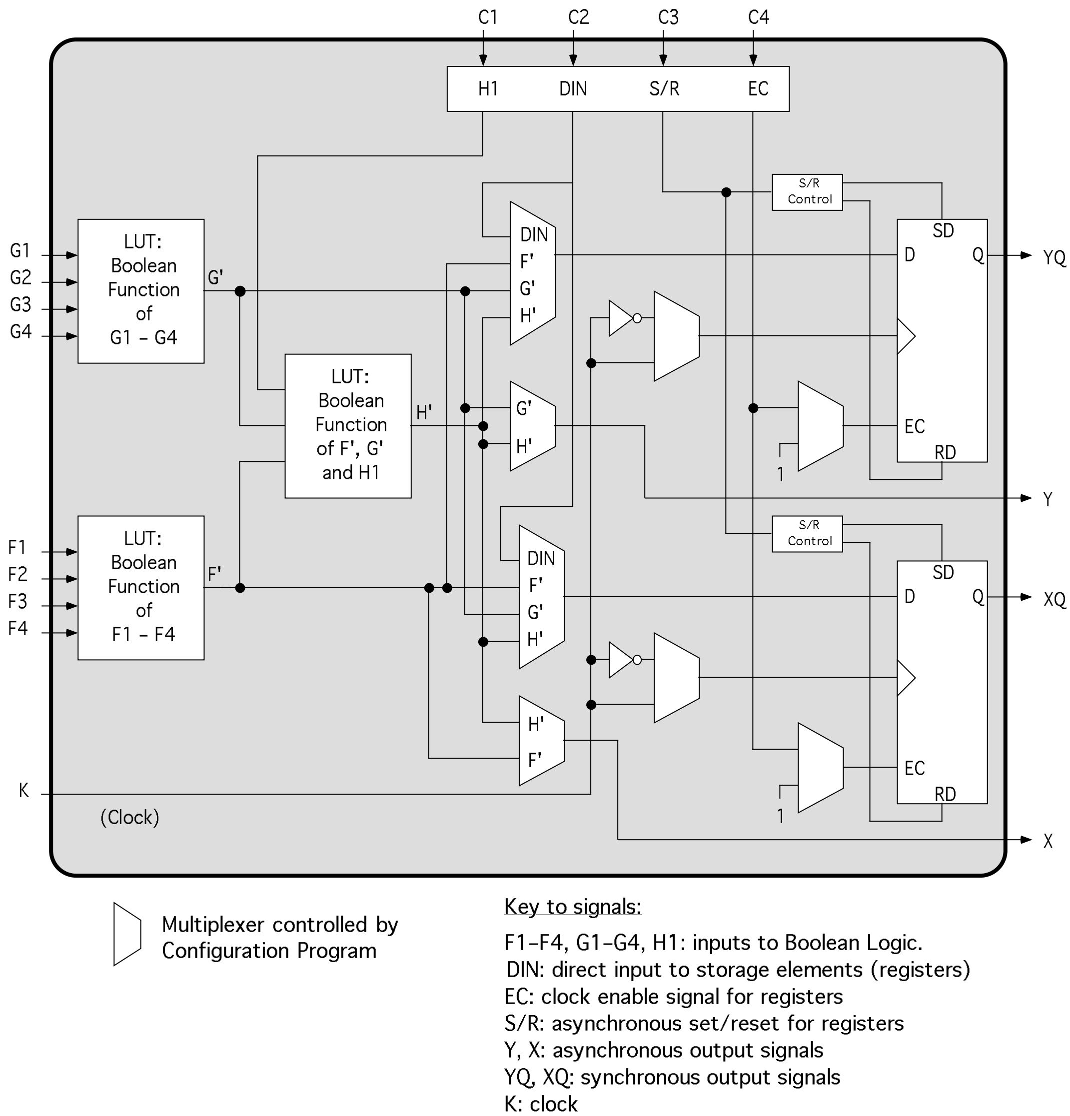


Figure 6.2. XC4000 series FPGA Logic Block.

Signals are transmitted between logic blocks via three types of configurable interconnection. *General purpose* connections form a grid surrounding the logic blocks. Intersection points on the grid are occupied by switching matrices which route signals according to the downloaded configuration. *Long line* connections also form a grid around the logic blocks, but these bypass the switching matrices. They are intended for use with high fan-out signals which must travel long distances with a minimum propagation delay. Additional *long lines* driven by a global buffer are provided for the distribution of a clock signal. The third type of connection, the *direct connection*, provides a high-speed link between adjacent logic blocks. The input and output connections of each logic block to the *direct connection* network, the *long lines*, and the local segments of the *general purpose* network are controlled by switches on that logic block which are set up during configuration.

Signals enter and exit the FPGA via the I/O blocks, which provide an interface between the internal logic and the device package pins. Each I/O block contains both input and output buffers and the downloaded configuration determines not only whether a pin is to act as an input or output, but also whether the voltage thresholds will be TTL or CMOS compatible. Before configuration, all I/O levels are TTL compatible.

The configuration of an FPGA can be performed at power-up or upon command. A variety of modes exist for the transfer of the configuration data, and setting a 3-bit pattern at the 'mode pins' of the FPGA informs the device of which mode is to be used. In 'Master Parallel' mode for example, the FPGA retrieves the data from an external memory, supplying both the memory addresses and the clock signal itself (using an internal crystal oscillator). Alternatively, in 'Peripheral' mode, a micro-processor can be used to supply both the configuration data and a clock signal to the FPGA, which provides a READY/BUSY signal as a handshake. A total of six configuration modes exist for the XC4008 device. The configuration bit-stream itself is generated by proprietary software, and is the end result of producing a logic design for an FPGA (as discussed in more detail in the following section). The length of the configuration bit-stream is independent of the complexity of the logic design and depends only upon the size of the device to be configured (147,505 bits for an XC4008).

## 6.4. Implementation the BCID algorithms in an FPGA

Developing a logic design for an FPGA is a process with two distinct stages. The first stage is to specify the functionality of the logic design. This may be done using either schematic capture, a hardware design language such as ABEL, or some combination of these two methods. The design at this stage may be kept independent of the technology upon which it is to be implemented, or such device-specific features as I/O pins may be allocated via the schematic or ABEL file etc.. Simulation may be performed before progressing any further to ensure that the logic specified is functionally correct. Once this has been completed, the next stage is to target the logic design towards FPGA technology. For most designs this is a largely automated process. Software supplied by the FPGA manufacturer can target the design towards a specific FPGA, or the smallest FPGA in a given family which will accommodate the specified logic. The routing of designs is optimised to maximise the speed of the logic and a configuration bit-stream is produced. At this stage the design can be simulated once again, to verify the functionality and the timing of the logic.

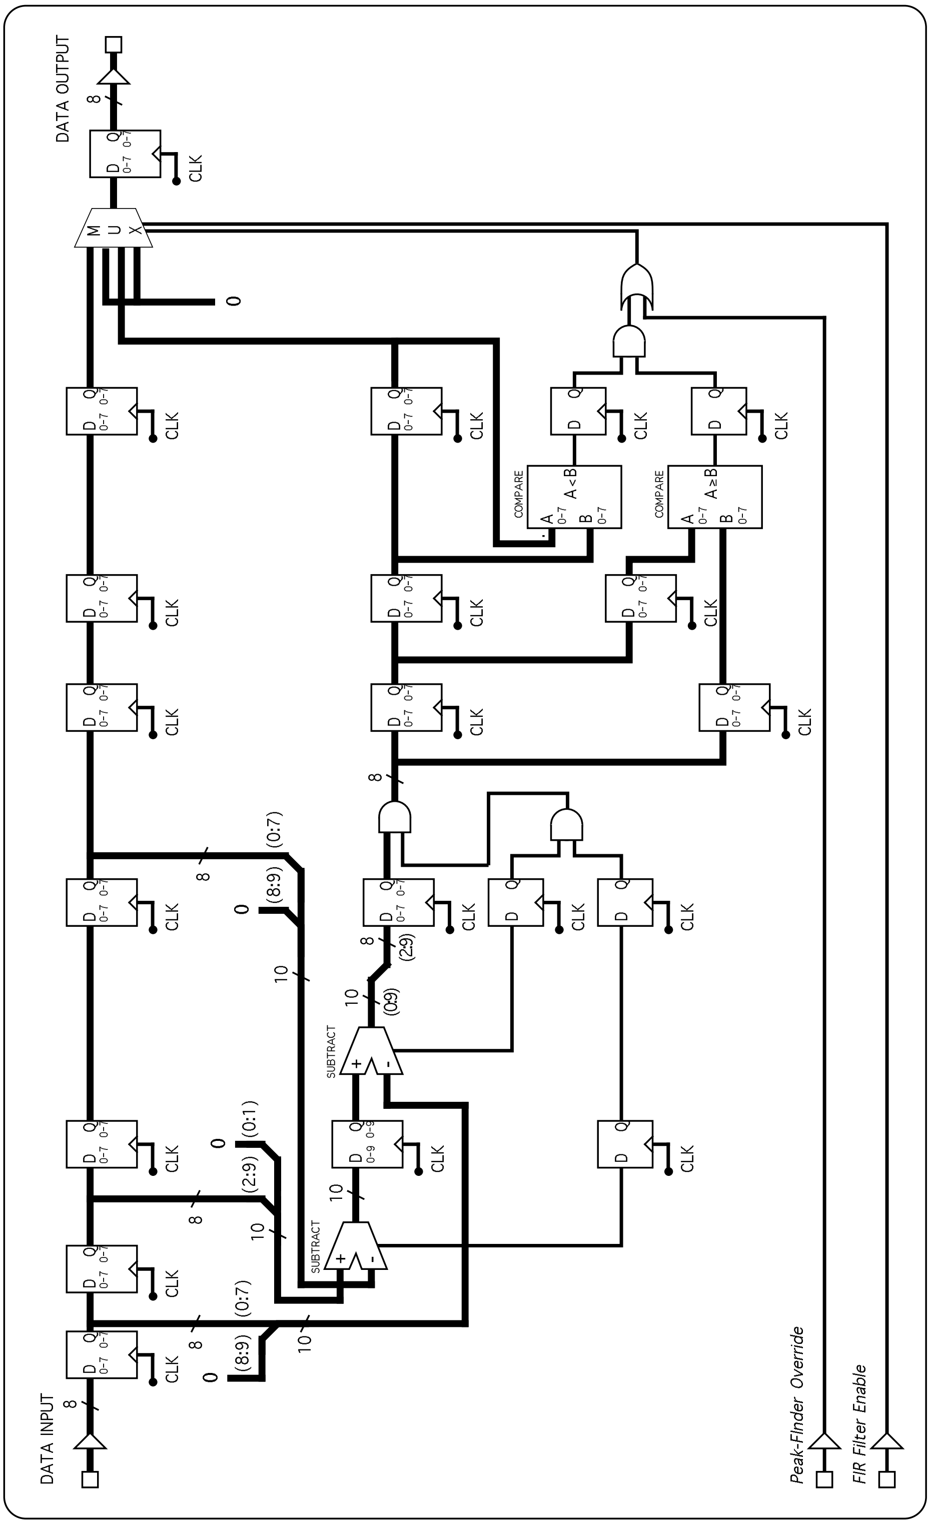


Figure 6.3. A schematic diagram of an FPGA-based BCID logic design. The FIR coefficients used in this design are –1, 0, 4, 0, –1 and the peak-finding criterion is $R_1 < R_2 \geq R_3$.

The BCID logic was designed using schematic capture. CAD software supplied by Viewlogic was used for this process, together with a components library supplied by Xilinx. Figure 6.3 shows one of the logic designs produced. Functionally, this design is a slight expansion of that shown in figure 6.1. For diagnostic reasons the range of data which can be output from the design has been expanded to include signals which are not visible in the normal mode of operation, such as the filtered waveform produced by the FIR filter. The data to be output are selected by setting the *Peak Finder Override* and *FIR Filter Enable* inputs, as shown in table 6.1. When operating in 'BCID mode', this design is functionally identical to that shown in figure 6.1. Appendix B fully explains the schematic diagram shown in figure 6.3.

An FPGA implementation of the logic design shown in Figure 6.3 was developed using the XACT software package, supplied by Xilinx. Once the functionality of the design had been verified by simulation, the primary concern was to ensure that a clock of period 25 ns could be used with the design. It was chosen to target the design towards FPGAs of the XC4000 family (the most advanced family available), which are quoted as being able to support clocks down to a minimum period of 20 ns. This figure assumes that a minimal amount of processing is performed per clock cycle. For a particular design the minimum clock period which can be used is actually determined by the complexity of the functions which must be carried out in one clock period. The first attempt at implementing the BCID logic in an FPGA produced a design which had a minimum clock period of 29 ns (as verified by simulation).

| **Status of *Peak Finder Override* input** | **Status of *FIR Filter Enable* input** | **Mode of Operation** |
|---|---|---|
| **0** | **0** | *BCID Mode*: full BCID is performed. |
| **0** | **1** | *FIR + BCID Inspection Mode*: the filtered waveform produced by the FIR filter is output from the logic, gated by the peak finder result. |
| **1** | **0** | *Transparent Mode*: data are passed through the logic unchanged |
| **1** | **1** | *FIR Inspection Mode*: the filtered waveform produced by the FIR filter is output directly from the logic (with the result of the peak finder being overridden). |

Table 6.1. Modes of Operation for the BCID logic.

There are various ways to optimise the timing of an FPGA logic design. At the schematic level, the maximum acceptable delay for a data path can be defined, and it is possible to specify the order in which nets are routed. The nets which have the most critical timing can therefore be routed first, taking a more direct path through the connection resources of the FPGA. When using the XACT automatic routing program, there are many constraints and parameters which can be used to optimise the routing of a particular design. Unfortunately all of these methods failed to produce a BCID design which could be used with a clock of period less than 26 ns, so manual routing had to be performed. The XACT design editor provides an interface which allows routed designs to be modified or created at the FPGA level, with the user directly controlling the configuration of logic blocks and connection resources. This is an arduous process and is best attempted only after all other methods of controlling the implementation of a design have been exhausted. For the BCID logic however, it produced the desired result, with the final design being able to accept clock signals with a minimum period of 21 ns. This is the guaranteed worst-case minimum clock period, as predicted by the Xilinx software.

A total of 10 different BCID algorithms were converted into FPGA designs. These algorithms are listed in table 6.2. It can be seen that only FIR filter coefficients which are powers of two have been employed, so that all of the necessary multiplication can be performed by bit-shift operations. The number of non-zero coefficients in a single design has also been limited to three, to minimise the size of FPGA required to implement the designs.

| **FIR filter coefficients** | **Peak-Finding criteria** |
|---|---|
| 0, 0, 1, 0, 0 | $R_1 < R_2 \geq R_3$ |
| 0, 0, 1, 0, 0 | $R_1 < R_2 > R_3$ |
| 0, 0, 1, 0, 0 | $R_1 \leq R_2 > R_3$ |
| 0, 0, 1, 0, 0 | $R_1 \leq R_2 \geq R_3$ |
| -1, 0, 2, 0, -1 | $R_1 < R_2 \geq R_3$ |
| -1, 0, 2, 0, -1 | $R_1 \leq R_2 > R_3$ |
| -1, 0, 2, 0, -1 | $R_1 \leq R_2 \geq R_3$ |
| -1, 0, 4, 0, -1 | $R_1 < R_2 > R_3$ |
| -1, 0, 4, 0, -1 | $R_1 \leq R_2 > R_3$ |
| -1, 0, 4, 0, -1 | $R_1 < R_2 \geq R_3$ |

Table 6.2. The BCID algorithms implemented as FPGA designs.

Each BCID design performs BCID on four channels of data in parallel on one Xilinx chip, so that logic shown in figure 6.3 represents only one quarter of a complete design. All of the designs have a latency of eight clock cycles. Some of the designs (most notably those which use the FIR coefficients 0, 0, 1, 0, 0) could be implemented with a shorter latency than this (five clock cycles), but it is desirable to have the same latency for all of the algorithms.

## 6.5. The BCID Board

In this section is presented the design of the BCID board. Three identical boards of this type have been constructed, for the purpose of testing and demonstrating the FPGA-based BCID logic system. The boards are wire-wrap, 9U x 40 cm in size, and fit into the crate used by the CFM demonstrator system.

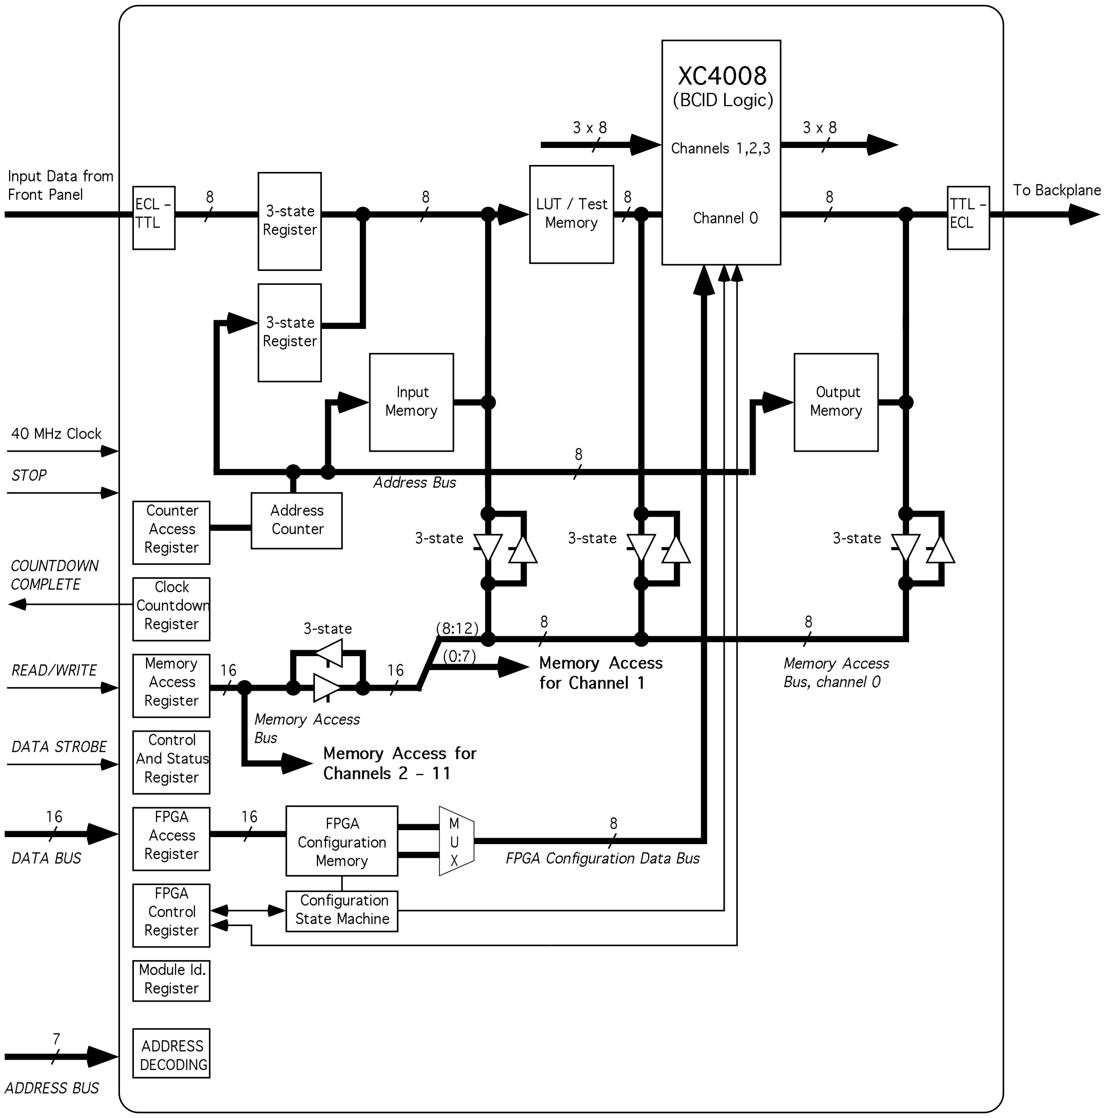


Figure 6.4. The BCID board. Logic is shown for channel 0 only.

Figure 6.4 shows a block diagram of the BCID board. Whilst the full board houses three FPGAs and processes twelve channels of trigger data, only one channel has been shown here for the sake of clarity. There are two modes of operation for the BCID board: ‘Beamline’ mode and ‘Test’ mode. In Beamline mode the board is designed to act as part of a larger demonstrator system, accepting input data from a collection of FADCs and transmitting output data to the CFM described in section 5.2. The data path in Beamline mode is this: twelve channels of 8-bit FADC data are received on to the front panel of the BCID board as differential ECL signals. These signals are converted to TTL levels and latched, before being used to address LUTs which subtract from the data the DC pedestal introduced by the FADCs. The resultant pedestal-subtracted data are then passed through the FPGAs, which perform BCID, and are converted to totem-pole ECL levels before being transmitted to the CFM across the crate backplane. Rolling memories capture the data at two stages on the BCID board: before pedestal subtraction and after BCID. At the end of each event the BCID board is halted by a stop signal from the clock module, so that the memories on the board can be read out by the DAQ.

The second mode of operation, Test mode, is designed to allow the functionality of the FPGA-based BCID logic to be tested in isolation from the rest of the demonstrator system. Test data are downloaded on to the BCID board from the crate controller and are stored in the memories normally used to perform pedestal subtraction. When the board is run in Test mode an incrementing counter is used to address these memories, and the test data are scrolled through the FPGAs at 40 Mbaud. The output of the FPGAs is recorded in the output memories for later examination.

### 6.5.1. Clock Distribution and Timing

Figure 6.5 shows the logic pertaining to clock distribution on the BCID board. A 40 MHz differential-ECL clock signal arrives on the board via the crate backplane. It is converted to TTL levels and gated by two signals: $\overline{\mathit{Stopped}}$ and *Clock Enable*. If either of these signals is low, then the distribution of the clock signal around the BCID board is disabled. The $\overline{\mathit{Stopped}}$ and *Clock Enable* signals are more fully explained in section 6.5.2. It can be seen from figure 6.5 that logic is included to ensure that no glitches are generated when the clock signal is enabled or disabled: the distribution of the clock signal always starts on a rising clock edge and stops on a falling clock edge. When distribution of the external clock is disabled, it is possible to generate single clock ticks internal to the BCID board by pulsing the *Single Tick* signal line. This operation is performed by the Control and Status Register and is intended for diagnostic purposes only.

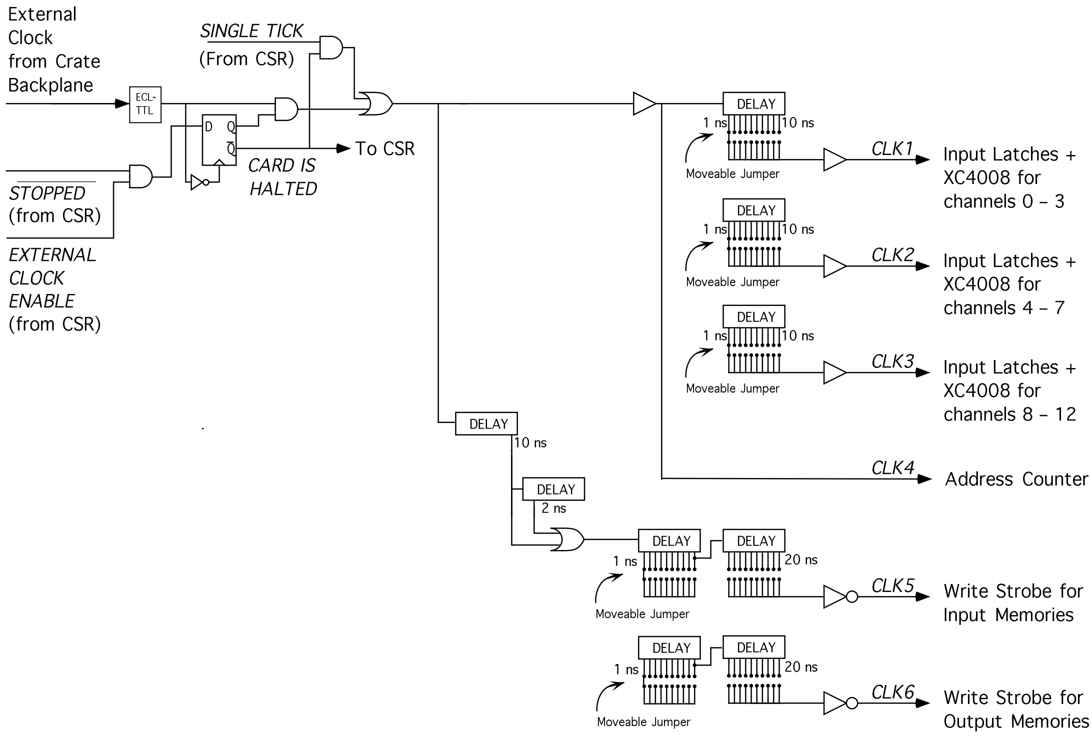


Figure 6.5. BCID board clock tree.

As shown in figure 6.5, the clock signal is fanned out six ways. Passive delay lines are used to correctly align the phase of the clock signals to different parts of the circuit. A delay line is also used to stretch the high-time of the clocks used to provide the $\overline{\textit{Write Enable}}$ signals for the input and output memories. These memories require a write strobe of at least 12 ns duration. The high-time of the received 40 MHz clock signal is 12.5 ns (theoretically), but this could easily be altered by signal reflections and the rise/fall-times of the buffers used in the clock circuit. The delay line is therefore used to stretch the high-time to 14 ns, which provides a margin of error of 1.5 ns.

All of the delay lines on the BCID board are 10 ns long, with taps every 1 ns. The moveable jumper connections shown in figure 6.5 allow different output taps to be used, so that the timing of the clock circuit can be optimised. Figure 6.6 shows the path followed by the data and address signals when the BCID board is running in either Beamline mode or Test mode. Various points on the data path have been labelled *A* to *C,* and figures 6.7 and 6.8 show the timing of the signals at these points. Figure 6.7 illustrates the fact that the same clock signal can be used to drive the front-end latches and the XC4008 devices, due to an internal delay in the distribution of the XC4008 clock. Figure 6.8 shows the delays necessary for the write strobes to the input and output memories (~18 and ~22 ns respectively).

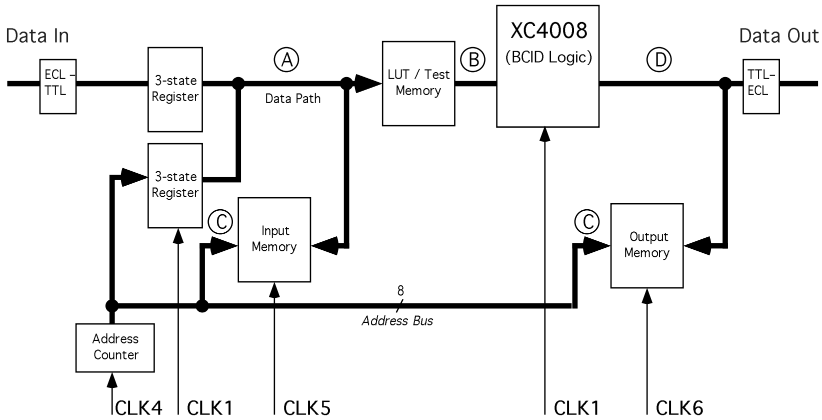


Figure 6.6. Key to timing diagrams.

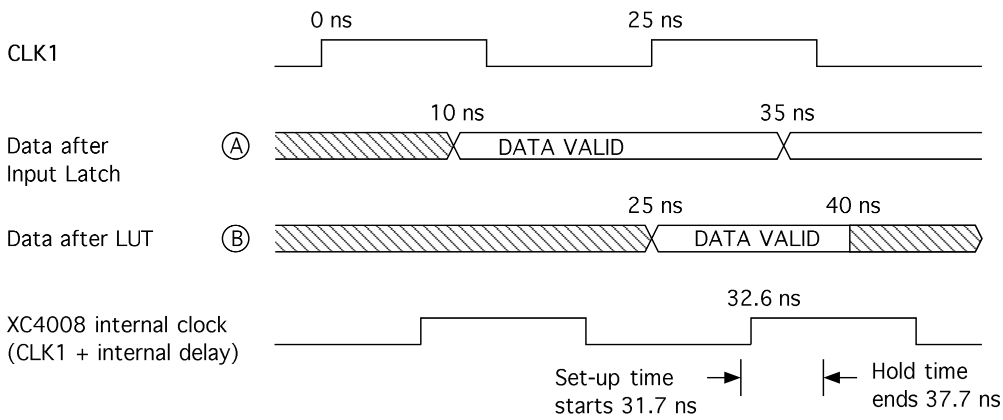


Figure 6.7. The timing of data from the front-end latches to the XC4008s.

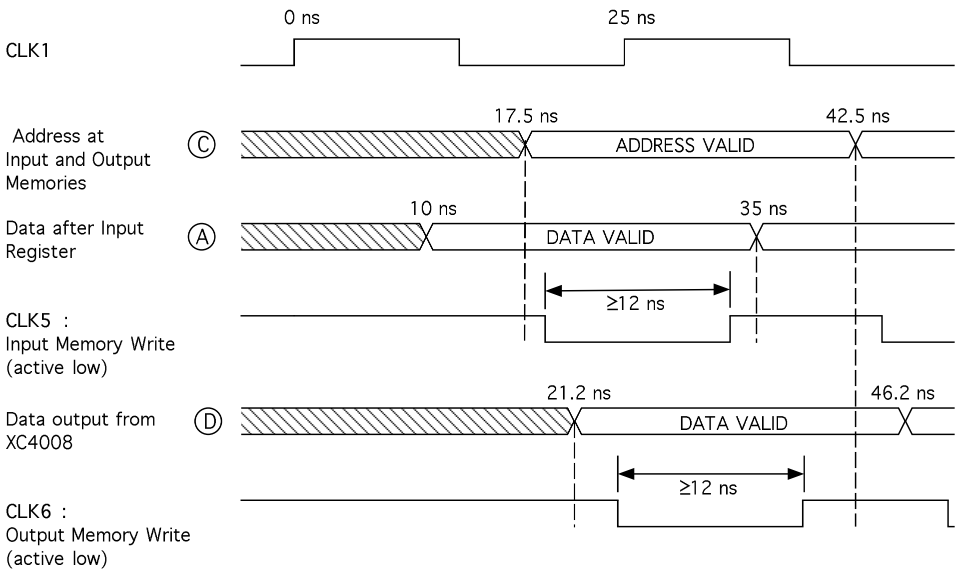


Figure 6.8. The timing of data and address signals to the input and output memories.

### 6.5.2. Slow Control and Data Readout

Slow control and readout of the BCID board are performed using a VME-based DAQ system, which interfaces with the BCID board via a local crate controller. The crate controller provides 16 bi-directional data lines, a 7-bit address bus, a read/write select line and a data strobe. The interface to the crate controller on the BCID board consists of the seven registers listed in table 6.3. With the exception of the read-only Module Identification Register, which contains a hard-wired bit-pattern denoting the type and serial number of the board, all of these registers can be both written to and read. The Control and Status Register (CSR) controls the mode of operation of the BCID module and displays its current status. Table 6.4 lists the bits in this register. It can be seen that there are two types of bits: those which are set by the user (via the crate controller) and those which are set by the board itself. To avoid any conflict in the information arriving at a register, none of the bits can be set by both of these sources. Some of the bits, such as the *Clear Stop* bit, can be repeatedly set high by the crate controller to repeatedly trigger the same operation. These bits are always read back as zero.

| Address location | Register name |
|---|---|
| 0 | Module Identification Register |
| 1 | Control and Status Register |
| 2 | Counter Register |
| 3 | Clock Countdown Register |
| 4 | Memory Access Register |
| 5 | FPGA Control Register |
| 6 | FPGA RAM Access Register |
| 7 | FPGA RAM Counter Register |

Table 6.3. Access registers on the BCID module

The *Stopped, External Clock Enable, Card Is Halted* and *Clear Stop* bits in the CSR, all determine or reflect the status of the 40 MHz clock signal on the BCID board. The *Stopped* bit is set when the board is halted by a *STOP* signal from the clock module. This state is cleared by writing a 1 to the *Clear Stop* bit. The function of the *STOP* signal is to halt all of the modules in the demonstrator system at the end of an event. An alternative method of halting the board is to set the *External Clock Enable* bit low (a situation which is reversed by setting the *External Clock Enable* bit high). If the board is halted for either of these reasons the *Card Is Halted* bit is set, indicating that the external clock is being ignored.

Two further bits in the CSR relate to the clocking of the BCID board: the *Generate Single Tick* and the *Clock Countdown Enable* bits. Both of these are inputs to the board. Writing a 1 to the *Generate Single Tick* bit whilst the BCID board is halted generates a single clock tick which is internal to the board. Setting the *Clock Countdown Enable* bit high whilst the board is running places the board in 'Clock Countdown' mode, such that with every clock tick the number contained in the Clock Countdown Register (see table 6.3) is decremented by one. When this number reaches zero the BCID board transmits a *Countdown Complete* signal via Lemo cable to the external clock module, which issues a stop signal halting all of the boards in the demonstrator system. Clock Countdown mode is designed for diagnostic purposes, to allow the entire demonstrator system to be run for a set number of clock ticks. The Clock Countdown Register can be loaded with any number between 1 and 255.

| Bit | Bit name | Bit type |
|---|---|---|
| 0 | *Channel Pair No.* (bit 0) | Input to card |
| 1 | *Channel Pair No.* (bit 1) | Input to card |
| 2 | *Channel Pair No.* (bit 2) | Input to card |
| 3 | *Clear Stop* | Input to card (Always 0 on read) |
| 4 | *Memory Type* (bit 0) | Input to card |
| 5 | *Memory Type* (bit 1) | Input to card |
| 6 | *Auto-Increment Enable* | Input to card |
| 7 | *Clear Protocol Error* | Input to card (Always 0 on read) |
| 8 | *Generate Single Tick* | Input to card (Always 0 on read) |
| 9 | *External Clock Enable* | Input to card |
| 10 | *Clock Countdown Enable* | Input to card |
| 11 | *Test Mode Enable* | Input to card |
| 12 | *Card is halted* | Output from card |
| 13 | Unused | – (Always 0 on read) |
| 14 | *Stopped* | Output from card |
| 15 | *Protocol Error* | Output from card |

Table 6.4. The Control and Status Register.

The remaining bits in the CSR relate to memory access. Access to the input memories, output memories and LUTs on the BCID board is provided by the Memory Access Register. This register is actually a window which opens on to the memory location specified by the Address Counter (see figures 6.4 and 6.6) and the *Channel Pair Number* and *Memory Type* bits in the CSR. The decoding scheme is shown in figure 6.9. As the input, output memories and LUTs are all 8-bits wide and the data interface to the crate controller is 16 bits wide, the memories are read out in pairs, being grouped according to channel number. Setting the *Channel Pair Number* to 0, for example, specifies channels 0 and 1 on the BCID board.

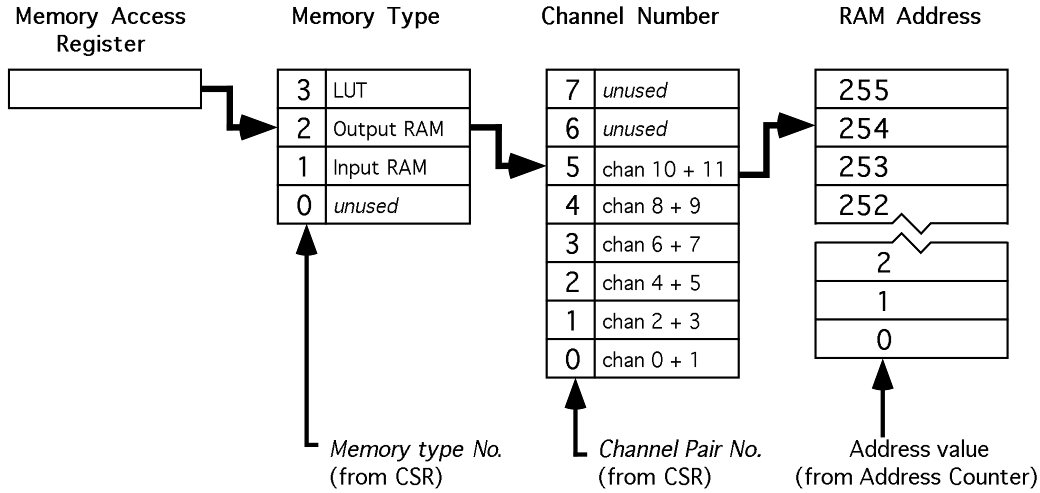


Figure 6.9. Memory Access.

The Address Counter can be pre-loaded via the Counter Access Register with any value from 0 to 255 (the address range of each memory), or it can be set to auto-increment after each memory access. This function, which is enabled by the *Auto-Increment Enable* bit in the CSR, makes it possible for the DAQ to read or write sequential blocks of memory without the need to reset the counter before each access.

It can be seen from figure 6.9 that the *Memory Type* and *Channel Pair Number* bits in the CSR can be used to select memories which do not physically exist on the BCID board. Any attempt to access a non-existent memory is illegal. In order to avoid the situation where the crate controller and the BCID board both try to read or modify the information in a given register at the same instant, it is also illegal to access any of the registers listed in table 6.3, other than the CSR and the Module Identification Register, whilst the BCID board is running. No information is written to the board by an illegal access and any information read will be undefined (i.e., unpredictable). To indicate that an illegal access has taken place the *Protocol Error* bit in the CSR is set high by the board. This bit can be cleared by writing a 1 to the *Clear Protocol Error* bit.

The FPGA-based logic on the BCID board is controlled by the FPGA Control and Status Register (FCR). Table 6.5 lists the bits in this register. As described above, each FPGA-based BCID design has two mode-select inputs: *Peak Finder Override* and *FIR Filter Enable* (see table 6.1). These inputs are tied to the two least significant bits of the FCR, such that setting the *Peak Finder Override* bit in the FCR asserts the *Peak Finder Override* signal to all of the FPGAs on the BCID board, and similarly for the *FIR Filter Enable* bit. The rest of the bits in the FCR are used to control or monitor the initial act of configuring the FPGAs with a given design.

| **Bit** | **Bit name** | **Bit type** | |
|---|---|---|---|
| 0 | *Peak Finder Override* | Input to card | |
| 1 | *FIR Filter Enable* | Input to card | |
| 2 | Unused | – | (Always 0 on read) |
| 3 | Unused | – | (Always 0 on read) |
| 4 | *FPGA Number* (bit 0) | Input to card | |
| 5 | *FPGA Number* (bit 1) | Input to card | |
| 6 | Unused | – | (Always 0 on read) |
| 7 | Unused | – | (Always 0 on read) |
| 8 | *Clear Program* | Input to card | |
| 9 | *Readback* | Input to card | |
| 10 | *RAM Clock Trigger* | Input to card | (Always 0 on read) |
| 11 | *RAM Clock Status* | Output from card | |
| 12 | *Not Initialised* | Output from card | |
| 13 | *All Loaded* | Output from card | |
| 14 | *Readback In Progress* | Output from card | |
| 15 | *Loaded* | Output from card | |

Table 6.5. FPGA Control and Status Register.

### 6.5.3. FPGA Configuration

The configuration bit-stream which determines the internal configuration of an FPGA is transferred from the DAQ system to the FPGA in a two-stage process. In the first stage the bit-stream is downloaded from the DAQ on to the BCID board and stored in the FPGA Configuration RAM (a 9k × 16-bit block of memory which is external to all of the FPGAs on the board). This memory is accessed via the FPGA RAM Access Register, which provides a window on to the memory location specified by the FPGA Counter Register. After each RAM access the address contained within the FPGA Counter Register automatically increments, to speed up the data transfer process. The FPGA Counter Register can be pre-loaded with any desired value.

Once the entire bit-stream has been downloaded to the FPGA Configuration RAM it can be used to configure any of the FPGAs on the BCID board. As mentioned in section 6.3, there are a number of different modes in which a bit-stream can be transferred to an XC4008 FPGA. ‘Synchronous Peripheral’ mode is used on the BCID board. In this mode the configuration bit-stream is supplied to the FPGA as a sequence of 8-bit parallel words, accompanied by a clock signal. The FPGA uses the clock to serialise the 8-bit words internally, so the clock must run at eight times the baud rate of the data [31].

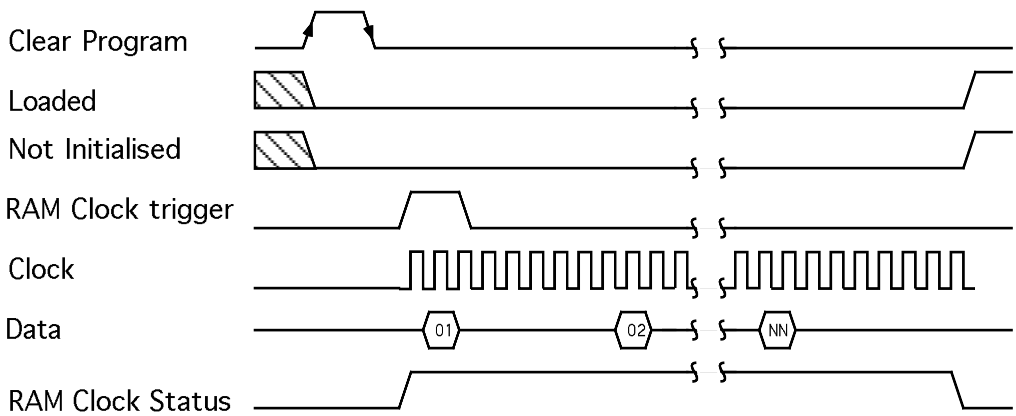


Figure 6.10. Signals used to control and monitor the FPGA configuration process.

On the BCID board the configuration process is initiated by the FCR, but controlled by a state machine. Figure 6.10 illustrates the process. The FPGA to be configured is selected via the *FPGA number* contained within the FCR (see table 6.5). *The Clear Program, Readback, Not Initialised, Readback In Progress*, and *Loaded* bits in the FCR all refer only to this FPGA, and not to the board as a whole. Configuration is initiated by setting the *Clear Program* bit in the FCR high. The effect of this is to pull the $\overline{prog}$ input to the selected FPGA low, putting the device in a loop where it continuously flushes its configuration memory, erasing any configuration previously held. In this state the FPGA also pulls down the *Loaded* bit in the FCR. When the

DAQ subsequently lowers the *Clear Program* bit, the FPGA is placed in a state where it awaits the arrival of the first rising edge of the configuration clock. This clock is provided by the state machine mentioned above. For reasons internal to the Xilinx FPGA there must be a 5 μs delay between the lowering of the *Clear Program* bit and the arrival of the first configuration clock edge. Rather than have this delay provided by the BCID board it is provided by the DAQ, which starts the state machine by pulsing the *RAM Clock Trigger* bit.

The state machine is driven by a 1 MHz crystal oscillator. In addition to providing the clock to the FPGA it provides all of the signals necessary to read the configuration bit-stream out of the FPGA Configuration RAM and time-multiplex the resultant 16-bit words on to the 8-bit bus used to transfer the bit-stream to the FPGA. Whilst running it monitors the *Not Initialised, Loaded,* and *Clear Program* bits in the FCR, halting if any of these bits are pulled high. If the *Loaded* bit is pulled high (by the FPGA), this indicates that the configuration process has been completed successfully. If the FPGA pulls the *Not Initialised* bit high, this indicates that an error has been found in the configuration bit-stream and the loading process has been aborted. The only way for the DAQ to stop the state machine is to pull the *Clear Program* bit high, which aborts the loading process and erases from the FPGA any configuration data already transferred. The status of the state machine is displayed in the FCR by the *RAM Clock Status* bit, which is pulled high whilst the state machine in running. Once all three FPGAs on a board have been configured, the *All Loaded* bit in the FCR is set high. It is pulled low again if any of the FPGAs have their configuration cleared.

In order to verify the configuration of an FPGA, the configuration bit-stream may be read back non-destructively by placing the FPGA in 'Readback' mode. This is done by setting the appropriate *FPGA number* in the FCR and pulling the *Readback* bit high. Once in Readback mode an FPGA, if supplied with an external clock, outputs its configuration bit-stream serially on a dedicated pin. On the BCID board, this serial bit-stream is assembled into 16-bit words and transferred to the FPGA Configuration RAM – a process which is controlled by the same state machine as the configuration process. The *Readback In Process* bit in the FCR is driven high by the FPGA whilst readback is occurring and the state machine halts when this bit is driven low, indicating that the complete bit-stream has been transferred to the FPGA Configuration RAM. If the *Readback* bit is cleared by the DAQ whilst readback is occurring, the process is aborted and the state machine is halted. If readback is allowed to run to completion the *Readback* bit remains high until cleared by the DAQ. The FPGA and the state machine are only sensitive to transitions in the state of this bit, so to initiate another readback the DAQ must clear and reset this bit.

### 6.5.4. Board Layout

The design entry and layout of the BCID board were performed using the Cadnetix CAD package. Figure 6.11 shows the board layout. The front panel is on the left of the diagram; it houses three 64-way ribbon-cable connectors through which the board receives twelve channels of 8-bit differential ECL data. There are a total of 207 IC components on the board. Most noticeable are the outlines of the three large FPGA packages: U25, U77 and U131. Most of the other components consist of RAM blocks to record data, tri-state latches to control bus access, and ECL/TTL converters. Data are output from the board to the CFM via the Teradyne back-plane connectors (PL1–4, PL9–12 and PL28–31). Access for slow-control and readout is provided by the back-plane connectors PL61 and PL62. Most of the slow control logic is implemented in the 40-pin PALs (EP910s) shown at the top of the board (U1, U2, U7, U8).

Figure 6.12 shows a photograph of a BCID board.

Figure 6.11. Layout diagram for the BCID board.

Figure 6.12. A photograph of a BCID board.

# 7. Beam Tests

## 7.1. Introduction

The BCID modules described in the previous chapter were designed to be tested in a CERN test beam, in conjunction with the already-existing cluster finding module (CFM). This complete demonstrator system was run in a test beam for two periods during 1995: June and September [32]. This chapter describes the work done during those periods.

## 7.2. Test Beam Set Up

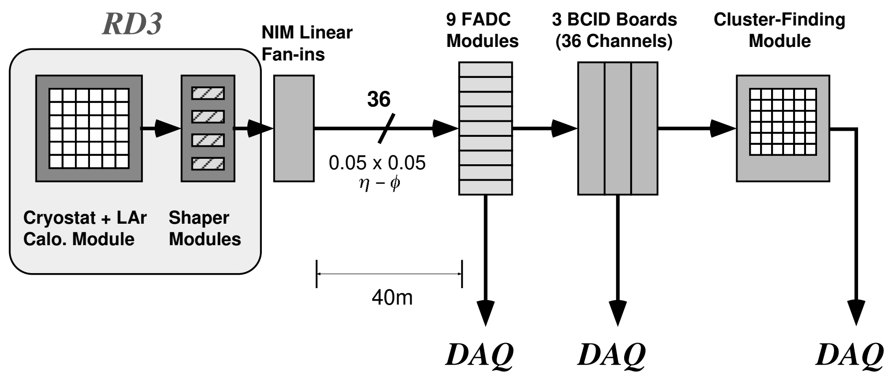


Figure 7.1. A block diagram of the BCID and trigger demonstrator system.

The demonstrator system used in the 1995 beam tests is shown schematically in figure 7.1. As in previous tests, calorimeter signals were taken from the RD3 prototype LAr calorimeter module. On this occasion the calorimeter was divided into seven regions, each of which employed different front-end electronics [33]. Only three of these regions produced signals which could be used by the trigger demonstrator (the remaining four regions featured a geometry which could not be mapped on to the $\eta$–$\phi$ geometry of the level-1 trigger). The three regions which were of use to the level-1 demonstrator are shown in figure 7.2. They are labelled according to the type and position of the pre-amps used to amplify the signals from the calorimeter read-out electrodes in that region:

- in the *Si* region, linear pre-amps were used, situated inside the LAr cryostat;
- in the *Si-0T* region, linear pre-amps were used, situated outside the cryostat;
- in the *Si-bi-gain* region, non-linear, bi-gain pre-amps were used, situated inside the LAr cryostat.

In order to implement $6 \times 6$ triggers cells in the area of calorimeter available, it was necessary to use a reduced trigger-cell granularity of $\Delta\eta \times \Delta\phi \approx 0.05 \times 0.05$ (c.f.

$\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ used in previous beam tests and proposed for the final system). With this granularity, each of the regions shown in figure 7.2 was $6 \times 4$ trigger cells in area. The connections between the calorimeter and the trigger demonstrator were moved frequently to make the region of the calorimeter targeted by the beam visible to the trigger demonstrator.

The FADCs shown in figure 7.1 were purpose-built for the level-1 demonstrator system by the University of Heidelberg [34]. There were nine FADCs in total, each processing four channels of data. The analogue signals were digitised using a free-running 40 MHz clock and a TDC was used to record the phase of the clock relative to the arrival time of the beam particle. The calibration of the FADCs was 0.9 GeV/count and the dynamic range was 8 bits. After digitisation the calorimeter signals were transmitted, via twisted-pair ribbon cables, to the BCID system.

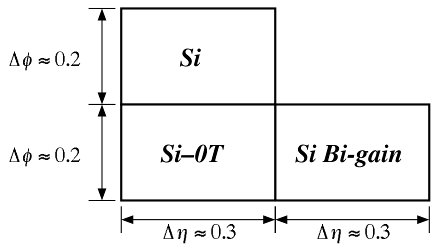


Figure 7.2. The area of calorimeter available to the trigger demonstrator.

The BCID system consisted of three boards of the type described in chapter 6. Thirty six channels of data were processed, with each BCID board receiving the data from three FADC boards. Once BCID had been performed the signals were transferred via a custom-built crate backplane to the CFM [26]. This module implemented a version of the e/γ cluster-finding algorithm which predated that described chapter 3.

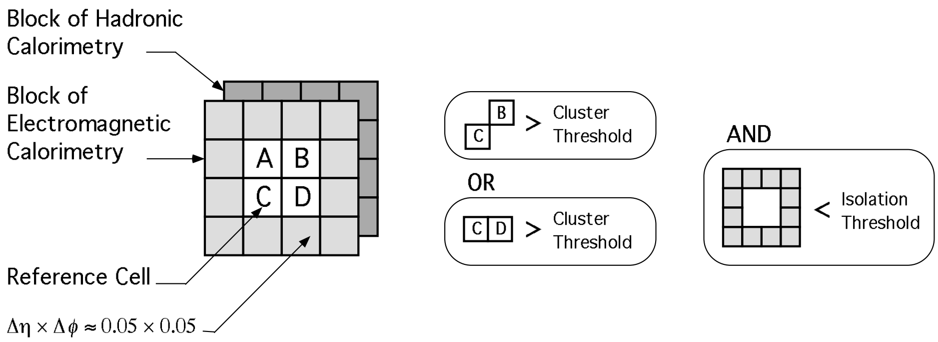


Figure 7.3. The e/γ trigger algorithm used by the CFM.

It was intended that the CFM should process two cluster windows: one in the $\eta$ direction and one in the $\phi$ direction. Instead, a slight design error in the CFM resulted in the cluster windows shown in figure 7.3 being implemented. The CFM processed a total area of $6 \times 6$ trigger cells, forming nine e/γ trigger windows. Each trigger window was formed by one of nine ASICs mounted upon the module. Two sets of isolation and cluster thresholds were provided by each ASIC (set A and set B) and the e/γ hit result was calculated for both sets of thresholds. In addition, each ASIC performed the initial stages of the missing-$E_T$ and jet triggers by calculating the energy sum over a $4 \times 4$ area of trigger cells. Together with the e/γ hit results, these energy sums were sent to FIFOs on the CFM. From there they were read out by the DAQ.

The data processed by the level-1 demonstrator were captured in memories positioned throughout the system. In addition to the FIFOs mentioned above, memories were positioned at the output of the FADCs, the input and output of the BCID boards, and the input of the CFM. It can be noted that some of these memories recorded identical data: the FADC memories and the BCID input memories for example. This duplication of recorded data was a valuable feature when commissioning the system, as it allowed faults in data transmission due to bad connections or inappropriately phased clocks to be tracked down easily.

The data recording process itself was almost identical to that described in section 5.2. It was initiated by a signal from a scintillator counter placed in the beamline, upstream of the calorimeter. Upon receipt of this signal the clock module produced a programmable number clock ticks which drove the FADCs, the BCID boards and the CFM. After the programmed number of ticks had elapsed the clock signal was suspended, freezing the system, and the various memories described above were read out by the VME-based DAQ. An exabyte storage system was used to record the data.

## 7.3. A sample Event

Figure 7.4 shows a sample of the data recorded for one event in the June test beam period. Eighty timeslices were recorded for this event, of which only sixty time slices are shown.

The uppermost plot in figure 7.4 shows the data for one trigger cell, as captured by the BCID input memory. The next plot down shows the data for the same trigger cell recorded in the BCID output memory. Between these two stages pedestal subtraction and bunch-crossing identification are performed, so the only non-zero timeslice in the BCID output data should be that corresponding to the peak of the calorimeter pulse. Two conditions should therefore be met by the peak in the output data. Firstly, it should appear eight clock cycles after the peak in the input data (corresponding to the latency of the BCID algorithm), and secondly, it should be of a height equal to that of the input pulse minus the constant pedestal value of thirty FADC counts. Inspection shows that both of these conditions are met. The BCID algorithm used here is a stand-alone peak finder with the criterion $R_1 < R_2 \geq R_3$.

Upon reaching the CFM, data from individual trigger cells are combined to form the $4 \times 4$–cell windows used by the e/γ trigger. The last two plots of figure 7.4 show two of the results obtained from a window built using data from the trigger cell shown in the previous plots. The results shown are the energy sum over the trigger window and the e/γ hit results for one of the two sets of thresholds implemented by the CFM. Note that these two processes are implemented with different latencies and their outputs appear in different timeslices.

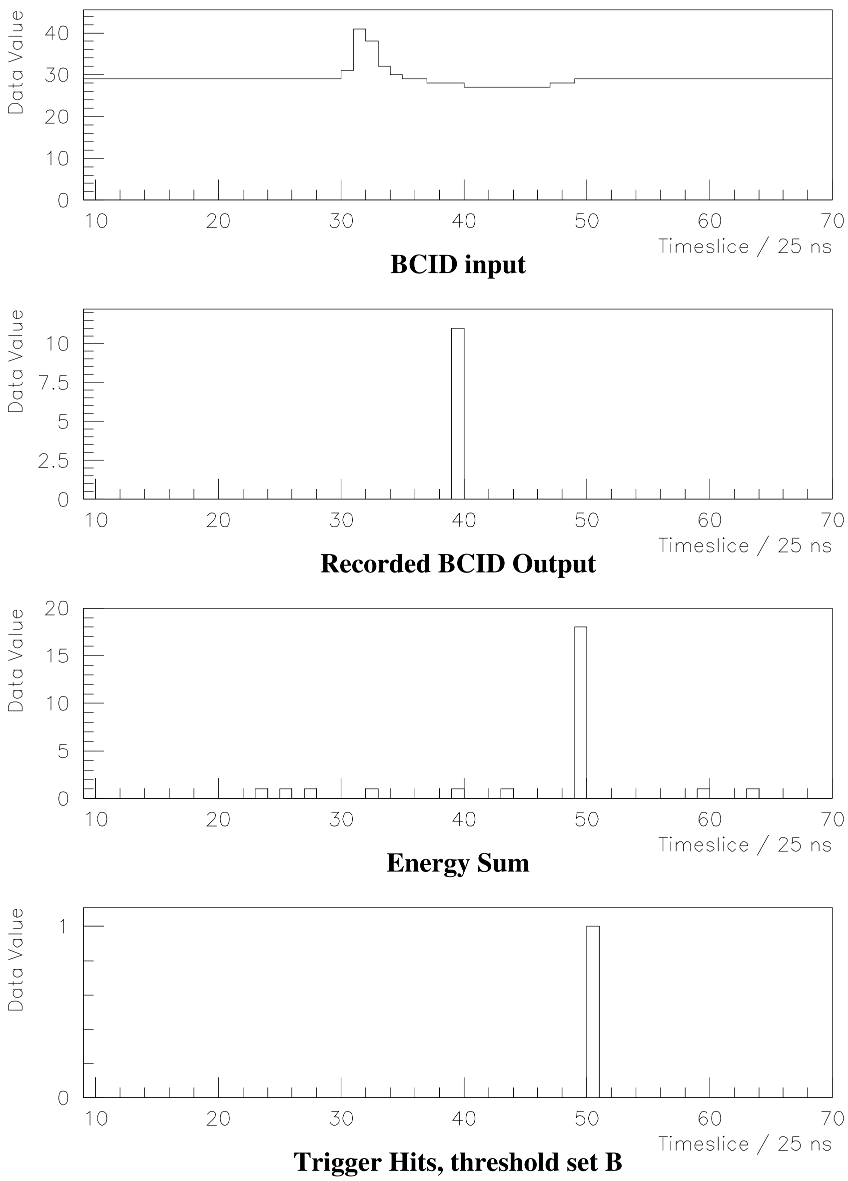


Figure 7.4. A sample of the data recorded by the demonstrator system for a single event.

As the height of the energy sum shown in the third plot of figure 7.4 is greater than that of the pulse shown in the previous plots, it obviously contains contributions from trigger cells other than that shown. It is encouraging that the energy sum is not spread over more than one timeslice, as this would indicate that BCID had been performed incorrectly for at least one of the trigger cells in the window. From the e/γ hit result shown in the last plot in figure 7.4, it can be deduced that the position of the e/γ shower relative to the trigger window under consideration was such that the trigger conditions were met. Namely, the energy in one of the cluster windows was greater than 10 FADC counts and the energy in the isolation region was less than 5 FADC counts. Again it is encouraging that the trigger condition was met for one timeslice only: that corresponding to the peak of the calorimeter pulse. Not shown in figure 7.4 are the contents of the FADC output memories and the CFM input memories. These memories would have recorded the same data as the BCID input and BCID output memories respectively, but over a slightly different time frame (due to the time taken to transfer data between boards).

## 7.4. Analysis of Test Beam Data

### 7.4.1. Hardware Tests

The programme to analyse the data recorded at the test beam had two goals: to establish that the hardware functioned to specification, and to evaluate the performance of the BCID algorithms. Initially, priority was given to the former of these tasks. This was for two reasons. Firstly, it was necessary to establish the integrity of the recorded data before any algorithmic evaluation could be conducted. Secondly, it was desirable to use the data recorded in the June test beam period to identify and correct any faults in the demonstrator system before the September run. Whilst it was known from on-line monitoring that the results produced by the demonstrator system appeared sensible, no detailed event-by-event analysis was possible on-line. This task was left to the off-line hardware checks.

The off-line hardware checks involved comparing the data recorded at all stages of the demonstrator system with that predicted by a software model of the system. This software model used the data recorded in the FADC memories as input. The vast majority of the events examined showed the hardware behaving as expected. For around 0.5% of events however, some inconsistency was found between the recorded data and the predicted data. There were found to be two causes of error in the hardware.

One cause was the timing set-up of the write strobe to the memories on the BCID board. Occasionally the write strobe would arrive at a memory out of time (by ~1 ns) with either the address or data signals and would either begin before these signals had properly been set up, or persist longer than they were valid. In such instances the data written to a particular memory location would be corrupted. Setting up the correct timing on a BCID board was a non-trivial task and no significant reduction in the rate at which these errors occurred could be made before the September run.

All errors in the recorded data not accounted for by the above problem were found to be caused by the memories on a BCID board occasionally being out of synchronisation with the memories on other modules. When this loss of synchronisation occurred all of the data recorded by one BCID board for one event appeared displaced by a single timeslice. This phenomenon was due to a slight flaw in the logic on the BCID board, as explained below.

At the end of each event the BCID boards received a *STOP* signal from the clock module. This signal prevented the memories on a module from recording any further data and inhibited the distribution of the clock signal around the module. Unfortunately, in the original design of the BCID module, whilst the memories stopped recording data instantly, the clock signal continued to be propagated until a falling clock-edge was reached (thus avoiding glitches in the clock). This meant that on occasion a BCID module would process one more clock cycle than it recorded, leading to a miss-alignment between the recorded data and the Address Counter on that board. When the data for that event were read out by the DAQ, the data from that BCID board would thus appear to be out of step with the rest of the system. Once this fault had been identified it was corrected with a minor wiring change which ensured that the data recording process was disabled synchronously with the clock signal being halted (on a falling clock edge).

The conclusions drawn from the hardware tests were positive. The BCID algorithms were shown to be behaving as they had been designed to do and the vast majority of the recorded data were shown to be reliable. Analysis of the BCID algorithms could therefore proceed with the knowledge that any failures in the BCID process discovered would most likely be due to design flaws in the algorithms themselves rather than hardware faults.

### 7.4.2. Evaluation of the BCID algorithms

Table 7.1 summarises the 1995 test-beam data which have been used to evaluate the performance of the BCID algorithms. Only data recorded in the *Si* and *Si-0T* regions of the calorimeter have been used for this purpose, as preliminary analysis of the data recorded in the *bi-gain* region shows that the pulses produced by the front-end electronics in that region of the calorimeter are unsuitable for testing the BCID algorithms. This is for two reasons.

| Algorithm Label | FIR coefficients | Peak-Finding Criteria | Number of events recorded |
|---|---|---|---|
| *bcid1* | 0, 0, 1, 0, 0 | $R_1 < R_2 \geq R_3$ | 232,000 |
| *bcid2* | 0, 0, 1, 0, 0 | $R_1 < R_2 > R_3$ | 95,000 |
| *bcid3* | 0, 0, 1, 0, 0 | $R_1 \leq R_2 > R_3$ | 15,000 |
| *bcid4* | 0, 0, 1, 0, 0 | $R_1 \leq R_2 \geq R_3$ | 21,000 |
| *bcid5* | -1, 0, 2, 0, -1 | $R_1 < R_2 \geq R_3$ | 294,000 |
| *bcid6* | -1, 0, 2, 0, -1 | $R_1 \leq R_2 > R_3$ | 15,000 |
| *bcid7* | -1, 0, 2, 0, -1 | $R_1 \leq R_2 \geq R_3$ | 8000 |
| *bcid8* | -1, 0, 4, 0, -1 | $R_1 < R_2 > R_3$ | 17,000 |
| *bcid9* | -1, 0, 4, 0, -1 | $R_1 < R_2 \geq R_3$ | 232,000 |
| *bcid10* | -1, 0, 4, 0, -1 | $R_1 \leq R_2 > R_3$ | 13,000 |

Table 7.1. Data used to evaluate the BCID algorithms.

Firstly, the pulses from the *bi-gain* region are much longer than those for which the BCID algorithms were designed. This is important for those algorithms which employ FIR filters, as the coefficients in those filters have been optimised for pulses of a specific length and the algorithms are known to fail for pulses of a longer duration than this.

Secondly, and most importantly, the pulses produced in the *bi-gain* region of the calorimeter have an amplitude-dependent shape. This is catastrophic for BCID, as such pulses cannot be processed adequately using linear-response algorithms. In the most optimistic (hypothetical) case, BCID can be performed on such pulses using a bank of filters, with the appropriate filter being selected according to the pulse height. This idea is complicated though by the fact that the trigger cells used by the level-1 trigger are formed from a summation of calorimeter cells. Each pulse seen by the trigger is therefore a composite of a number of smaller pulses, and its shape will depend not upon its total height, but upon the height of each individual pulse from which it is composed. This means that there will be no unambiguous relationship between the height of a pulse, as seen by the trigger, and its shape. To attempt BCID for such pulses is beyond consideration. In the final ATLAS system only linear calorimeter pulses will be acceptable to the level-1 trigger.

The data shown in table 7.1 were recorded in both the June and September periods. Data were collected for all of the BCID algorithms, with the number of events collected per algorithm ranging from 8000 to over 200,000. The algorithms which were targeted with the most events were those which simulation studies indicated would give the best performance. As had been done in previous beam tests, a TDC was used in the June and September runs to record the time between the arrival of the incident beam particle and the first rising edge of the demonstrator clock. In order to analyse properly the performance of the BCID demonstrator, the TDC value has been used to make a cut on the data considered.

Only events which feature a calorimeter pulse being sampled within ± 2.5 ns of its peak have been used as test data. This reduces the number of events in the test data to one fifth of that shown in table 7.1. The number of test pulses which can be extracted from each event however, is greater than one. Each shower generally occupies four trigger cells (due to the reduced trigger-cell size employed in the 1995 test-beam period). As the noise in those four cells is uncorrelated, each of the four resultant pulses can be considered separately when analysing the data. Furthermore, as the energy in the shower is not (in general) distributed evenly between all four trigger-cells, the four pulses are generally of differing amplitudes.

The performance of each BCID algorithm has been evaluated by examining the contents of the input and output memories on the BCID modules. In this process the height of a pulse is measured from the input memories, and the output memories are then used to examine how BCID has proceeded for that pulse. The results have been sorted into four categories:

- *Correct* results are those where the peak of a pulse has been gated through the BCID hardware to appear in the output data in the correct timeslice and at the correct height, whilst all other timeslices in the immediate vicinity of the peak (± 2 timeslices) have been suppressed to zero.
- *Incorrect* results are those where a sample other than the peak of the pulse has been gated through the hardware. It is important to note that some of the BCID algorithms can gate through more than one sample per pulse. If a single pulse is found to have produced two non-zero timeslices in the output data, one corresponding to the peak of the pulse and one to some other sample, then that result is counted twice: once in the *correct* category and once in the *incorrect* category. The *incorrect* category is designed to count incorrect identifications produced by inefficiencies in the BCID algorithm, as opposed to hardware faults such as missing bits and dead output channels.
- *Missing bit* events are those where the peak of the pulse is present in the output data in the correct timeslice, but at the wrong height.
- *Dead* events are those where no data is output within 5 samples of the expected output peak.

These last two categories are intended to highlight any cases where the BCID output data may be incorrect due to a hardware fault rather than an algorithmic inefficiency. They are intended as guidelines only and some incorrect events are open to miss-interpretation. For example, the presence of no output data for a single input pulse could have a number of causes, including an algorithmic inefficiency. A large number of results of this type for a given channel though would definitely indicate some sort of hardware problem. In practice, less than 0.1% of the incorrect BCID results fall into either the *missing bit* or *dead* categories.

Figure 7.5 shows some of the results obtained from the BCID algorithm tests. All of the algorithms displayed in this plot use a transparent FIR filter, so only the peak-finding criteria vary between algorithms. The plot shows the percentage of events classified as *correct* for each algorithm as a function of pulse height. Note that as in the earlier software simulation of the BCID algorithms (chapter 5), the analysis of the results has concentrated on pulses of small amplitude.

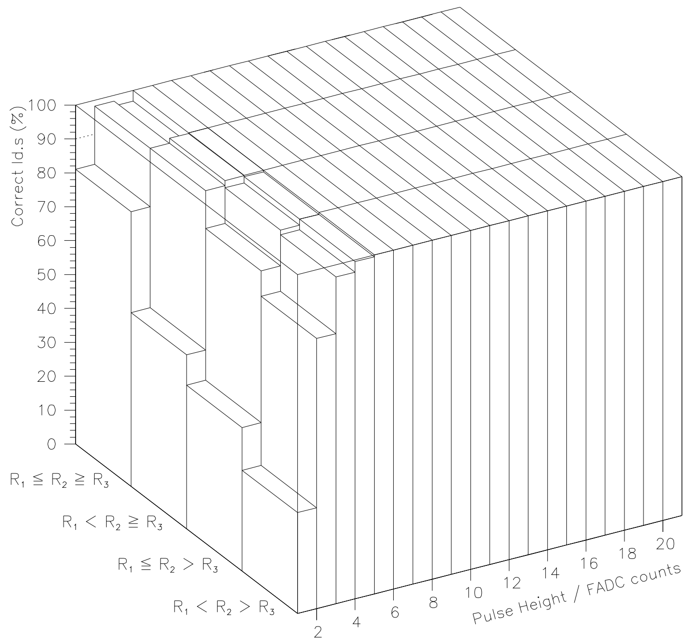


Figure 7.5. Percentage of correct identifications produced by the stand-alone peak finders. Results are shown for the peak-finding criteria indicated.

The general form of the results in figure 7.5 is as expected: the efficiency of all of the algorithms improves with the amplitude of the pulse, and full efficiency is reached for pulses of a few FADC counts height (five in the worst case). It can also be seen that the efficiency varies between different peak-finding criteria as expected. The peak-finding criterion $R_1 < R_2 > R_3$ has the lowest rate of correct identifications for pulses of small amplitude, as it is the most sensitive to noise. The peak-finding criteria $R_1 < R_2 \geq R_3$ and $R_1 \leq R_2 \geq R_3$ have produced the best results, as they are less sensitive to noise on the falling edge of the pulse (as discussed in chapter 5). The apparent drop in the efficiency of the $R_1 \leq R_2 \geq R_3$ peak finder for pulses of 3 FADC counts height is due to poor statistics for this particular algorithm: one pulse out of a total of 63 was incorrectly identified at this height. By comparison, the algorithms with the peak-finding criteria $R_1 < R_2 > R_3$ and $R_1 \leq R_2 > R_3$ were tested with ~200 events per pulse height, and the algorithm with the peak-finding criterion $R_1 < R_2 \geq R_3$ was tested with ~2000 events per pulse height.

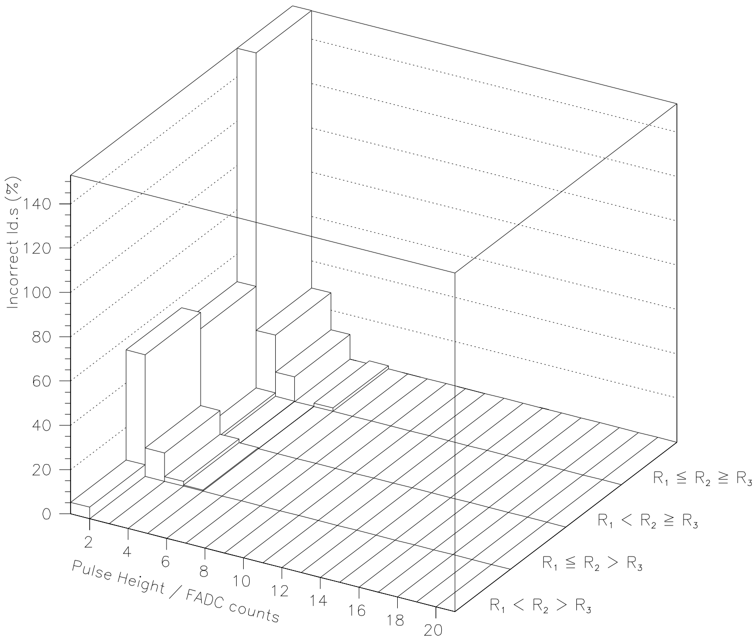


Figure 7.6. The incorrect identifications produced by the stand-alone peak finders, expressed as a percentage of the total number of test pulses of the height indicated.

Not all of the information needed to judge the relative performance of the algorithms shown in figure 7.5 is contained within that plot. In addition to examining the rate of correct identifications produced by each algorithm, the number of incorrect identifications must be studied. These are shown in figure 7.6.

Most of the algorithms displayed in figures 7.5 and 7.6 can produce only one identification per pulse, so for these algorithms the percentage of incorrect identifications and the percentage of correct identifications total to 100%. The exception to this is the algorithm which uses the peak-finding criterion $R_1 \leq R_2 \geq R_3$. It can be seen from figure 7.6 that for pulses of height one FADC count, this algorithm has a rate of producing incorrect identifications of greater than 140%. This means that for most pulses of this height the algorithm produced at least one incorrect identification, and for some pulses of this height it produced two incorrect identifications. These incorrect identifications will have been produced in addition to any correct identifications that may have been made for the same pulse.

Looking back to figure 7.5 it can be seen that the percentage of correct identifications produced by the peak-finding criteria $R_1 < R_2 \geq R_3$ and $R_1 \leq R_2 \geq R_3$ are approximately equal (discounting the results for peaks of just 1 FADC count in height). From figure 7.6 however it can be seen that the more relaxed of the two peak-finders, $R_1 \leq R_2 \geq R_3$, has a much higher rate of producing off-peak identifications for low-energy pulses. As mentioned in chapter 5, it is better to accept a loss in BCID efficiency of a few percent for small pulses than to generate large numbers of spurious identifications at these energy levels. It can therefore be concluded that of the peak-finding criteria shown in figures 7.5 and 7.6, $R_1 < R_2 \geq R_3$ is the most suitable criterion for use in the level-1 calorimeter trigger.

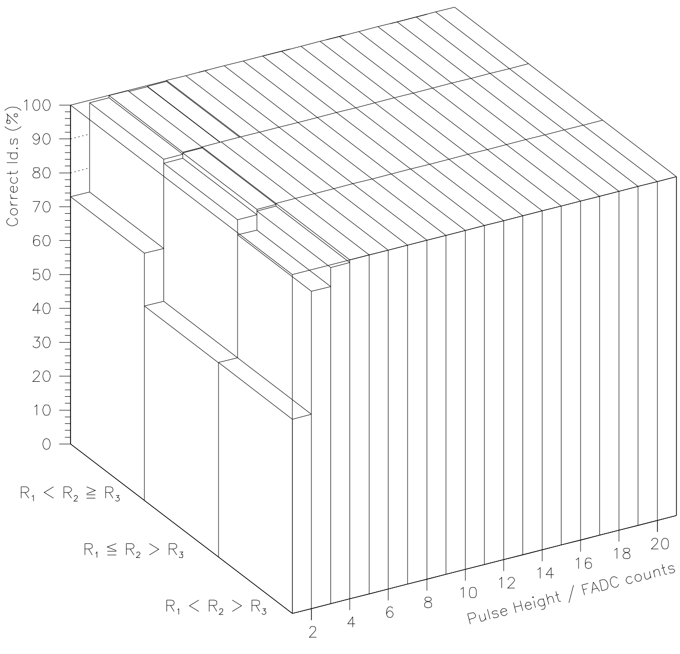


Figure 7.7. Percentage of correct identifications produced by all algorithms using filter coefficients -1, 0, 4, 0, -1. Peak-finding criteria are as indicated.

Figure 7.7 shows more of the results obtained from the beam tests. The general format is the same as figure 7.5: the percentage of correct identifications is shown as a function of pulse height for a variety of algorithms. This time all of the algorithms shown use the FIR filter coefficients -1, 0, 4, 0, -1. As before, all algorithms achieve full efficiency for pulses above a few FADC counts in height (four counts in the worst case), and the BCID efficiency for pulses below this height improves when more suitable peak-finding criteria are used. The improvement in efficiency obtained when using the peak-finding criterion $R_1 < R_2 \geq R_3$, as opposed to $R_1 < R_2 > R_3$ though is less noticeable here than when a transparent filter is used. This is because using a filter with the coefficients given leads to a reduction in the number of pulses which have flat tops (i.e., no single peak sample). This greatly improves the performance of the peak-finding criterion $R_1 < R_2 > R_3$, but has a much smaller effect on the efficiency of the peak-finding criterion $R_1 < R_2 \geq R_3$.

The results obtained for the algorithms which use the filter coefficients -1, 0, 2, 0, -1 are almost identical to those obtained for the algorithms which used the coefficients -1, 0, 4, 0, -1, and for this reason they are not shown here. In each case, irrespective of the filter coefficients used, the peak-finding criterion which produces the best efficiency is $R_1 < R_2 \geq R_3$. This was the expected result and it was for this reason that these algorithms were used the most extensively during the beam tests. Figure 7.8 shows the results for the three algorithms which use this peak-finding criterion. It can be seen that the use of a non-transparent filter produces only a small gain in BCID efficiency, mainly due to the fact that the efficiency when no filter is used is already very impressive: greater than 99% efficiency for all pulses of 2 FADC counts height or greater.

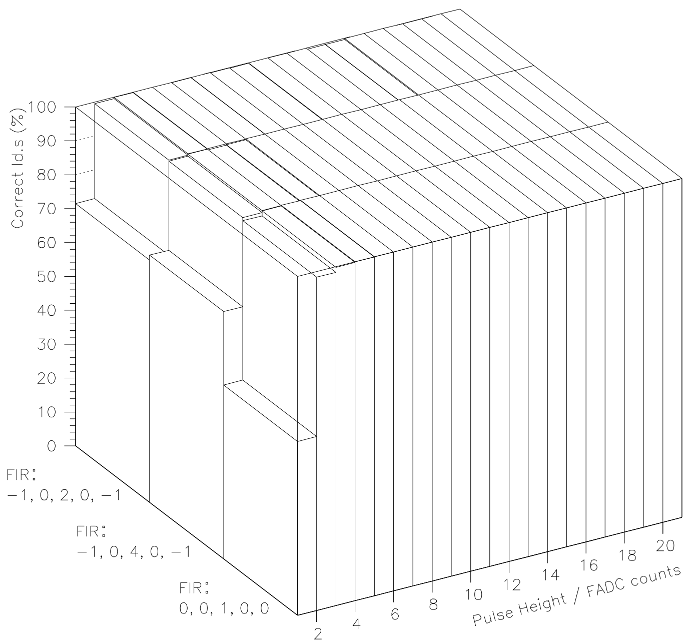


Figure 7.8. percentage of correct events for all algorithms using the peak-finding criterion $R_1 < R_2 \geq R_3$. FIR filter coefficients as indicated.

As all of the algorithms shown in figure 7.8 can produce only one bunch-crossing identification per pulse, the rate at which off-peak identifications are produced by each algorithm is equal to the deficit in the rate at which correct identifications are produced. It can therefore be seen that the performance of all of the algorithms shown in figure 7.8 is comparable. For the data observed in this test beam then, the use of an FIR filter (which adds complexity and latency to the BCID algorithm) is not justified by a significant increase in efficiency.

Figure 7.9 shows the efficiency of one BCID algorithm acting on data recorded in 1993, and 1995. The results for the 1993 data have been taken directly from chapter 5. The difference in efficiency between the two data sets is due to the difference in the noise level on the recorded data. The data recorded in 1993 has an R.M.S. noise level of ~0.8 counts, whilst that recorded in 1995 has an R.M.S. noise level of ~0.2 counts. This reduced noise level can be attributed to two factors: the superior quality of the FADCs used for the 1995 beam tests and the reduced trigger-cell granularity used for those tests. Reducing the number of calorimeter cells which contribute towards a trigger cell inevitably reduces the level of (predominantly electronic) noise present in that cell.

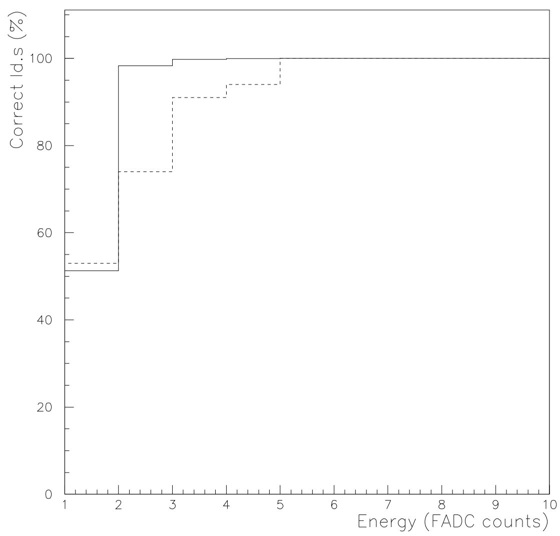


Figure 7.9. Results for the stand-alone peak finder ($R_1 < R_2 \geq R_3$) with data recorded in 1993 (dotted line) and as performed in real-time in 1995 (solid line). Both sets of data were recorded using FADCs of sensitivity 0.9 GeV/count.

The ATLAS technical proposal quotes an expected noise R.M.S. of 400 MeV for the electromagnetic calorimeter. This figure is for a $3 \times 3$ area of calorimeter cells; it includes electronic noise and pile-up noise calculated at the LHC design luminosity. To a first approximation, these contributions scale as the square root of the area of calorimeter considered [3]. The R.M.S. noise level for a trigger cell composed of $4 \times 4$ calorimeter cells will therefore be in the order of 500 MeV, or ~0.5 FADC counts. This is somewhere between the level of noise observed on the 1993 data (over a $4 \times 4$ area calorimeter cells) and the level of noise observed on the 1995 data (over a $2 \times 2$ area calorimeter cells). It can therefore be predicted that if the BCID algorithms were tested with a noise level (and granularity) matching that expected at the LHC, the efficiency observed would be somewhere between that observed for the 1993 data and that observed for the 1995 data.

### 7.4.3. CFM Output

As the BCID demonstrator system was used to perform BCID on a $6 \times 6$ area of trigger cells, the effect of BCID on the performance of the trigger algorithms used by the CFM demonstrator can be investigated.

The CFM demonstrator outputs three results for each of the nine trigger windows which it processes. These are:

- an e/γ hit decision for threshold set A;
- an e/γ hit decision for threshold set B;
- the energy sum over the area of the trigger window (for the proposed use of the jet and missing-$E_T$ triggers in the final system).

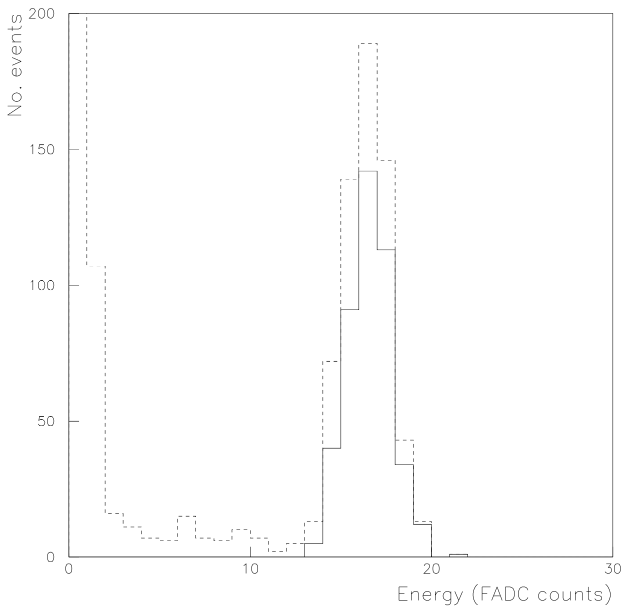


Figure 7.10. The energy measured by the CFM over a window of $4 \times 4$ trigger cells for a sample of 981 events. Distributions are shown for all events (dotted line) and for those events which produced a positive e/γ hit result (solid line).

Figure 7.10 collates the CFM demonstrator output for a run of 981 events recorded at CERN with a 15 GeV electron beam. BCID has been performed on all of the events shown using algorithm *bcid1* (see table 7.1). The horizontal axis of the plot denotes the energy of the incident beam particle as measured by the CFM (i.e., the energy sum described above), whilst the vertical axis indicates the number of particles which were measured as having this energy.

Two sets of data are shown in figure 7.10. The first set of data (shown as a dotted line) contains the energy measurements for all 981 events recorded. Looking at these data, two types of incident beam particle can be distinguished: electrons and pions. The electrons, which deposit the majority of their energy (~97% [35]) in the RD3 calorimeter, are responsible for the Gaussian distribution of events centred at 17 FADC counts (15 GeV). The pions, which deposit only a small fraction of their energy in the (electromagnetic) calorimeter, are responsible for the much broader spread of events at energies below this.

The second data set shown in figure 7.10 (as a solid line) contains the energy measurements for only those events which produced an e/γ trigger hit from the CFM, for either set of thresholds. The thresholds used were:

- threshold set A: cluster threshold = 14 FADC counts,<br>isolation threshold = 4095 FADC counts;
- threshold set B: cluster threshold = 10 FADC counts,<br>isolation threshold = 5 FADC counts.

It can be seen from figure 7.10 that whilst no pion events produced an e/γ hit (as expected), there were also some electron events which failed to produce an e/γ hit, even though they produced an energy deposit greater than the lowest cluster threshold (10 FADC counts). This was not due to the hits being vetoed by the isolation requirement, but was a result of the decreased granularity of trigger cells used for these beam tests ($\Delta\eta \times \Delta\phi \approx 0.05 \times 0.05$ compared to $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ proposed for the final system). This reduced granularity led to a poor lateral containment of showers within the $2 \times 1$ cluster windows implemented by the e/γ trigger. In general, each shower was spread over a $2 \times 2$ area of trigger cells, so only around half of the energy of a shower was captured by any one cluster window. Figure 7.11, for example, shows an event which has produced a total energy deposit of 16 FADC counts – greater than the lowest cluster threshold of 10 FADC counts. The largest amount of energy which has been deposited in a single cluster window however, is only 9 FADC counts. The event has therefore failed to produce an e/γ hit. All of the electron events which failed to produce a hit did so for a similar reason. No hits were vetoed by the isolation requirement.

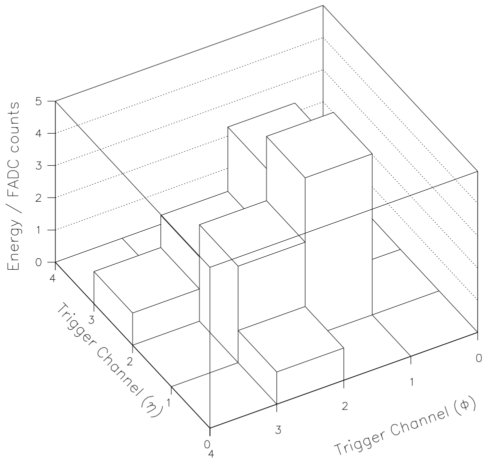


Figure 7.11. A map of the energy deposits produced in the calorimeter for a single shower.

Figure 7.12 shows the same 981 events as figure 7.10. Here however, each event has been plotted as a function of the largest energy deposit it produced in any cluster window – i.e., the cluster energy as measured by the e/γ trigger. This measurement has been reconstructed off-line using the recorded CFM input data. From figure 7.12, it can be seen that there is a clear cut between those events which produce an e/γ hit and those which do not, and that this cut occurs at 10 FADC counts as expected. From this it can be concluded that the CFM functioned correctly during the beam tests.

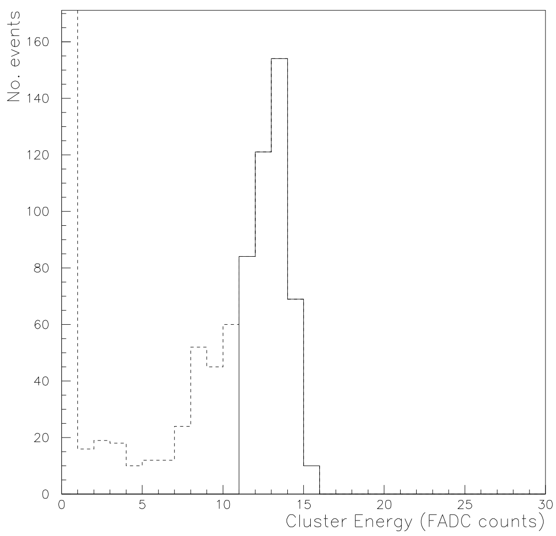


Figure 7.12. The largest energy deposit in a CFM cluster window for a sample of 981 events. Distributions are shown for all events (dotted line) and for those events which produced a positive e/γ hit result (solid line).

A more important conclusion can be drawn from figure 7.12 if the cluster energy is reconstructed not from the CFM input data, but from the FADC or BCID input data. If this is done the resultant plot is identical to that shown. As the FADC and BCID input data were recorded before BCID was performed, it can be concluded from this that for these 981 events at least, BCID did not have a detrimental effect on the performance of the e/γ trigger.

In order to understand fully the effect of BCID on the e/γ trigger a more thorough investigation is required. Specifically, it is important to look at the effect of BCID on the sharpness of the trigger threshold. As described in section 3.2.1, this is one of two figures of merit for the e/γ trigger (the other being the trigger rate produced when the thresholds are set at a level efficient for the desired physics).

To examine the effect of BCID on the sharpness of the e/γ cluster threshold, a software model of the e/γ trigger has been created. The algorithm used by this model is not the exact e/γ trigger algorithm but is a modification of it designed to cope with the reduced trigger cell granularity. A cluster window of $3 \times 3$ cells is used. This window is large enough to contain 98% of the energy deposited in the calorimeter by all showers. The modified algorithm compares the energy deposited in this window to a single cluster threshold. No isolation requirement is imposed.

The software e/γ trigger has been tested on data recorded at the test beam. By using data recorded either before or after BCID was performed, it is possible to examine the effect of BCID on the trigger. Figure 7.13[1] shows how the efficiency of the trigger varies as a function of the cluster threshold, both with and without BCID having been performed. The BCID algorithm used in this instance is *bcid10* (see table 7.1) and the data are from a 200 GeV electron beam (325 events). It can be seen that the effect of BCID is to shift the threshold curve down by 1–2 FADC counts, without degrading the sharpness of the curve. This is because BCID reduces the value of the energy in the cluster window by an average of ~ 1 FADC count (see below). If such a situation were replicated at the LHC, the e/γ trigger thresholds would have to be lowered by one or two counts to maintain efficiency for the events of interest. As BCID reduces *all* energy deposits by ~1 GeV, the background trigger rate would not be increased by this action (compared to no BCID and a higher threshold). This situation is therefore acceptable. A more detrimental situation would be if BCID lead to a decrease in the sharpness of the threshold. The necessary lowering of the thresholds in this case would lead to an increased background trigger rate.

[1] Figure 7.13 is courtesy of Adam Connors.

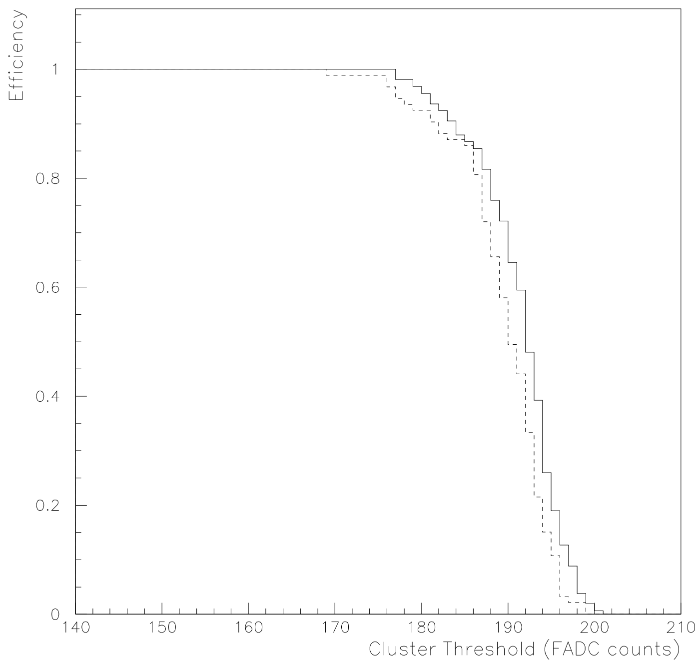


Figure 7.13. The efficiency of the e/γ trigger as a function of the cluster threshold when BCID is not performed (solid line) and performed (dotted line). The e/γ trigger algorithm used to generate this plot is described in the text.

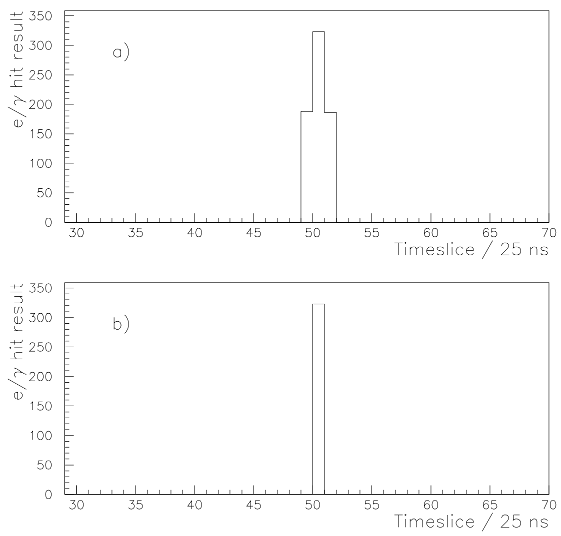


Figure 7.14. Hit results from the e/γ trigger as a function of time. Plots are shown for a) data recorded before BCID and b) data recorded after BCID.

The need for BCID is illustrated in figure 7.14. This plot shows the e/γ hit results for those events used to generate figure 7.13 as a function of time, both before and after BCID. The e/γ-trigger algorithm used to generate this plot is that implemented by the CFM demonstrator. This has been performed in real-time by the CFM itself for the post-BCID data, and by a software model for the pre-BCID data. The cluster threshold in both cases is set at 40 FADC counts. It can be seen that when BCID is not performed, over half of the events produce e/γ hit flags for more than one consecutive timeslice. After BCID, all events produce a hit result in one timeslice only: that corresponding to the peak.

In addition to studying the effect of BCID on the e/γ trigger, it is also important to study its effect on the resolution of the CFM energy measurements. This resolution can be extracted from figure 7.10 (for example), by fitting a Gaussian curve to the distribution of electron events (all electron events, not just those which produce e/γ hit). If this is done, the figure obtained for the resolution is $\sigma/E = (7.4 \pm 0.3)\%$.

The resolution of the RD3 calorimeter is quoted as

$$\frac{\sigma}{E} = \frac{10\%}{\sqrt{E}} \oplus \frac{28\%}{E} \oplus 0.3\% \oplus \frac{0.1\%}{E}$$

where $E$ is in GeV. It should be noted that this resolution is for the full calorimeter, which is 25 radiation lengths in depth[14]. As the level-1 demonstrator system was connected to only the first two sampling depths of the calorimeter (18 radiation lengths in total) the actual resolution observed will be worse than that predicted by this equation, due to the leakage of shower energy. This matter will be dealt with below.

Each term in the above equation represents a contribution to the calorimeter resolution from a different source. The first term is the sampling term, which accounts for fluctuations in the number of electrons produced and detected for a shower (hence the variation with $E^{-1/2}$). The second term is the contribution to the resolution from both pile-up noise and electronic noise. The third term is due to instrumental effects such as differences in calibration between calorimeter cells, and the fourth term is due to quantisation noise on the digital signals.

For particles of energy 15 GeV (17 FADC counts), the above equation predicts a resolution of ~4% for the calorimeter. The resolution extracted from figure 7.10, which is the resolution of the CFM energy measurement, is almost twice this figure. This is a common result: by fitting Gaussian distributions to electron events for a variety of beam energies, the resolution of the CFM energy measurement can be measured over a range of energies. In each case the measured CFM resolution is 2-3 times worse than that predicted by the above equation. It is obviously necessary to identify all of the causes of this poor resolution and to establish that BCID is not responsible.

Figure 7.15 clearly shows that BCID is not a cause of the poor resolution of the CFM energy measurement. This plot shows the resolutions obtained from actual CFM energy measurements (post-BCID) and from energy measurements reconstructed off-line using FADC data (pre-BCID data). It can be seen that in each case the resolutions of these two measurements are comparable, and that they are a factor of two worse than the predicted calorimeter resolution. Further investigation shows there to be two causes of discrepancy between the resolution as predicted above and the resolution observed. These two causes are the leakage of shower energy (as mentioned above), and a poor calibration between trigger cells.

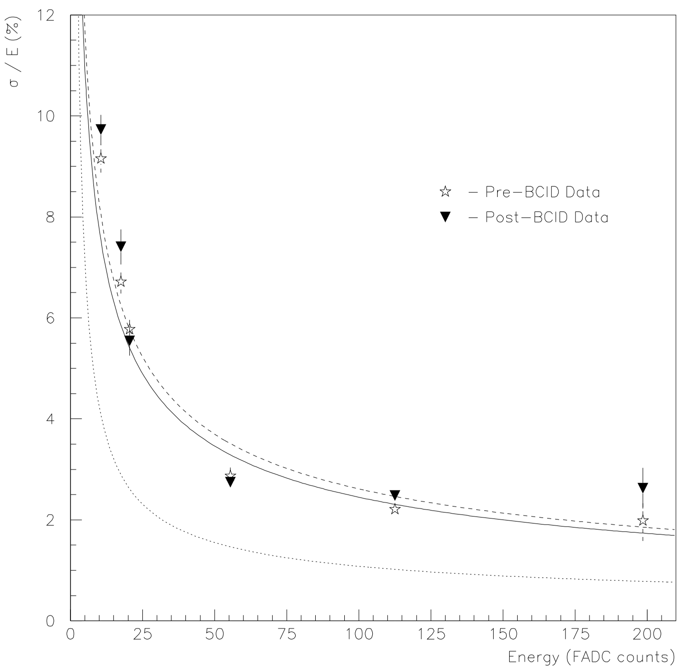


Figure 7.15. The resolution of the CFM energy measurement before and after BCID. Also shown is the predicted calorimeter resolution (lower dotted line).

In the test beam the BCID demonstrator system took signals from only the first two sampling depths of the calorimeter. This covered a total depth of 18 radiation lengths. For showers of 200 GeV (for example) only 93% of the energy of the shower will be deposited within this length [35]. The effect of this on the resolution of the calorimeter can be judged by Monte Carlo simulation. According to such methods an extra term of 1.49 % should be added (in quadrature) to the predicted resolution of the calorimeter for electrons of 200 GeV. The predicted resolution for such electrons becomes

$$\frac{\sigma}{E} = \frac{10\%}{\sqrt{E}} \oplus \frac{28\%}{E} \oplus 0.3\% \oplus \frac{0.1\%}{E} \oplus 1.49\% = 2.0\%.$$

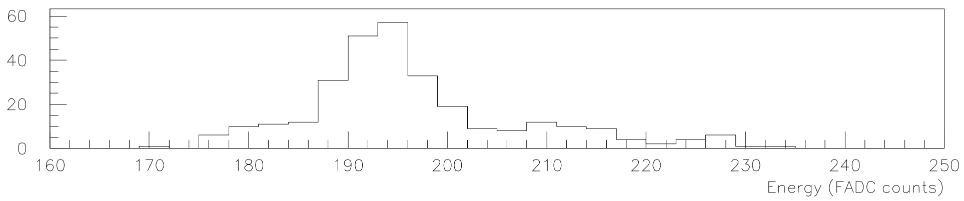


**a) All Events**

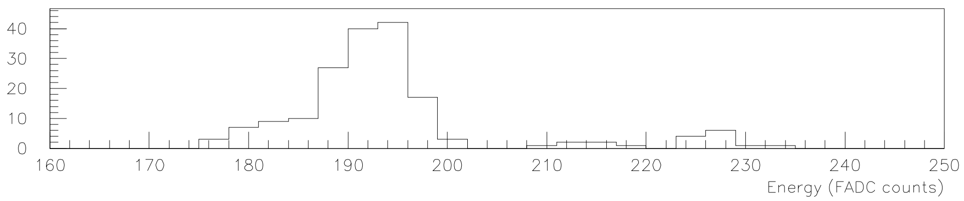


**b) Events which deposit 75% of their energy within a single cell.**

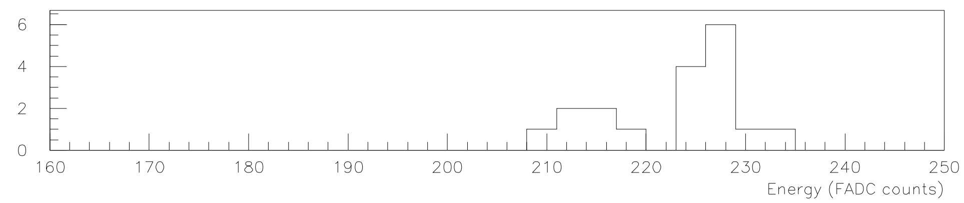


**c) Events which deposit 75% of their energy within cell (3,4).**

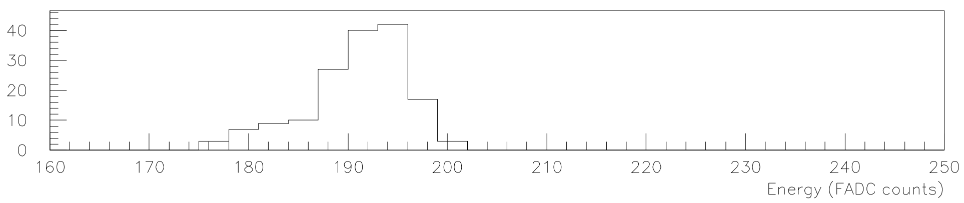


**d) Events which deposit 75% of their energy within cell (3,3).**

Figure 7.16. The Energy distribution for a sample of 297 events with the cuts indicated.

The poor calibration observed between trigger cells at the test beam may have occurred either in the FADCs, or before them in the RD3 front-end electronics. The data available cannot be used to distinguish between these two cases. Figure 7.16[2] shows the effect of poor calibration on the resolution of the CFM energy measurement. Plot (a) shows the distribution of the CFM energy measurements for 500 events of 200 GeV electrons. The energy measurements have been made over a $4 \times 4$ area of cells. It can be seen that a wide spread of measurements is obtained and that the distribution of these measurements is not Gaussian. Plot (b) shows a subset of the data from plot (a). A cut has been made such that only showers which deposited > 75% of

[2] Figure 7.16 is courtesy of Adam Connors.

their energy within a single trigger cell are shown. This has resolved the distribution of events into two peaks, centred at ~195 FADC counts and ~225 FADC counts. Closer examination shows that these two peaks correspond to the showers being centred on two different trigger cells (see plots (c) and (d)). It can therefore be concluded that the energy measurements from these two trigger cells must have a different calibration.

By fitting a Gaussian curve to the events shown in figure 7.16(d), a value can be obtained for the resolution of the CFM energy measurement without any degradation due to the difference in calibration between trigger cells. The Gaussian curve should be fitted only to those events of energy greater than 186 FADC counts, not to the long low-energy tail of the distribution which is due to leakage. The resolution obtained is $\sigma/E = (2 \pm 0.2)\%$. Thus, when leakage and calibration are taken into account, there is a good correlation between the predicted calorimeter resolution of 2.0% and the measured CFM resolution of $(2.0 \pm 0.2)\%$. This means that the quality of the energy measurement has not been degraded by any process between the calorimeter and the CFM. Specifically, it has not been degraded by BCID. This is an important conclusion.

Despite the above results, BCID obviously does have some effect on the energy measured by the CFM. This is shown in figures 7.13 and 7.15, wherein the plots constructed using data recorded before and after BCID are not identical. Figure 7.17 shows the mean discrepancy between energy measurements made before and after BCID. The energy measurements made after BCID are taken directly from the recorded CFM output, whilst the energy measurements made before BCID have been reconstructed off-line using the recorded FADC data. All measurements are over a $4 \times 4$ area of trigger cells ($\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$). It can be seen that on average BCID reduces the measured energy by ~1 FADC count, irrespective of the original energy of the event. The cause of this reduction can be traced back to inefficient BCID for pulses of 1-2 FADC counts in height. Most of the observed events generated at least one pulse of this height in a trigger cell on the periphery of the shower. As BCID is only ~50% efficient for pulses of such height, some energy is inevitably lost. This effect is greatly exaggerated by the decreased trigger-cell granularity, which results in showers being spread over a larger number of trigger cells than would otherwise be expected. Thus showers of 200 GeV (for example) still manage to produce a significant number of pulses of ~1 FADC counts height in individual trigger channels.

If trigger cells of granularity $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ are used, the above effect becomes insignificant for all but the smallest of pulses. Figure 7.18 shows this. It has been generated using the same 1000 events as figure 7.17. Here however, a trigger cell granularity of $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ has been created by summing the recorded FADC data over the appropriate trigger cells. BCID has then been applied using a software model and the energy has been measured over the same area as before ($\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$). It can be seen that for all energy deposits above 10 FADC counts in magnitude, the mean energy loss due to BCID is immeasurably small.

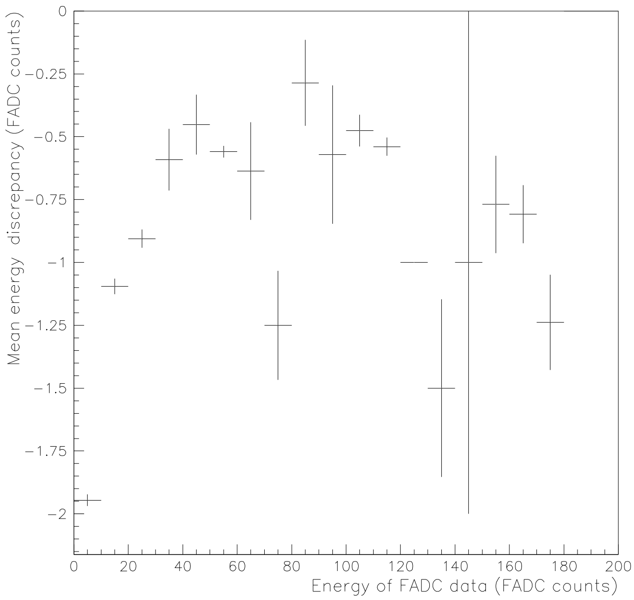


Figure 7.17. The mean discrepancy between energy measurements made before and after BCID, as a function of the energy measured before BCID. All measurements are made over an area of $\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$, using a trigger cell granularity of $\Delta\eta \times \Delta\phi \approx 0.05 \times 0.05$

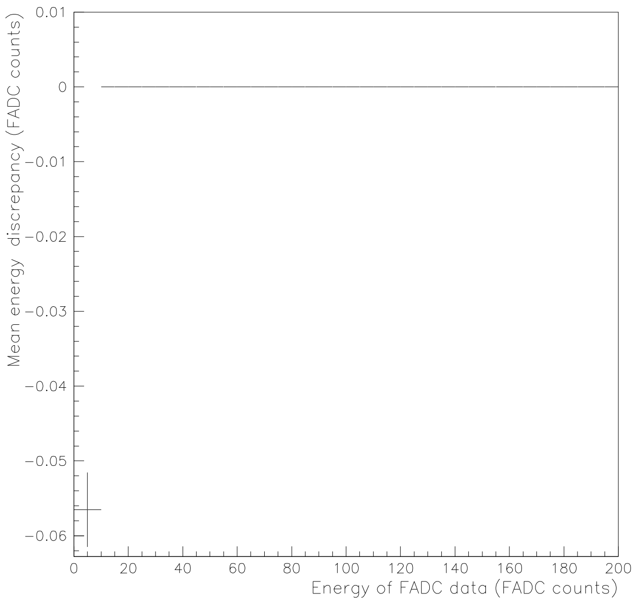


Figure 7.18. The mean discrepancy between energy measurements made before and after BCID, as a function of the energy measured before BCID. All measurements are made over an area of $\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$, using a trigger cell granularity of $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$

In order to judge properly the effect of BCID on the level-1 calorimeter trigger, it is necessary to simulate BCID and the level-1 trigger algorithms together under LHC conditions (namely, at the LHC design luminosity). This has been done by parameterising the BCID performance in terms of the percentage of correct identifications made as a function of pulse height, and fitting these parameters into the existing GEANT [36] simulation code for the ATLAS level-1 trigger[3]. Two sets of BCID efficiencies have been created: *BCID1* and *BCID2* . The parameters for these sets of efficiencies are listed in table 7.2. *BCID1* is a pessimistic estimation of the BCID efficiency – it is comparable to the level of efficiency observed for BCID with the 1993 data (see chapter 5). *BCID2* is a more optimistic estimation of BCID efficiency, comparable to the results obtained from the 1995 beam tests.

| BCID label | Percentage of correct identifications given for pulses of the height indicated | | | | |
|---|---|---|---|---|---|
| | 1 GeV | 2 GeV | 3 GeV | 4 GeV | ≥ 5 GeV |
| ***BCID1*** | 50 % | 70 % | 85 % | 95 % | 100 % |
| ***BCID2*** | 64 % | 94 % | 98% | 100 % | 100 % |

Table 7.2. The BCID efficiency as parameterised for the Monte Carlo GEANT simulations.

The first step towards judging the effect of BCID on triggering is to determine the tightest trigger thresholds (i.e., the highest cluster threshold and the lowest isolation thresholds) which can be used to achieve the required trigger efficiency. These thresholds will produce the lowest background trigger rate for that given efficiency. Table 7.3 lists the tightest thresholds which can be used with the e/γ trigger to produce a 95% efficiency for electrons of 40 GeV. The results in this table have been generated using a sample of 500 electron events. The e/γ trigger algorithm used here is the full algorithm described in section 3.2.1. The results are shown for the two BCID efficiencies (*BCID1* and *BCID2*), and for the case where BCID is performed with 100% efficiency.

In the first column of table 7.3, the cluster threshold shown is the highest which can be applied whilst maintaining an efficiency of ≥ 95% for 40 GeV electrons. This is when no isolation requirement is imposed. The second column of the table shows the tightest thresholds which can be applied when the full e/γ trigger algorithm is used, and the third column shows the lowest isolation thresholds which can be used when the cluster threshold is set at zero. This latter case is equivalent to setting the efficiency of the

[3] All of the Monte Carlo results presented here are courtesy of Alan Watson.

cluster threshold to 100%. This is appropriate for modelling the efficiency of the trigger for electrons which are more than a few GeV above the cluster threshold.

Table 7.3 shows that for the e/$\gamma$ trigger, inefficient BCID has the largest effect on the isolation thresholds. This is to be expected, as BCID is only inefficient for small pulses. These small pulses make a more significant contribution towards the isolation sums than they do towards the cluster sum, hence it is the isolation thresholds which are most affected by inefficient BCID.

| **BCID Efficiency** | **Thresholds required to produce 95 % efficiency for e/$\gamma$ trigger.** | | |
|---|---|---|---|
| | **Cluster Threshold only** | **Full e/$\gamma$ trigger algorithm** | **Isolation thresholds only** |
| **Perfect BCID** | > 32 | em cluster > 32<br>+ em isolation < 6<br>+ had. isolation < 4<br><br>( em cluster > 32<br>+ em isolation < 7<br>+ had. isolation < 3 ) | em isolation < 5<br>+ had. isolation < 3<br><br>( em isolation < 6<br>+ had. isolation < 2 ) |
| ***BCID1*** | > 32 | em cluster > 32<br>+ em isolation < 4<br>+ had. isolation < 3<br><br>( em cluster > 31<br>+ em isolation < 4<br>+ had. isolation < 2 ) | em isolation < 3<br>+ had. isolation < 2 |
| ***BCID2*** | > 32 | em cluster > 32<br>+ em_isolation < 5<br>+ had. isolation < 3 | em isolation < 4<br>+ had. isolation < 2 |

Table 7.3. Trigger thresholds required to produce 95 % trigger efficiency for electrons of 40 GeV with the BCID efficiencies indicated. Where two sets of thresholds produce identical results, the alternative set of thresholds is shown in parentheses.

Having obtained the above results, it is important to determine what trigger rates result from the threshold levels listed in table 7.3. As mentioned previously, the trigger rate of the e/$\gamma$ trigger is dominated by the QCD jet background. A Monte Carlo study has therefore been made using a sample of 1289 QCD jet events, with a pile-up of 18 events per bunch-crossing and an energy spectrum matching that expected at the LHC (see figure 3.3 for example). Table 7.4 shows the number of triggers generated by the e/$\gamma$ trigger with each of the BCID efficiencies. The isolation thresholds set in each case are those from the last column of table 7.3. The cluster threshold is set at 15 GeV. This low cluster threshold has been used in order to increase the number of triggers generated, and thus to increase the statistical significance of any difference in trigger rates observed between the three BCID efficiencies. The resultant trigger rates are therefore not realistic figures for the final ATLAS e/γ trigger. They do however, show the relative effect of different BCID efficiencies on the trigger rates. In the most pessimistic case, the performance of BCID with a finite efficiency increases the trigger rate by ~15%. This compares to an error in the predicted ATLAS level-1 trigger rates of ~10% [19]. At this stage then, the increase in the trigger rate due to BCID is insignificant when compared to the general uncertainty in the predicted trigger rates.

| BCID efficiency | Number of triggers | Corresponding trigger rate |
|---|---|---|
| **Perfect BCID** | 51 (52) | (74 ± 10) kHz |
| ***BCID1*** | 59 | (86 ± 11) kHz |
| ***BCID2*** | 57 | (83 ± 11) kHz |

Table 7.4. The trigger rates produced by the isolation thresholds specified in table 7.3 for a Monte Carlo simulation of 1289 QCD jet events at $L = 10^{34}\text{cm}^{-2}\text{s}^{-1}$.

Table 7.5 shows the effect of BCID on the jet trigger. This table has been generated using ~500 events each of 40 GeV and 300 GeV QCD jets. Shown in each case is the jet threshold required to produce a trigger efficiency of 90%. It can be seen that to maintain efficiency when non-perfect BCID is performed it is necessary to lower the jet threshold by a few GeV, irrespective of the energy of the jets being considered. Such a situation would be perfectly acceptable for the final ATLAS level-1 trigger. As there is no significant background to the jet trigger, lowering the jet threshold to compensate for imperfect BCID would not lead to an increase in the jet trigger rate.

| BCID efficiency | Threshold required for 90 % efficiency for 40 GeV QCD jets | Threshold required for 90 % efficiency for 300 GeV QCD jets |
|---|---|---|
| **Perfect BCID** | > 30 GeV | > 272 GeV |
| ***BCID1*** | > 24 GeV | > 264 GeV |
| ***BCID2*** | > 27 GeV | > 268 GeV |

Table 7.5. The effect of BCID (with a finite efficiency) on the jet-trigger thresholds.

The above results show that the level of efficiency achieved by the BCID demonstrator system would be acceptable for the final ATLAS level-1 trigger system.

# 8. Conclusion

The high luminosity and large centre-of-mass energy of the LHC will provide an opportunity to explore many new areas of physics, most notably, the electroweak Higgs sector. The proposed ATLAS detector is a multi-purpose pp experiment, able to exploit the full discovery potential of the LHC. The high luminosity and short bunch-crossing period of the LHC will present a considerable challenge to the ATLAS triggering system.

Bunch-crossing identification is an essential part of the ATLAS level-1 calorimeter trigger. Investigation shows that in most instances bunch-crossing identification can be performed with good accuracy using relatively simple, digital algorithms. There are three types of pulses for which these algorithms are less accurate:

- small pulses (less than ~ 5 GeV), due to the increased signal-to-noise ratio;
- saturated pulses;
- overlapping pulses (instances of pile-up).

Of these three types of pulses, only pulses of the first type have been considered in this thesis. It has been shown that for small pulses, BCID can be performed to an accuracy acceptable to the level-1 calorimeter trigger using simple, digital algorithms. The feasibility of implementing these algorithms in hardware has been proven by a technical demonstrator.

The accuracy of a BCID algorithm when dealing with overlapping pulses is not considered to be an important figure of merit for the level-1 BCID system. Monte Carlo studies have been performed to determine the extent of pile-up expected at the LHC. Assuming pulses of three bunch-crossings duration are observed by the calorimeter trigger, there is a ~20% chance of any one pulse overlapping with a second pulse of 1 GeV height (averaged over the whole detector). The chances of any pulse overlapping with a second pulse of 2 GeV height are ~2%, and the chances of any pulse overlapping with a pulse of 3 GeV height are ~0.5%. Extrapolating from this, it can be concluded that the probability of two significant pulses overlapping is negligible [37]. Overlapping pulses are therefore not a concern of level-1 BCID.

Whilst the issue of saturated pulses has not been addressed in this thesis, it is believed that the problems associated with such pulses are understood. A second BCID algorithm is required, operating in parallel to the one which performs BCID for non-saturated pulses. This second algorithm is required to recognise saturated pulses and extract timing information from them. A number of simple digital algorithms are under consideration as possible methods of performing BCID for saturated pulses. One candidate algorithm is the constant fraction discriminator.

The architecture of the level-1 BCID system has not yet been finalised; it could be implemented using either analogue or digital technology. The requirements for a digital BCID system however are well understood, and some guiding design principles can be specified:

- If a digital BCID system is to be used, it must be implemented in ASIC technology. Whilst FPGAs provide an ideal solution for a BCID demonstrator system, they are too large and costly to be used in the final system.
- Two parallel BCID algorithms are required, to deal with saturated and non-saturated pulses.
- The BCID logic should be programmable at the level of parameters (FIR coefficients, etc.) to allow it to be optimised for the different pulse shapes which are expected in different areas of the calorimeter.
- The BCID logic should be of minimal latency. Current latency calculations for the level-1 trigger assume a latency of four bunch crossings.

The investigation of both digital and analogue methods of bunch-crossing identification is continuing. The technology choice for BCID (and for the level-1 calorimeter trigger as a whole) will be frozen in 1997. At this point detailed design of the BCID system will begin, with the proposed prototype modules ready in 1999. The LHC is currently scheduled to start running in 2004.

# Appendix A.

## A Note on the Conventions used to Describe FIR Filters in this Thesis.

As noted in chapter 4 (section 4.2.1), the behaviour of an FIR filter is described by the equation

$$y(n) = \sum_{i=0}^{N-1} h(i)x(n-i)$$

where $h(i) = a_i$ if $0 \le i \le N\text{-}1$

and $h(i) = 0$ elsewhere.

This can also be written as a matrix equation:

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} a_1 & 0 & 0 & \cdots \\ a_2 & a_1 & 0 & \cdots \\ a_3 & a_2 & a_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \end{bmatrix}$$

where $y_n$ is the output from the FIR filter in the $n$th timeslice.

Throughout this thesis, a graphic representation of the FIR coefficients ($a_i$) is used. For example, a filter which contains the following coefficients:

$a_1 = 1,$

$a_2 = -X,$

$a_i = 0$ where $i \ge 3$,

can be represented by the pictogram shown in figure A.1.

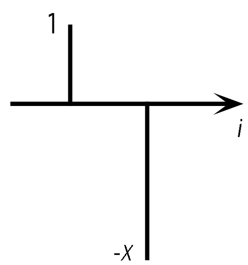


Figure A.1. A graphic representation of the FIR filter coefficients ($a_i$) described in the text.

The operation of the FIR filter shown in figure A.1 can be described by the following matrix equation.

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ \vdots \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & \ldots \\ -X & 1 & 0 & 0 & \ldots \\ 0 & -X & 1 & 0 & \ldots \\ 0 & 0 & -X & 1 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \end{bmatrix}$$

In section 4.3.4 (and elsewhere) it is stated that two FIR filter operations can be combined and performed by a single FIR filter. This is easy to show using matrices. Consider an output, $y_n$, which is the product of two consecutive FIR filters (A and B) acting on the input $x_n$. This is described by the equation

$$Y = B.A.X$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} b_1 & 0 & 0 & \ldots \\ b_2 & b_1 & 0 & \ldots \\ b_3 & b_2 & b_1 & \ldots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} a_1 & 0 & 0 & \ldots \\ a_2 & a_1 & 0 & \ldots \\ a_3 & a_2 & a_1 & \ldots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \end{bmatrix}$$

which can be seen to be equivalent to the equation

$$Y = C.X$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} b_1a_1 & 0 & 0 & \ldots \\ b_2a_1 + b_1a_2 & b_1a_1 & 0 & \ldots \\ b_3a_1 + b_2a_2 + b_1a_3 & b_2a_1 + b_1a_2 & b_1a_1 & \ldots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \end{bmatrix}$$

hence the two FIR filters, A and B, can be replaced by a single FIR filter with the following coefficients:

$$c_1 = b_1a_1,$$

$$c_2 = a_1b_2 + b_1a_2,$$

$$c_3 = a_1b_3 + a_2b_2 + a_3b_1, \quad \text{etc..}$$

It is this process of combining coefficients from consecutive FIR filters which is shown in figure A.2. This figure is a reproduction of figure 4.8 from chapter 4. It shows that the equation

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ \vdots \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 1 & 1 & 0 & 0 & \dots \\ 1 & 1 & 1 & 0 & \dots \\ 1 & 1 & 1 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} X-1 & 0 & 0 & 0 & \dots \\ -1 & X-1 & 0 & 0 & \dots \\ -1 & -1 & X-1 & 0 & \dots \\ -1 & -1 & -1 & X-1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \end{bmatrix}$$

is equivalent to

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ \vdots \end{bmatrix} = \begin{bmatrix} X-1 & 0 & 0 & 0 & \dots \\ X-2 & X-1 & 0 & 0 & \dots \\ X-3 & X-2 & X-1 & 0 & \dots \\ X-4 & X-3 & X-2 & X-1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} . \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \end{bmatrix}$$

which follows from simple matrix mathematics.

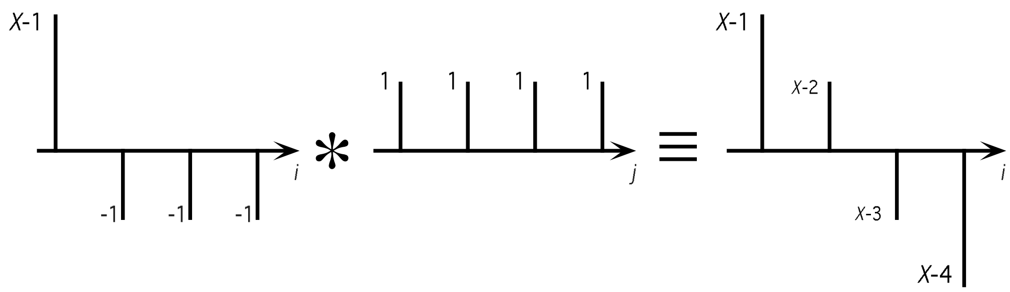


Figure A.2. A graphic illustration of how two consecutive FIR filters can be combined and performed by a single FIR filter. The example shown here is taken from section 4.3.5: The Comparison of the Pulse with its Integral.

# Appendix B.

## An Explanation of the Schematic Diagram for the FPGA-Based BCID Logic.

Each FPGA-based BCID logic design contains four identical blocks of logic, processing four channels of calorimeter data in parallel. The logic required to process one channel of calorimeter data is shown as a block diagram in figure B.1. In this logic, incoming data are stored in a pipeline memory whilst BCID is performed. Data samples are tapped from this pipeline by the FIR filter, which performs a weighted sum of the samples and outputs the result to a second pipeline memory. This memory contains the filtered waveform. The peak finder then examines this filtered waveform and upon finding a peak gates out the corresponding data sample from the original pipeline memory. If no peak is found in the filtered waveform for a given timeslice, then the output from the BCID logic for that timeslice is suppressed to zero.

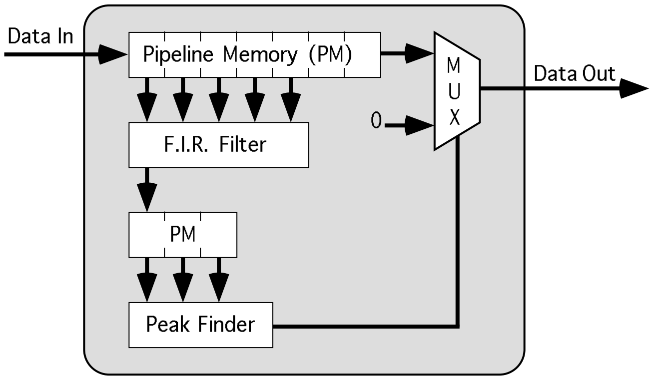


Figure B.1. Block Diagram of the BCID Logic.

Figure B.2 shows a more fundamental schematic diagram of a BCID logic design. As with figure B.1, the logic shown is that required to process one channel of calorimeter data – a complete design would contain four blocks of logic identical to this. The correspondence between the schematic diagram of figure B.2 and the block diagram of figure B.1 is shown in figure B.3.

It can be seen from figure B.3 that the two pipeline memories used in the BCID design are built from chains of registers. Further registers are positioned throughout the design to ensure the synchronous operation of the logic.

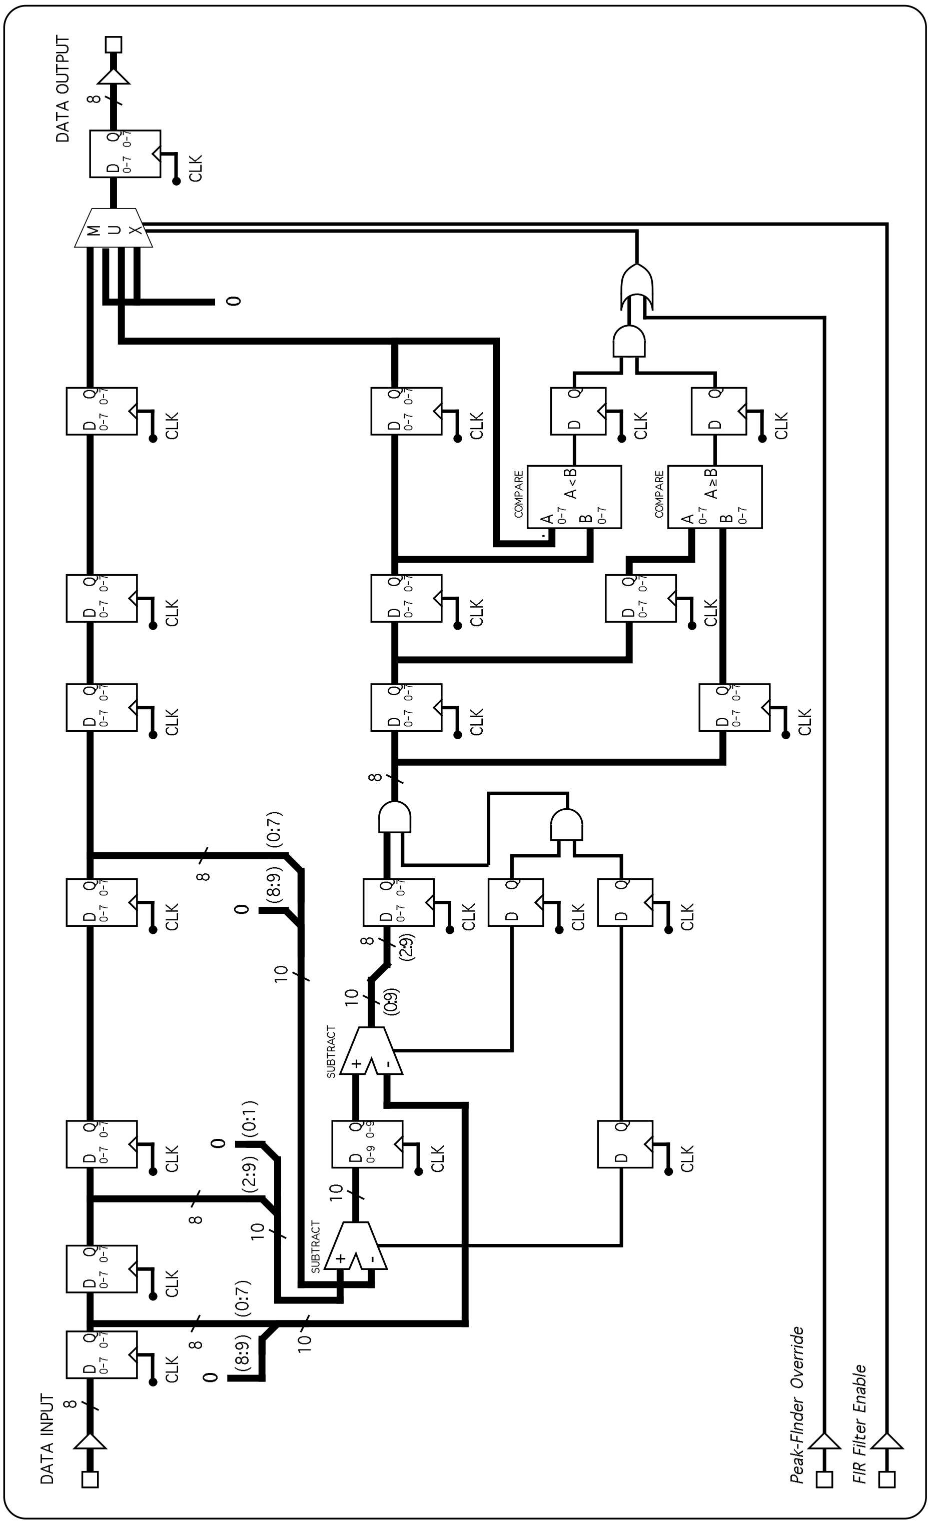


Figure B.2. A schematic diagram of an FPGA-based BCID logic design. The FIR coefficients used in this design are –1, 0, 4, 0, –1 and the peak-finding criterion is $R_1 < R_2 \geq R_3$.

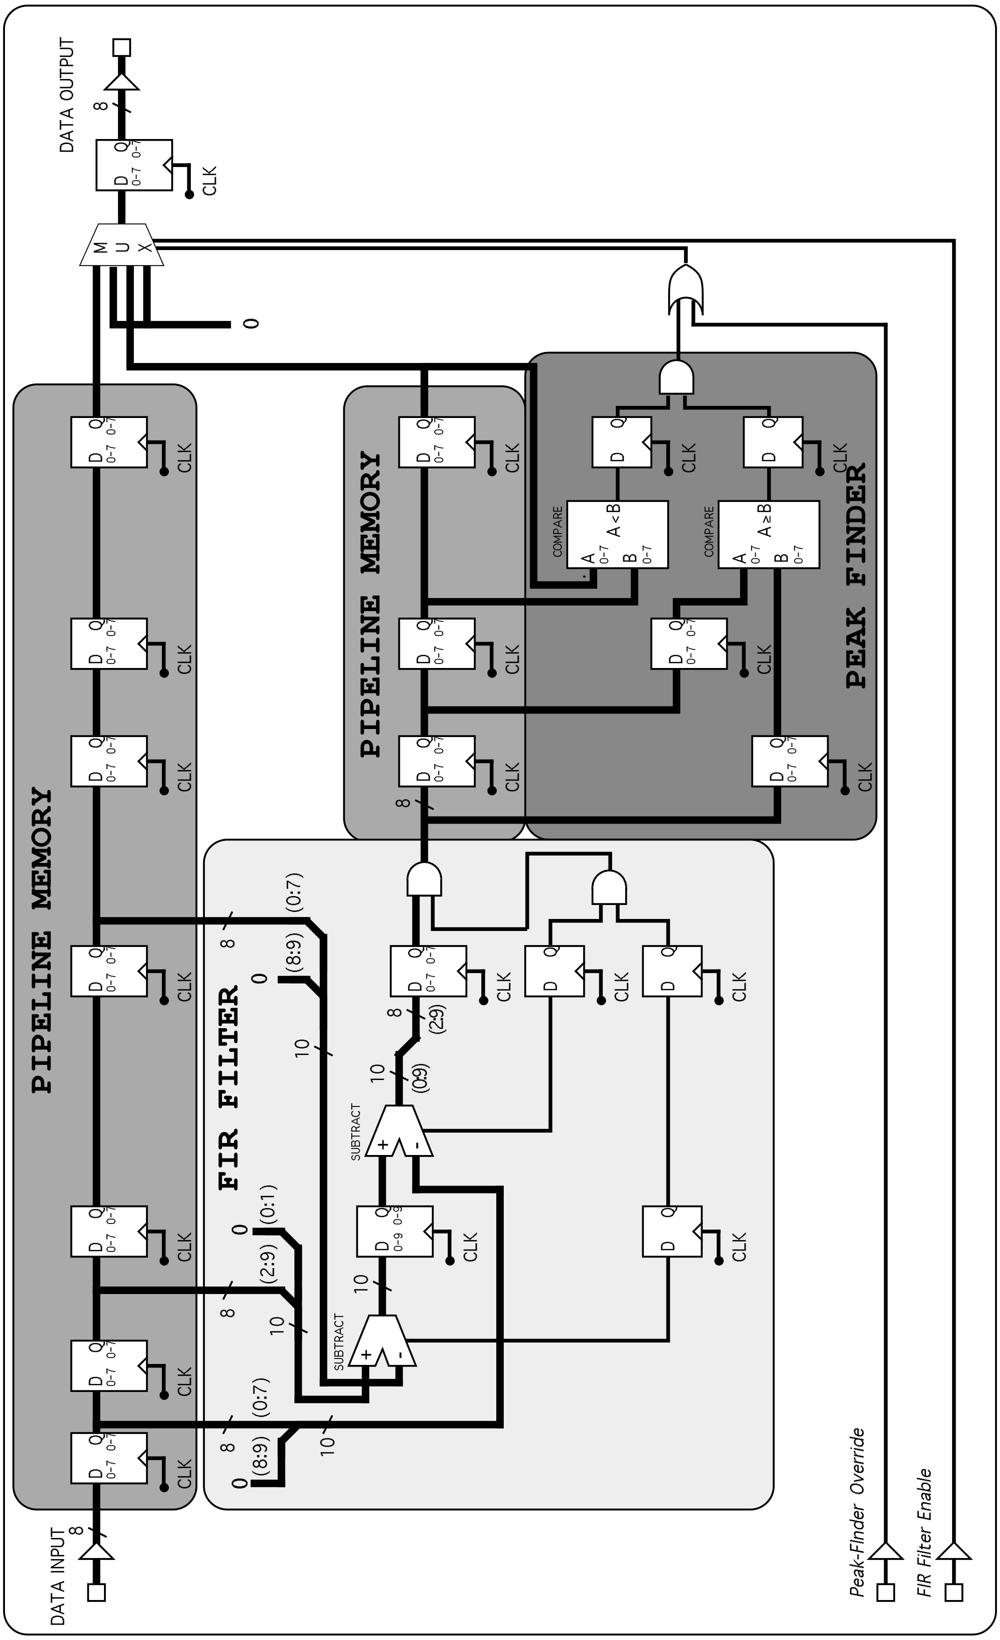


Figure B.3. The correspondence between the BCID logic-block diagram shown in figure B.1 and the more fundamental schematic diagram shown in figure B.2.

The FIR filter indicated in figure B.3 performs four tasks:

1. The data samples extracted from the pipeline memory are weighted. For the design shown this means that the samples are multiplied by a factor of either one or four, according to the position in the pipeline from which they are extracted. Any multiplication necessary is done by bit shifting: null bits are added to the least significant end of the data bus to increase the value of the data transferred. Where no multiplication is necessary, null bits are added to the most significant end of the data bus. This means that all of the buses are expanded from eight bits in width to ten bits in width by the weighting process.
2. The weighted samples are summed. In the design being considered here, it is subtracters which are used for this purpose as negative filter coefficients are required (-1, 0, 4, 0, -1).
3. The ten-bit wide output from the subtracters is rounded to eight bits.
4. Any negative values output from the subtracters are suppressed to zero. Negative values are identified via a flag produced by the subtracters.

Tasks 3 and 4 are necessary to simplify the design of the BCID logic. Without this simplification it is not possible to produce a BCID design capable of processing four channels of calorimeter data per XC4008 FPGA.

The data output from the FIR filter are fed into the second pipeline memory and from there into the peak finder. It can be seen that the peak finder consists of two comparators whose results are combined using an AND gate.

The output from the peak finder is combined with the *Peak Finder Override* slow-control input using an OR gate, and the resultant signal forms one of the two control inputs to a multiplexer. The other control input is the *FIR Filter Enable* slow-control signal. The multiplexer determines what data (if any) are output from the BCID logic. This is shown in table B.1. As described in section 6.4, in normal operation both the FIR Filter Enable and the Peak-Finder Override inputs should be set low.

| Inputs to multiplexer: | | | Data Output from BCID Logic: |
|---|---|---|---|
| *Peak-Finder Override* (slow control) | *FIR Filter Enable* (slow control) | Peak Finder Result | |
| 0 | 0 | 0 | Value of data output = zero. |
| 0 | 0 | 1 | Data output from pipeline containing unfiltered calorimeter data. |
| 0 | 1 | 0 | Value of data output = zero. |
| 0 | 1 | 1 | Data output from pipeline containing output of FIR filter. |
| 1 | 0 | 0 | Data output from pipeline containing unfiltered calorimeter data. |
| 1 | 0 | 1 | Data output from pipeline containing unfiltered calorimeter data. |
| 1 | 1 | 0 | Data output from pipeline containing output of FIR filter. |
| 1 | 1 | 1 | Data output from pipeline containing output of FIR filter. |

Table B.1. The control inputs to the multiplexer shown in figure B.3.

# List of References


[1] S. L. Glashow, Nucl. Phys. 22 (1961) 579,

S. Weinberg, Phys. Rev. Lett. 19 (1967) 1264,

A. Salam, 'Elementary Particle Theory', (Ed.) N. Svartholm, (Almqvist and Wikell, Stockhom, 1968) 367.

[2] F. Halzen and A. D. Martin, 'Quarks and Leptons: An Introductory Course in Modern Particle Physics', (Wiley, 1984).

[3] ATLAS letter of intent, CERN–LHCC/92–4.

[4] Design Study of the Large Hadron Collider (LHC), CERN/91–03.

[5] R. P. Feynman 'Quantum Electrodynamics', (Benjamin, New York, 1962).

[6] W. Marciano and H. Pagels, Phys. Reports 36C (1978) 137.

[7] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. Japan 49 (1973) 652.

[8] P. W. Higgs, Phys. Rev. Lett. 13 (1964) 508, Phys. Rev. 145 (1964) 1156.

[9] M. 'Quirós, Bounds on the Higgs mass in the Standard Model and Minimal Suspersymmetric Standard Model', CERN–TH.7507/94.

[10] H. B. Jensen, 'Evidence for Top Quark Production in 1.8 TeV $p\bar{p}$ Collisions', Proceedings of the 27th International Conference on High Energy Physics, Glasgow, July 1994.

F. Abe *et al.*, 'Observation of Top Quark Production in $p\bar{p}$ Collisions with the CDF Detector at Fermilab', FERMILAB–PUB–95/022–E.

[11] M. Carena *et al.*, 'Higgs Physics at LEP2', CERN/HEP–PH/9602250.

[12] H. Georgi and S. L. Glashow, Phys. Rev. Lett., 32 (1974) 438.

[13] D. H. Perkins, 'Introduction to High Energy Physics', 3rd edition, (Addison and Wesley, 1987).

[14] ATLAS technical proposal, CERN–LHCC/93–43.

[15] CMS technical proposal, CERN–LHCC/94–38.

[16] ALICE technical proposal, CERN–LHCC/95–71.

[17] LHC-B letter of intent, CERN–LHCC/95–5.

[18] I. P. Brawn *et al*, 'A Demonstrator Programme for the ATLAS Level-1 Trigger', ATLAS DAQ note DAQ–No–31 (1994).

N. Ellis *et al*., 'Level-1 trigger system', ATLAS DAQ note DAQ–No–5 (1992).

[19] I. P. Brawn *et al*, 'The Level-1 Calorimeter Trigger System for ATLAS', ATLAS DAQ note DAQ–No–30 (1994).

[20] H. Alexanian et al., Nucl. Instr. and Meth. A357 (1995) 306, Nucl. Instr. and Meth. A357 (1995) 318.

[21] RD3 proposal, CERN–DRDC/90–23.

[22] S. Gadomski *et al*., 'The Deconvolution Method of Fast Pulse Shaping at Hadron Colliders', CERN–PPE/92–23.

[23] P. W. Nicholson, 'Nuclear Electronics', (Wiley, 1974).

[24] GEM technical design report, GEM–TN–93–262.

O. Benary *et al*., 'Precision Timing with Liquid Ionisation Calorimeters', BNL–481459.

[25] I. P. Brawn *et al*., Nucl. Instr. and Meth. A349 (1994) 356.

[26] R. Staley *et al*., 'Implementation of a First-Level Calorimeter Trigger for use at the LHC at CERN', RD27 note 17 (1993).

[27] Spider write-up, CN Division, CERN.

[28] W. Zimmerman *et al*., 'A 16-channel VME flash ADC system (F1001–FADC)', H1 internal report, DESY, Hamburg (1989).

[29] I. Brawn, 'Preliminary Investigations into Bunch-Crossing Identification for the Level-1 Trigger', RD27 note 11 (1993).

[30] I. P. Brawn *et al*., 'Beam Tests of a Single-Channel Bunch-Crossing Identification Module for the Level-1 Trigger', RD27 note 31 (1993).

[31] The Programmable Logic Data Book, Xilinx Inc., 1993.

[32] I. P. Brawn *et al.*, ‘Bunch-Crossing Identification for the ATLAS First-Level Calorimeter Trigger’, Proceedings of the 1st Workshop on Electronics for LHC Experiments, Lisbon, September 1995.

[33] D. V. Camin *et al.*, ‘Front-End Electronics in III – V Technology for Liquid Argon Calorimetry’, Proceedings of the 1st Workshop on Electronics for LHC Experiments, Lisbon, September 1995.

R. L. Chase *et al.*, Nucl. Instr. and Meth. A343 (1994) 598.

[34] P. Hanke *et al.*, ‘A flash ADC for Trigger Tests with ATLAS Calorimetry’, RD27 note 41 (1995).

[35] J. Bishop *et al.*, ‘DULAC, a Compact Electromagnetic Calorimeter for the High Rate SSC Environment’, Proceedings of the 1st International Conference on Calorimetry in HEP (1990).

[36] GEANT detector simulation package, CERN program library long write-up, W5013.

[37] Private email from Alan Watson.