HD-IHEP 99-11

HD-ASIC 50-0899

A Compact Pre-Processor System for the ATLAS Level-1 Calorimeter Trigger

Institut für Hochenergiephysik
Universität Heidelberg

Dissertation
of
Ullrich Pfeiffer

DISSERTATION
submitted to the
Combined Faculties for the Natural Sciences and for Mathematics
of the Rupertus Carola University of
Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

A Compact Pre-Processor System for the ATLAS Level-1 Calorimeter Trigger

presented by
Ullrich Pfeiffer

Heidelberg, October 19, 1999

Zusammenfassung

Diese Arbeit beschreibt Untersuchungen und Entwicklungen für den Bau eines kompakten Pre-Processor Systems, das als Teil des Level-1 Kalorimeter Triggers zur Vorverarbeitung von Ereignissen am ATLAS Experiment beiträgt. Die Entwicklungsbeiträge beziehen sich auf die Funktionsweise und die Architektur des Pre-Processors. Dazu zählt die Entwicklung eines Multi-Chip Moduls (PPrD-MCM), das eine kompakte Signalverarbeitung von vier analogen Kalorimetersignalen (trigger-towers) ermöglicht. Die Signalverarbeitung umfaßt die Digitalisierung der Signale mit einer Genauigkeit von 8-bit, die zeitliche Identifikation zu einem Teilchenpaket (BCID), Transversalenergie Kalibrierung, die Auslese von Ereignisdaten und die serielle Datenübertragung zu den Triggerprozessoren.

Das Multi-Chip Modul besteht aus 9 Halbleiter Chips die auf einer Fläche von 15.9 cm^2 untergebracht sind. Das MCM wurde mit einer Strukturgröße von 100 μm in einem MCM-L Prozeß der Firma Würth Elektronik gefertigt. Ein applikationsspezifischer integrierter Schaltkreis wurde für das MCM entwickelt (Finco). Er wurde im 0.8 μm BiCMOS Prozeß der Firma Austria Micro Systems (AMS) gefertigt, er wird Flip-Chip Kontaktiert und ermöglicht unter anderem eine Verdopplung der seriellen Übertragungsrate.

Für saturierte trigger-tower Signale wurde eine effizient Methode entwickelt, die zeitliche Zugehörigkeit zu einem Teilchenpaket zu identifizieren. Ein Übertragungsformat (BC-mux) hat die effektive Bandbreite der seriellen Übertragung verdoppelt. Das PPrD-MCM wurde als Teil eines modularen Pre-Processor Testsystems in Betrieb genommen, wobei Auslese- und Übertragungstests gezeigt haben, daß die kompakte Verarbeitung von 64 trigger-tower Signalen auf einer VME Karte möglich sein wird.

Abstract

This thesis describes the research whose aim is to develop a compact Pre-Processor system for the ATLAS Level-1 Calorimeter Trigger. Contributions to the performance and the architecture of the Pre-Processor were made. A demonstrator Multi-Chip Module (PPrD-MCM) was developed and assembled which performs most of the preprocessing of four analogue trigger-tower signals. The preprocessing includes digitisation to 8-bit precision, identification of the corresponding bunch-crossing in time (BCID), calibration of the transverse energy, readout of raw trigger data, and high-speed serial data transmission to the trigger processors.

The demonstrator Multi-Chip Module has a size of 15.9 cm^2 and it consists of 9 dies. The MCM was designed with a smallest feature size of 100 μm and it was fabricated in a laminated MCM-L process offered by Würth Elektronik. A Flip-Chip interconnection ASIC (Finco) was developed for the PPrD-MCM and fabricated in a 0.8 μm BiCMOS-process offered by Austria Micro Systems (AMS). This ASIC has doubled the serial link speed and it is Flip-Chip mounted on the MCM substrate.

A BCID algorithm for saturated trigger-tower signals and a bunch-crossing multiplexing scheme (BC-mux) which doubles the effective bandwidth of the high-speed serial data transmission were developed. The PPrD-MCM was tested as part of a modular Pre-Processor test system, where transmission and readout tests have shown the feasibility of building a compact Pre-Processor system aimed at processing 64 trigger tower signals per VME board.

Contents

Introduction

The main objectives of the ATLAS experiment at the Large Hadron Collider (LHC) at CERN are the search for the Higgs boson, which is the only missing particle of the Standard Model (SM), the verification of the Standard Model, and the search for particles predicted by theories beyond the Standard Model in an energy range up to a few TeV. Observation of new particles may occur in a wide variety of decay channels, most of which suffer from their small production rate or from large QCD backgrounds. Therefore the ATLAS experiment requires a reliable trigger system, with high selectivity paired with optimal efficiency to retain interesting particle decays.

The Level-1 Trigger, which is the first-level of the ATLAS trigger system, achieves an event rate reduction from the 40 MHz LHC bunch-crossing rate down to the first level accept rate of 75 kHz. This is done by searching for muons, isolated electrons and photons, hadrons, jets of particles, and by calculating global energy sums of the calorimeter, within a total latency of 2.0 μs.

At the input to the calorimeter part of the Level-1 Trigger a Pre-Processor system performs the preprocessing of about 7200 analogue trigger-tower signals. The preprocessing includes digitisation, identification of the corresponding bunch-crossing in time (BCID), calibration of the transverse energy, rate monitoring, readout of raw trigger data, and high-speed data transmission to the following processors. The preprocessing of a large number of analogue signals requires a compact Pre-Processor system with fast hard-wired algorithms implemented in application-specific integrated circuits (ASICs).

The first chapter presents the LHC accelerator complex, the physics issues at the ATLAS experiment, and the ATLAS detector system. Next, an overview of the ATLAS trigger system is given with emphasis on the Level-1 Calorimeter Trigger architecture. The third chapter describes the Pre-Processor system of the ATLAS Level-1 Calorimeter Trigger. The contributions which were made to its bunch-crossing identification for saturated trigger-tower signals and its bunch-crossing multiplexing are described in Chapter 4. In Chapter 5 the development of a demonstrator Multi-Chip Module for the Pre-Processor (PPrD-MCM) is explained. It includes simulation and measurement results of the MCM temperature behaviour and the signal integrity. Chapter 6 is dedicated to the development of a Flip-

Chip interconnection ASIC (Finco) and Chapter 7 presents system measurements of the PPrD-MCM using a modular Pre-Processor test system. The Multi-Chip Module design experience obtained, and the bunch-crossing multiplexing scheme (BC-mux) which has contributed to the Level-1 Calorimeter Trigger architecture, have been published in parts within the Technical Design Report [TDR98] of the ATLAS Level-1 Trigger.

Chapter 1

The ATLAS experiment at the LHC

- The Large Hadron Collider
- Physics issues at the LHC
- The ATLAS detector

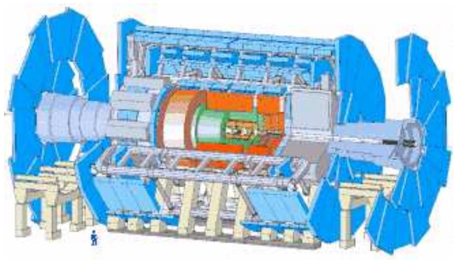

1.1 The Large Hadron Collider at CERN

The Large Hadron Collider (LHC) will reach higher energies than ever achieved in collider experiments before and thus open a new field of research in particle physics. It is now under construction at the European Laboratory for Particle Physics (CERN) and will provide proton–proton collisions in the year 2005 with a centre-of-mass energy of 14 TeV and a luminosity[1] of $\mathcal{L} = 10^{34}\text{cm}^{-2}\text{s}^{-1}$. In order to achieve this it must operate with 2961 bunches per beam at a very high intensity. The LHC will also operate for heavy-ion (Pb) physics at a luminosity of $\mathcal{L} = 10^{27}\text{cm}^{-2}\text{s}^{-1}$. The major challenge of the LHC machine is the superconducting magnet system with a dipole field of 8.3 T, operating in superfluid helium at 1.9 K. In contrast to electron–positron colliders the magnetic field is the limiting factor for a further increase of the LHC beam energy. The European Laboratory for Particle Physics is located near Geneva in Switzerland. See Figure 1.1 (a) for a geographical map of the LHC location and the CERN site.

This section starts with a short description of the LHC machine layout and the experiment locations. Because of their importance to the LHC experiments, the bunch structure, the beam luminosity, and the beam energy will be described next.

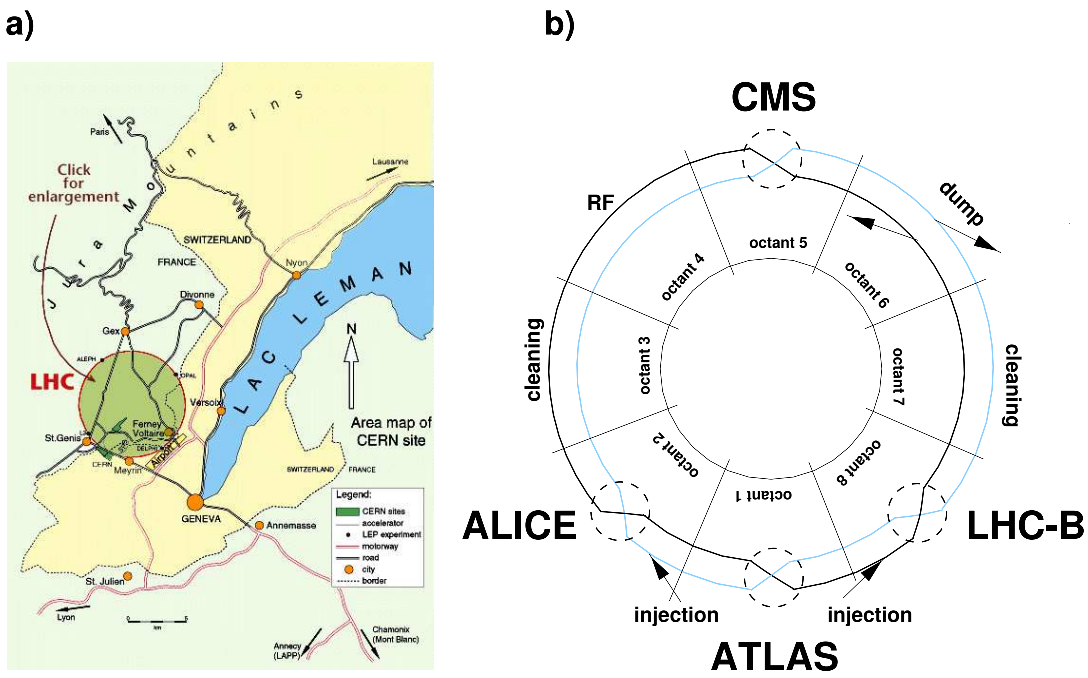

Figure 1.1: Geographical map of the LHC collider (a) and illustration of the experiment locations (b) [LHC99].

[1]The luminosity is the number of interactions per area and time.

Machine layout and experiment locations

The LHC accelerates bunches of protons or heavy ions in two separate circular beams, in opposite directions. It will be installed inside the existing LEP[2] tunnel in order to use most of the existing accelerator complex at CERN. The basic machine layout mirrors that of LEP, with eight long straight sections, each approximately 500 m in length, available for experimental insertions or utilities. The beams cross from one ring to the other at only four locations. Two 'high-luminosity' insertions are located at diametrically opposite straight sections. The ATLAS[3] experiment is located in octant 1 and the CMS[4] experiment is located in octant 5, see Figure 1.1 (b). A third experiment, optimised for 'heavy-ion' collisions (ALICE[5]), will be located at octant 2. A fourth experiment (LHCb[6]) has now been approved and will be located at octant 8.

The two general-purpose experiments, ATLAS and CMS, require a substantial amount of new civil engineering infrastructure, whilst the other two will be integrated into existing LEP caverns. Two 450 GeV/c beams from the SPS[7] are injected at octants 2 and 8. The other four long straight sections do not have beam crossings. Octant 3 and 7 are practically identical and are used for collimation of the beam halo in order to minimise the background in the experiments as well as the beam loss in the cryogenic parts of the machine. Finally, octant 6 contains the beam abort system, where the two beams are extracted using fast pulsed magnets and are transported to the external beam dumps.

The LHC will have two magnetic channels integrated into a common magnet (two-in-one magnet design), which is an alternative to two separate rings and allows enough free space in the existing (LEP) tunnel for a possible future re-installation of a lepton ring for electron–proton physics. In order to achieve the design energy of 7 TeV per beam with a dipole field of 8.3 T, the superconducting magnet system must operate in superfluid helium at 1.9 K. A summary of LHC performance parameters is given in Table 1.1.

Bunch structure

At the LHC the particles are stored in bunches, where each bunch contains about 10^{11} particles. Along the ring a total number of 2961 bunches are stored with rms bunch length of 0.075 m. Events can only come in time with the bunch-crossing guaranteeing a minimum time interval between two collisions of 25 ns. This time

[2]LEP: Large Electron Positron Collider.
[3]ATLAS: A Toroidal LHC Apparatus
[4]CMS: Compact Muon Solenoid
[5]ALICE:A Large Ion Collider Experiment
[6]LHCb: Study of CP violation in B-meson decays at the LHC
[7]SPS: Super Proton Synchrotron

LHC performance parameter	
Energy	7 TeV
Circumference	27 km
Dipole field	8.4 T
Coil aperture	56 mm
Distance between apertures	194 mm
Luminosity	$10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$
Beam-beam parameter	0.0034
Injection energy	450 GeV
Circulating current/beam	0.54 A
Bunch spacing	25 ns
Particles per bunch	10^{11}
Stored beam energy	334 MJ
Normalised transverse emittance	3.75 μm
r.m.s. bunch length	0.075 m
β-values at I.P. in collision	0.5 m
Full crossing angle	200 μrad
Beam lifetime	22 h
Luminosity lifetime	10 h
Energy loss per turn	6.7 keV
Critical photon energy	44.1 eV
Total radiated power per beam	3.6 kW

Table 1.1: LHC performance parameters [LHC]

separation corresponds to 40 MHz, which is the characteristic frequency for the LHC and its associated experiments and electronics.

The bunch structure of the LHC beams is determined by the injection and extraction systems of the accelerator chain. The bunches are grouped into 'trains' of 81 bunches, which in turn are grouped into 12 batches each containing three bunch trains. Figure 1.2 illustrates the LHC bunch structure. At full luminosity, an average of about 23 interactions will occur in each bunch-crossing every 25 ns. These interactions result in about 10,000 tracks observed within 100 ns, the typical duration of electronic signals in the detectors. The task of the LHC experiments is to identify and select the interesting events on top of this so-called 'pile-up' background. These pile-up events impose stringent requirements to the design and performance of the LHC detectors. Fast detector signals are required in order not to integrate the signals from pile-up events over too many bunch-crossings, and the detector signals must be highly granular in order to minimise the contribution of pile-up in a given detector cell.

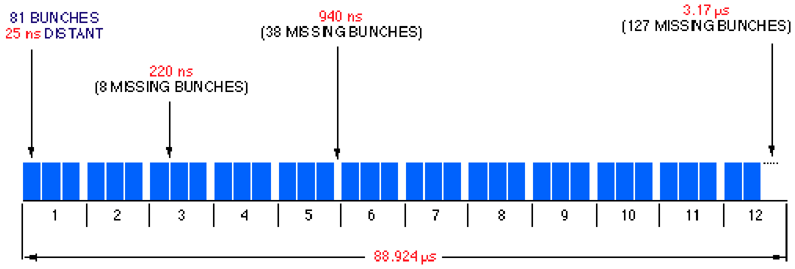

Figure 1.2: LHC bunch structure [TDC99].

Beam luminosity

New particles are produced by the inelastic interaction of two particles from the colliding bunches. Since the expected cross-section for interesting new particles σ reaches the level of at most $\sigma/\sigma_{tot} = 10^{-11}$ of the total cross-section ($\sigma_{tot} = \sigma_{el} + \sigma_{inel}$), the aim of the LHC is to squeeze the beams to small transverse sizes so that the luminosity $\mathcal{L}$, and hence the interaction rate, is sufficiently large to make rare processes accessible. Given the energy-dependent inelastic cross-section, σ_{inel}, and the beam luminosity $\mathcal{L}$, the particle production rate is:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \sigma_{inel}\mathcal{L}.$$

The luminosity $\mathcal{L}$ and the integrated luminosity L is calculated from the number of bunches B, the number of particles N_i in the colliding bunches ($i = 1, 2$), and the revolution frequency f_0 (11.2455 kHz) [Pov94]:

$$\mathcal{L} = \frac{B \cdot N_1 \cdot N_2 \cdot f_0}{A_{eff}} \qquad \text{and} \qquad L = \int \mathcal{L}\,\mathrm{d}t,$$

where A_{eff} is the effective bunch cross-section, which depends on the bunch shape and bunch-crossing angle. During the first 3 years the LHC will operate at a luminosity of $\mathcal{L} = 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, which is referred to as *low luminosity*. Later the luminosity will be increased by a factor of ten ($\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$), which is referred to as *high luminosity*.

The inelastic cross-section at the LHC beam energy is of the order of 70 millibarn (mb[8]). The inelastic production rate at high luminosity is therefore of the order of $7 \cdot 10^8$ Hz with about 23 inelastic collisions per bunch-crossing. For physics analysis, one requires measurements of the integrated luminosity as precise as possible, in order to convert an observed number of events to a cross-section. Typically a 5-10 % precision for the luminosity determination is assumed for measurements at

[8] The unit barn is defined as $1\ \mathrm{b} = 10^{-28}\,\mathrm{m}^2$, $1\ \mathrm{mb} = 10^{-27}\,\mathrm{cm}^2$

ATLAS. In addition to the integrated-luminosity measurement, monitoring of the instantaneous luminosity, possibly bunch-by-bunch, will also be done at ATLAS, e.g. for correction of the effect of pile-up on physics measurements [TDR99]. The high flux of particles from the proton–proton interactions places the detector and associated electronics in a high-radiation environment. Only radiation resistant detector components and read-out electronics can be used inside the detector.

Beam energy

Particle physics experiments at the LHC require a large beam energy for either the production of new particles with rest mass M *(mass scale)* or to probe the internal structure of nucleons *(length scale)*.

The kinetic energy of the colliding proton beams at the LHC will provide a centre-of-mass energy ($\sqrt{s}$) of 14 TeV, which ideally allows the creation of new particles which fulfil the relation:

$$\sqrt{s} \geq Mc^2.$$

The center-of-mass energy $\sqrt{s}$ of two colliding beam particles ($i = 1, 2$) with four momenta $p_i = (E_i, \mathbf{p}_i)$ is [Gre92]:

$$\sqrt{s} = \sqrt{(p_1 + p_2)^2}.$$

In case of two collinear colliding relativistic particles with the same rest mass M and the same beam energy E one gets:

$$\sqrt{s} = 2E.$$

Interacting proton constituents carry only a fraction of the proton momentum (x_a, x_b). Due to their momentum distributions, only a fraction of the beam energy $\widehat{s}$ can be used for creation of new particle masses:

$$\widehat{s} = x_a \cdot x_b \cdot s.$$

Hence the physics at the LHC is often referred to as 'TeV physics', where the main contribution is roughly at $\sqrt{x_a \cdot x_b} \approx 0.1$.

The length scale of the internal structure which one wants to explore in an experiment is related to the momentum $\mathbf{p}$ of the probing beam particles by de Broglie's equation:

$$\lambda = \frac{2\pi\hbar}{|\mathbf{p}|}.$$

The momentum $\mathbf{p}$ is given by the rest mass M and the relativistic energy E of the beam particles by:

$$E^2 = \mathbf{p}^2c^2 + M^2c^4.$$

1.2 Physics issues at the LHC

The Standard Model (SM) is a description of the known phenomena in the world of particles and their interactions, excluding gravity. It allows a very detailed description of today's high-energy physics experiments. However, there are many things the Standard Model does not explain and for this reason it is not regarded as a complete theory. An extension is needed to make the SM complete and theoretically consistent. Today's SM could then be seen as the limit of a new theory at low energy.

A new field of research in particle physics will be opened by the LHC, which will allow testing of the SM in a higher energy range where it has not been challenged before. This section starts with a short description of the Standard Model, followed by a summary of the physics issues at the LHC.

Standard Model

The Standard Model consists of a set of fundamental particles and their interactions via three of the four forces: electromagnetism, weak force, and strong force. Gravity is by far the weakest of the four forces and is not described by the SM.

Three types of particles are described: spin-1/2 fermions (matter particles), spin-1 gauge bosons, and a spin-0 Higgs boson. To make up 'normal' matter the spin-1/2 fermions, up- (u) and down- (d) quark, an electron (e), and the electron neutrino (ν_e) are required. For reasons unknown, this pattern is repeated three times, which leads to three 'generations'. Each generation consists of two quarks, a lepton and its neutrino. The three quark generations are:

$$\begin{pmatrix} u \\ d \end{pmatrix} \quad \begin{pmatrix} c \\ s \end{pmatrix} \quad \begin{pmatrix} t \\ b \end{pmatrix},$$

and the associated leptons:

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix} \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \quad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}.$$

The three fundamental forces are described by means of gauge theories and are mediated by the exchange of one or more bosons summarised in Table 1.2.

Force	Gauge symmetry	Bosons	Symbol	Strength
Weak	SU(2)	vector bosons	$W^{\pm}, Z^0$	$\alpha_{weak}= 0.03$
Electromagnetic	U(1)	photons	γ	$\alpha_{em}=1/137$
Strong	SU(3)	gluons	g	$\alpha_s \approx 0.12 .. 1$

Table 1.2: Fundamental forces and their associated bosons.

The electroweak symmetry-breaking mechanism of the SM predicts the existence of the spin-0 Higgs boson. This particle has not been observed so far and is the only missing piece of the Standard Model. The existence of the Higgs particle would then be responsible for the $W^{\pm}$ and Z^0 mass, and all masses of the SM could then be described. In a minimal supersymmetric extension of the Standard Model (MSSM), three neutral (h, H, A) and two charged ($M^{\pm}$) Higgs bosons are predicted instead of one.

Physics Issues

The physics issue at the LHC is, first and foremost, the problem of mass: is there an elementary Higgs boson, which is the origin of the spontaneous symmetry-breaking mechanism in the electroweak sector of the Standard Model? Next, the problem of flavour: why are there just three particle 'generations', what is the origin of their mass ratio and the generalised Cabibbo mixing angles, and what is the origin of CP violation? Finally, the problem of unification raises the question about neutrino masses and existence of proton decay. The ATLAS experiment is designed to resolve these questions by addressing the following physics programms: Higgs-boson physics, B physics, and supersymmetry.

Higgs-Boson Physics

What is the origin of different particle masses? Giving an answer to this question within the Standard Model means looking for the Higgs boson in a range from the expected limit of $m_H > 105$ GeV, at the and of LEP2, up to the theoretical upper limit of the order of 1 TeV. The observation of the Higgs boson at the LHC may occur in a wide variety of decay channels, most of which suffer from their small production rate or from very large QCD backgrounds. Excellent detector performance in terms of energy and momentum resolution and unprecedented particle-identification capabilities are required. Depending on the detector limits, some of the background events can be rejected, but others are irreducible. The ATLAS experiment will be able to explore the full range of possible Standard Model

Higgs boson masses. The following table lists the most promising decay channels and their associated Higgs mass range.

Decay channel				Higgs mass range [GeV]
1	$h \to bb$	and	$H \to bb$	$80 < m_H < 110$
2	$h \to \gamma\gamma$	and	$H \to \gamma\gamma$	$95 < m_H < 130$
3	$H \to ZZ^* \to 4l$			$130 < m_H < 700$
4	$H \to ZZ \to ll\nu\nu$	and	$H \to WW \to l\nu jj$	$600 < m_H < 1000$

The choice of which decay channel is used depends on the signal rates and the signal-to-background ratios in the various mass regions. The total Higgs boson cross-section has contributions from various subprocesses, of which gg fusion and WW fusion are the most important ones. The stringent requirements on the ATLAS detector system for the discovery of the numbered decay channels are as follows:

Decay 1: This decay requires the identification of b-jets, using mainly vertex but also soft-lepton tags.

Decay 2: This decay channel requires excellent energy and angular resolution in the electromagnetic calorimeter, and rejection against jets faking photons.

Decay 3: Z-decay into electrons and muons requires excellent charged-lepton momentum/energy resolution and a very good electron and muon identification capability.

Decay 4: This channel can be observed above background from W/Z plus jet production, provided the calorimeters have excellent hermetic coverage down to small angles with respect to the beam.

Figure 1.3 shows that ATLAS can discover a Standard Model Higgs boson in the mass range $90 < m_H < 1000$ GeV after a few years of operation at the LHC. The significance for each mass point is defined as $S/\sqrt{B}$, where S and B are the numbers of accepted signal and background events in a chosen mass window.

The plot in Figure 1.4 (a) shows the ATLAS discovery potential in the Higgs boson sector of the MSSM after a few years of operation at the LHC. A large fraction of the MSSM parameter space (m_H and $\tan\beta$) will be covered using a variety of Higgs boson decay channels. Figure 1.4 (b) shows that as much as ten years of operation will be needed to cover the MSSM parameter range completely.

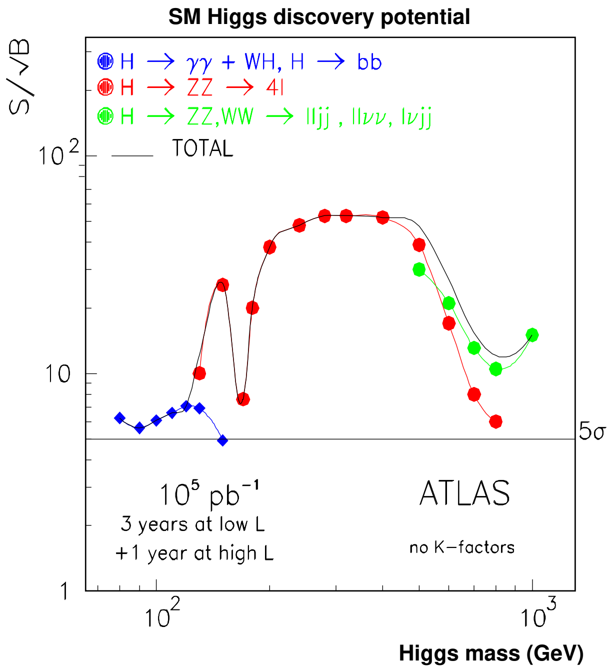

Figure 1.3: ATLAS discovery potential (signal over background events) for a Standard Model Higgs boson after three years of operation at low luminosity and one year of operation at high luminosity [TDR99].

B Physics

What is the reason for three particle 'generations'? Answers to this question could came from B-physics. The main emphasis will be on the precise measurement of CP violation in the B_d^0 system and the determination of the angles in the Cabibbo-Kobayashi-Maskawa unitarity triangle. CP violation has so far been observed only in the neutral kaon system, but the origin of CP violation is not yet understood. By observing the CP violation in B-meson decays the origin of CP violation could be explained. In addition, $B\bar{B}$ mixing in the B_s^0 system and B decays can be studied.

Search for Supersymmetry

What does a further unification look like? Grand Unified Theories (GUT) share some of the Standard Model problems: too many arbitrary parameters, and no integration with gravity. However, they are genuine unified field theories because they have only one coupling constant. Furthermore, they make the prediction that the proton will decay.

Supersymmetry (SUSY) is a theory which relates the properties of bosons to those of fermions, such that each particle acquires a supersymmetric partner with a spin

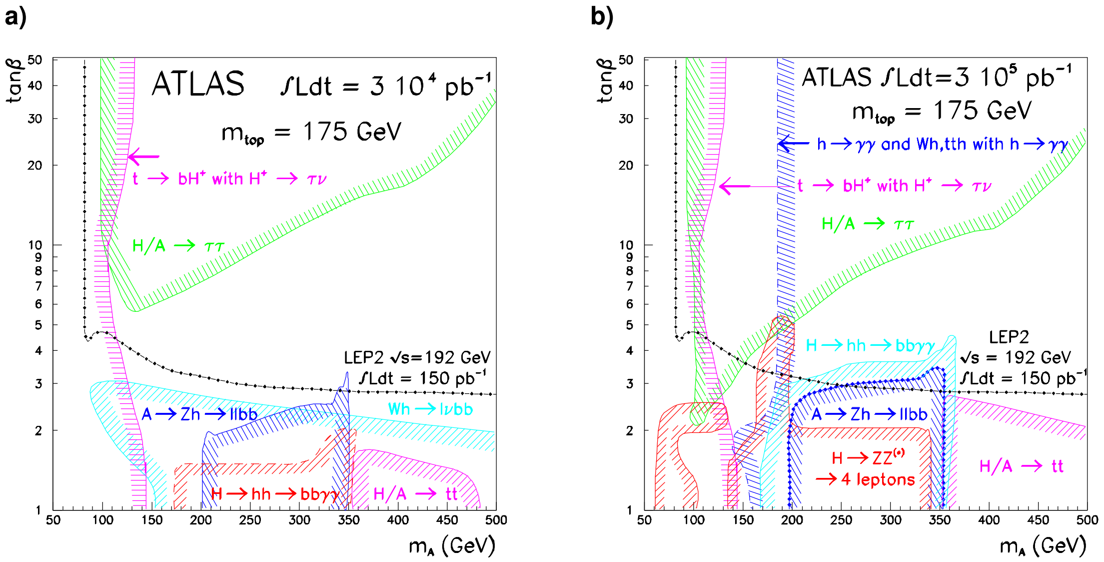

Figure 1.4: Discovery region of the MSSM parameter space covered by searches for Higgs-boson signatures with ATLAS after three years of operation at low luminosity (a). The discovery region after several years of operation at high luminosity is shown in (b) [TDR99].

differing by 1/2. The supersymmetry particle spectrum is therefore quite rich. It consists of spin-0 (squarks, sleptons) and spin-1/2 (gluino, charginos, neutralinos) particles. The lightest supersymmetric particle (LSP) is supposed to be stable and weakly interacting, implying signatures with large E_T^{miss}. The R-parity is of importance for SUSY experiments, which implies baryon (B) and lepton (L) conservation. The R-parity is defined as: $R = (-1)^{3 \cdot B + L + 2 \cdot S}$, where $R = +1$ is valid for particles and $R = -1$ for sparticles, the supersymmetric partners of particles. Therefore sparticles can only be created in pairs and they can only decay to an odd number of sparticles.

1.3 The ATLAS detector

The spectrum of physics studies described in Section 1.2 led to the ATLAS detector concept. It was first presented in the Letter of Intent [LOI92] and later in the Technical Proposal [TP94]. Since then the design has evolved guided by detailed physics performance studies and experience from research and development (R&D) programs. The ATLAS experiment will be designed, constructed and operated by a world-wide collaboration of scientists and engineers (~1,800 members) from 146 institutions and 33 countries. The basic design criteria of the detector, required to meet the physics goals, can be summarised as follows:

- Very good *electromagnetic calorimetry*, complemented by full coverage *hadronic calorimetry*. The electromagnetic calorimetry is used for electron and photon identification and measurements. The hadronic calorimeter is used for accurate jet and missing transverse energy measurements.

- A high-precision *muon momentum measurement*, with the cabability of accurate measurement of the bending of high-p_T muon tracks at high luminosity.

- An efficient *tracking* system for high-p_T lepton-momentum measurements, electron and photon identification, τ-lepton and heavy-flavour identification at high luminosity, and full event reconstruction at low luminosity.

- A *large acceptance* in pseudorapidity η with full coverage of the azimuthal angle ϕ everywhere. The azimuthal angle is measured around the beam axis, whereas pseudorapidity[9] relates to the polar angle θ ($\eta = -\ln\tan(\theta/2)$).

- And finally, a highly efficient *trigger system* for triggering and measurement of particles at low-p_T thresholds, providing high efficiencies for most of the physics processes of interest at LHC.

In the following section a brief description of the sub-detectors of the ATLAS experiment is given. Most of the information is taken from the Technical Design Reports of the sub-systems of the experiment [TP94], [LAr98].

1.3.1 Detector description

Depending on the physics goals and the need to keep the cost at a reasonable level, different design philosophies have been adopted among the four experiments proposed for the LHC. The ATLAS experiment uses a large air-core toroid system for its muon spectrometer. The electromagnetic (EM) calorimetry uses the liquid-argon technique. In the barrel, an iron-scintillator hadronic calorimeter is used and in the endcap the hadronic calorimeter is of the liquid-argon type. In front of the barrel EM calorimeter a superconducting-solenoid coil provides a 2 T field, integrated in the same cryostat. The inner-tracking system consists of semiconductor detectors in the inner-most part and straw-tubes in the outer part. The overall ATLAS detector layout is shown in Figure 1.5. Each of the shown sub-detector components are described in the following.

[9]For $\beta \to 1$ one gets the pseudorapidity η as $\lim_{\beta\to 1} Y = -\ln\tan(\theta/2)$, where Rapidity Y is defined as $Y = \frac{1}{2}\ln\frac{E+p\cdot\cos\theta}{E-p\cdot\cos\theta}$. Rapidity is introduced because differences are invariant under Lorentz transformation. In general, light particle 'boosts' are found in forward or backward direction ($|\eta| > 3$), having a large cross-section. Physics in this η coverage is often referred to as 'minimum bias' physics, where only few samples can be used out of the full event rate. Heavier particles with a smaller cross-section are produced centrally at $\eta \approx 0$.

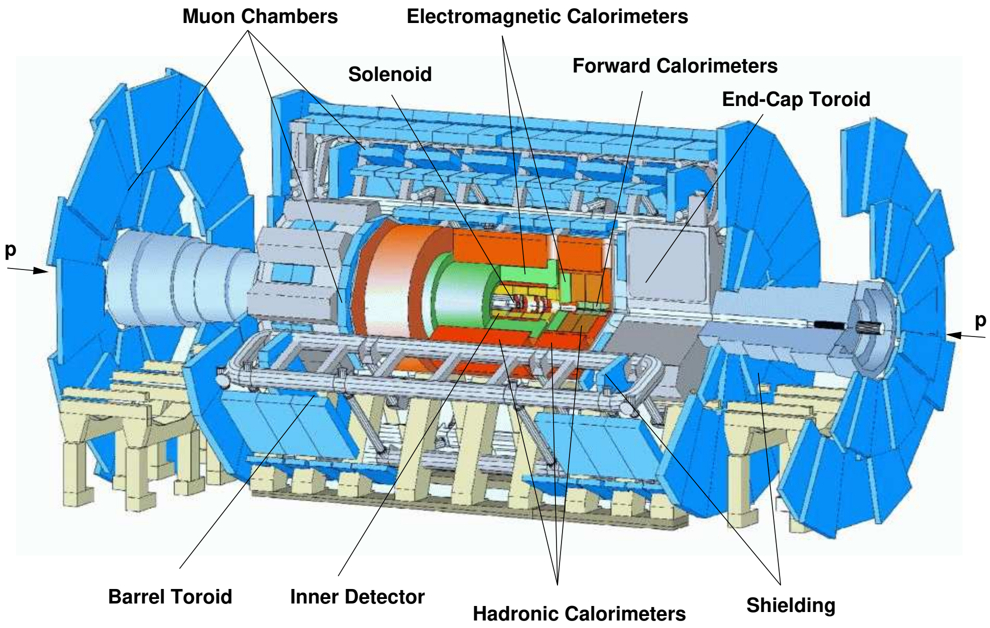

Figure 1.5: The ATLAS detector.

Magnet configuration

The ATLAS magnetic configuration consists of two magnetic systems. A thin solenoid magnet with a field of 2 T surrounds the inner detector cavity, and large air-core toroids consisting of independent coils are arranged outside the calorimeter.

Inner Detector

The task of the Inner Detector is to reconstruct the tracks and vertices in the event with high efficiency. It contributes to the electron, photon and muon identification and it provides the signature for short-lived particle decay vertices. Its acceptance covers the range $|\eta| < 2.5$. The Inner Detector is contained within a cylinder of a length of about 7 m and a radius of 1.15 m. It consists of high-resolution semiconductor pixel and strip detectors in the inner part, and continuous straw-tube tracking detectors with transition radiation capability in the outer part. The requirements for the Inner Detector are to be very fast, to be radiation hard, to have a fine granularity and good momentum resolution. In addition, the amount of material should be kept as small as possible, otherwise the momentum resolution in the tracker itself will be degraded and the energy measurement in the calorimeters will be spoiled.

Calorimetry overview

The calorimeters will play a crucial role at the LHC. They are required to measure the energy and direction of photons, electrons, isolated hadrons and jets, as well as the missing transverse energy. The calorimeters will be the leading detectors for many measurements in physics channels of prime interest. Combined with the inner tracker, calorimeter measurements are used for electron and photon identification. A fast detector response (<50 ns) and fine granularity are required to minimise the impact of pile-up on the physics performance. High radiation resistance is also needed for the given high particle flux expected over a period of operation of at least ten years. The ATLAS calorimetry covers a range of $|\eta| < 5$ using different techniques as best suited to the different requirements and radiation environments. The rapidity coverage and the basic granularity of the calorimeters is summarised in Table 1.3. The EM calorimeter system is contained in a cylinder which has a radius of 2.25 m and a total length of 6.65 m along the beam axis. The barrel hadronic calorimeter system has an outer radius of 4.23 m and a total length of about 12 m. The total weight of the calorimeter system is about 4,000 T. The end-cap EM, Hadronic and Forward calorimeter are hosted in the same cryostat.

Calorimeter System	η coverage	Granularity ($\Delta\eta \times \Delta\phi$)
EM Barrel	$\|\eta\| < 1.475$	0.003×0.1 (s1) 0.025×0.025 (s2) 0.05×0.025 (s3)
EM End-Cap	$1.375 < \|\eta\| < 3.2$	0.003–0.1 $\times$ 0.1
Presampler	$\|\eta\| < 1.8$	0.025×0 .1
Hadronic Barrel	$\|\eta\| < 1.7$	0.1×0.1
Hadronic End-Cap	$1.5 < \|\eta\| < 2.5$ $2.5 < \|\eta\| < 3.2$	0.1×0.1 0.2×0.2
Forward Calorimeter	$3.1 < \|\eta\| < 4.9$	$\sim 0.2 \times 0.2$

Table 1.3: The ATLAS calorimeter granularity

Electromagnetic Calorimeter

The Electromagnetic Calorimeter is a lead liquid/argon (LAr) detector with accordion geometry. In the pseudorapidity range $|\eta| < 1.8$ it is preceded by a presampler detector, installed behind the cryostat cold wall, which is used to correct the energy loss in the material (inner detector, cryostats, coil) upstream the calorimeter. The total thickness of the calorimeter is $\sim$25 radiation lengths (X_0) in the barrel and $\sim 26\,X_0$ in the end-caps. The total number of channels is about 200,000. The calorimeter energy resolution can be parametrised as follows:

$$\frac{\Delta E}{E} = \frac{a(\%)}{\sqrt{E}} \oplus \frac{b}{E} \oplus c(\%)$$

where E is the energy in GeV, a is the sampling term, b is the noise term, and c the constant term. The electron energy resolution obtained from a test programme of a detector prototype module over the full η range is (9.90–10.4)% for the sampling term, (280–520) MeV for the noise term, and (0.27–0.35)% for the local constant term. The overall constant term is 0.7 %, which puts stringent requirements on the detector construction, the dead material in front and the calibration precision. The constant term dominates at high energies and is therefore very important for the LHC experiments.

Tile Hadronic Calorimeter

The Tile Hadronic Calorimeter covers the range $|\eta| < 1.6$, and it consists of a barrel and two extended barrel cylinders. The hadronic barrel calorimeter has an inner radius of 2.28 m and an outer radius of 4.23 m, divided into three sections: the central barrel and two extended barrels. It is based on a sampling technique with plastic scintillator plates (tiles) embedded in an iron absorber matrix and read out by wave-length shifting fibres. The tiles are placed in planes perpendicular to the beam axis and staggered in depth, simplifying the mechanical construction. The calorimeter is segmented in three layers with different thicknesses measured in multiples of hadronic interaction length λ. The layers are approximately 1.4, 4.0 and 1.8 λ thick at $\eta = 0$. The total number of channels is of the order of 10,000. The Hadronic Calorimeter is placed behind the EM Calorimeter and the solenoid coil, resulting in a total active calorimeter thickness (EM + Tile) of 9.2 λ. The total amount of material in front of the muon system, including the support structure of the Tile calorimeter, is about 11 λ at $\eta = 0$.

Liquid Argon Hadronic Calorimeter

The Liquid Argon Hadronic Calorimeter covers a range of $1.5 < |\eta| < 4.9$. In the end-caps, which extend to $|\eta| < 3.2$, the hadronic calorimeter also makes use of the LAr technology, and shares the cryostat with the EM end-caps. The same cryostat also houses the special LAr forward calorimeters, which extend the pseudorapidity coverage to $|\eta| = 4.9$. Each hadronic end-cap calorimeter consists of two, equal diameter, independent wheels. The first wheel is built out of 25 mm thick copper plates, while the second one uses 50 mm plates. The active part of the end-cap calorimeter is $\sim 12\,\lambda$ deep.

Forward Calorimeter

The ATLAS Forward Calorimeter is integrated in the end-cap cryostat. Its front face is at about 5 m from the interaction point. This makes the forward calorimeter a particularly challenging detector due to the high level of radiation. The use of the Forward Calorimeter reduces the effects of the crack in the transition region around $|\eta| = 3.1$. A clear benefit comes from the continuity of the η coverage up to 4.9, with advantages to the forward jet tagging and the reduction of the tails in the E_T^{miss} distribution. The Forward Calorimeter has to accommodate at least 9 λ of active detector material in a rather short longitudinal space. Thus it is a high density detector, consisting of three longitudinal sections. The first section is made of copper and the other two of tungsten. Each of them consists of a metal matrix with regularly spaced longitudinal channels filled with rods. The sensitive medium is Liquid Argon which fills the gap between the rod and the matrix.

Muon Spectrometer

High-momentum final-state muons are among others the most promising and robust signatures of physics at the LHC. The discovery potential of the spectrometer has been optimised on the basis of selected benchmark processes. These processes are, in particular, Standard Model and supersymmetric Higgs decays and new vector bosons. The performance of the apparatus for low transverse momenta is of interest for B physics and CP violation. Important parameters that need to be optimised for maximum physics reach are: resolution, second-coordinate measurement, rapidity coverage of track reconstruction, trigger selectivity and trigger coverage. To exploit this potential, the ATLAS collaboration has designed a high-resolution muon spectrometer with stand-alone triggering and momentum measurement capability.

The muon spectrometer exploits the magnetic deflection of muon tracks in a system of three large superconducting air-core toroid magnets. The spectrometer is instrumented with separate-function triggers and high-precision tracking chambers. Over most of the pseudorapidity range, a precision measurement of the track coordinates in the principal bending direction of the magnetic field is provided by Monitored Drift Tubes (MDT). In order to provide a finer granularity, which is required to cope with the demanding rate and background conditions, Cathode Strip Chambers (CSC) are used in the first station of the end-cap region ($|\eta| < 2$).

Muon Trigger Chambers

The Trigger Chambers for the ATLAS muon spectrometer are based on two different types of detectors. Resistive Plate Chambers (RPC) are used in the barrel ($|\eta| < 2.4$) and Thin Gap Chambers in the end-cap region. The trigger chambers cover a total area of about 3650 m^2 in the barrel and 2900 m^2 in the end-cap region. The RPC is a gaseous detector providing a typical space-time resolution of 1 cm $\times$ 1 ns with digital readout. The basic RPC unit is a narrow gas gap formed by two parallel resistive plates, separated by insulating spacers. The TGC chambers are designed in a way similar to multiwire proportional chambers.

Analogue trigger tower signals

The ATLAS Level-1 Calorimeter Trigger, which will be described next, requires reduced-granularity data from the calorimeters as input. It will work with values of transverse energy E_T and it requires a dynamic range of at least 250 GeV. A coarser granularity for pseudorapidity η and azimuthal angle ϕ is formed in 'trigger-towers'. The output from one trigger-tower is an analogue signal, which in general represents an area of 0.1×0.1 ($\Delta\eta\times\Delta\phi$). The 'building' of trigger-towers is done by the calorimeter electronics, separately for electromagnetic and hadronic. A trigger tower is build from up to 60 calorimeter cells. The signal from each cell is fed through a number of amplifiers, shapers, buffers, delay lines, and summing circuits before reaching the first processor of the ATLAS Level-1 Calorimeter Trigger.

Chapter 2

The ATLAS trigger system

- Trigger system overview
- Level-1 trigger architecture

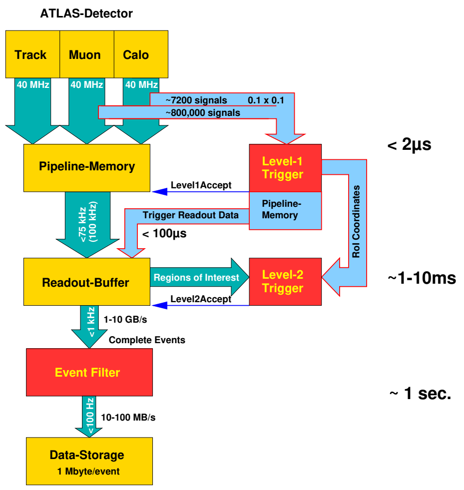

2.1 Trigger system overview

The high interaction rate of the LHC puts stringent requirements on the trigger and data-acquisition systems. Only a tiny fraction of interactions can be recorded for off-line analysis, requiring a trigger selectivity of about one interaction in 10^7. Furthermore, massive amounts of data have to be transmitted to and stored in buffer memories, while the trigger system performs its calculations. Since it is not possible to make a trigger decision within 25 ns between bunch-crossings, the so-called 'pipelined' readout has to be used, where data are stored in pipeline memories within electronics mounted on the detector.

The overall ATLAS trigger system is divided into three levels of event selection: the Level-1 Trigger, the Level-2 Trigger and the Event Filter. Each trigger level contributes to a rate reduction of selected events consisting of data from the individual sub-detectors. Figure 2.1 shows the connection of the detector readout with the ATLAS trigger system. Data from the tracking system, the muon detectors, and the calorimeters are stored in pipeline memories for each LHC bunch-crossing at a rate of 40 MHz.

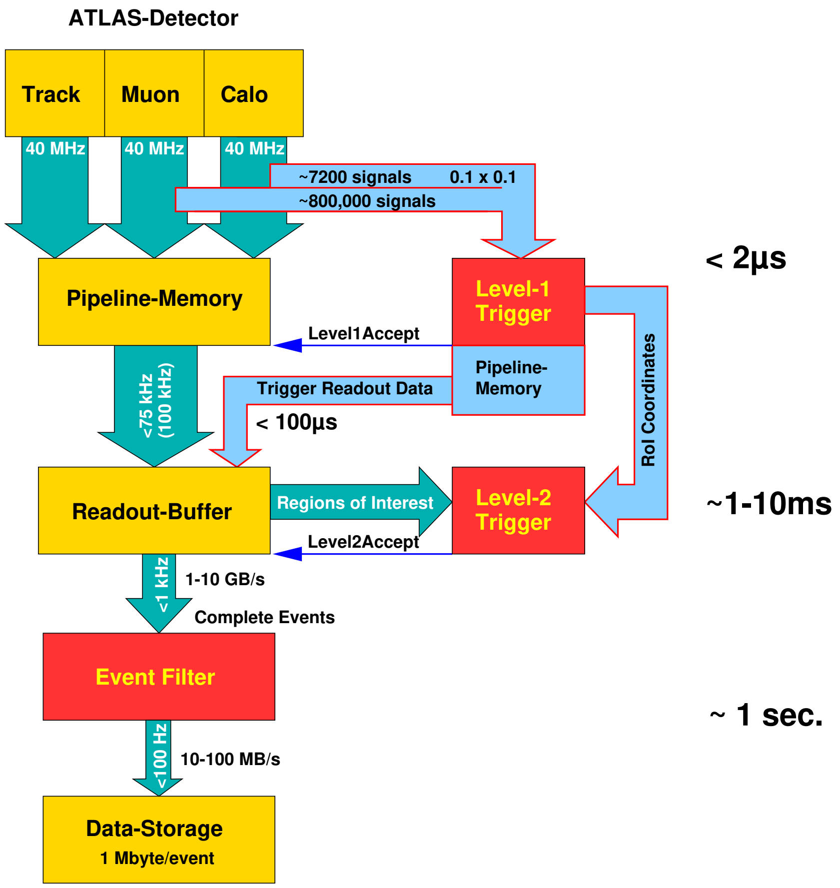

Figure 2.1: Connection of the Atlas Detector readout with the Atlas Trigger System.

The first rate reduction is achieved by the Level-1 Trigger, which is designed to reduce the event rate to 75 kHz, upgradable to 100 kHz. In order to reduce the size, and hence cost of pipeline memories, the Level-1 Trigger has to make a decision as fast as possible. Therefore the total latency (including cable delays) for the Level-1 Trigger is limited to 2 μs (80 bunch-crossings) with a contingency of 0.5 μs. The decision of the Level-1 Trigger is indicated by a level-1 accept signal, which is distributed to the detector front-end electronics, where it forces the electronics to keep only accepted event data. The level-1 accepted event data are sent from the detector front-end electronics to the Readout-Buffer memories (ROBs).

Coordinates of detector regions (Regions of Interest, RoIs) inside which the Level-1 Trigger has identified local energy depositions (e.g. hadronic clusters and/or electromagnetic clusters) or muons are given to the Level-2 Trigger. Using the RoI information, the Level-2 Trigger selectively accesses data from the ROBs, moving only data that are required in order to make the level-2 decision. The Level-2 Trigger has access to all of the event data, if necessary with the full precision and granularity. Based on additional track information it confirms the level-1 decision, that electrons/photons, jets, hadrons, or muons were found. The latency of the Level-2 Trigger is variable from event to event and is expected to be in a range 1–10 ms.

Remaining event data are sent to the Event Filter with an event rate of about 1 kHz ($\sim$1 GB/s). The Event Filter makes decisions on the basis of full event data, similar to offline analysis, and causes the data to be permanently stored with an event rate of about 100 Hz ($\sim$100 MB/s). Each event which is stored for offline analysis has a size of about 1 MB. In addition to the detector readout, data from each trigger level are recorded in order to keep the raw information on which the event was selected.

Level-1 Trigger

The Level-1 Trigger is a fast pipelined system for the selection of rare physics processes. Its selectivity achieves an event rate reduction from the 40 MHz LHC bunch-crossing rate down to the first level accept rate of 75 kHz. This is done by searching for isolated electrons and photons, hadrons, jets of particles, muons, and by calculating calorimeter global energy sums within 2.0 μs latency. This latency is taken up partly by cable delays. Hence, fast hard-wired algorithms implemented in application-specific integrated circuits (ASICs) and field programmable gate arrays (FPGAs) are required. For that reason the Level-1 Trigger is also referred to as 'hardware' trigger.

The Level-1 Trigger itself is divided into three subsystems, the Calorimeter Trigger, the Muon Trigger and the Central Trigger Processor (CTP). Figure 2.2 shows a block diagram of the overall Level-1 Trigger architecture.

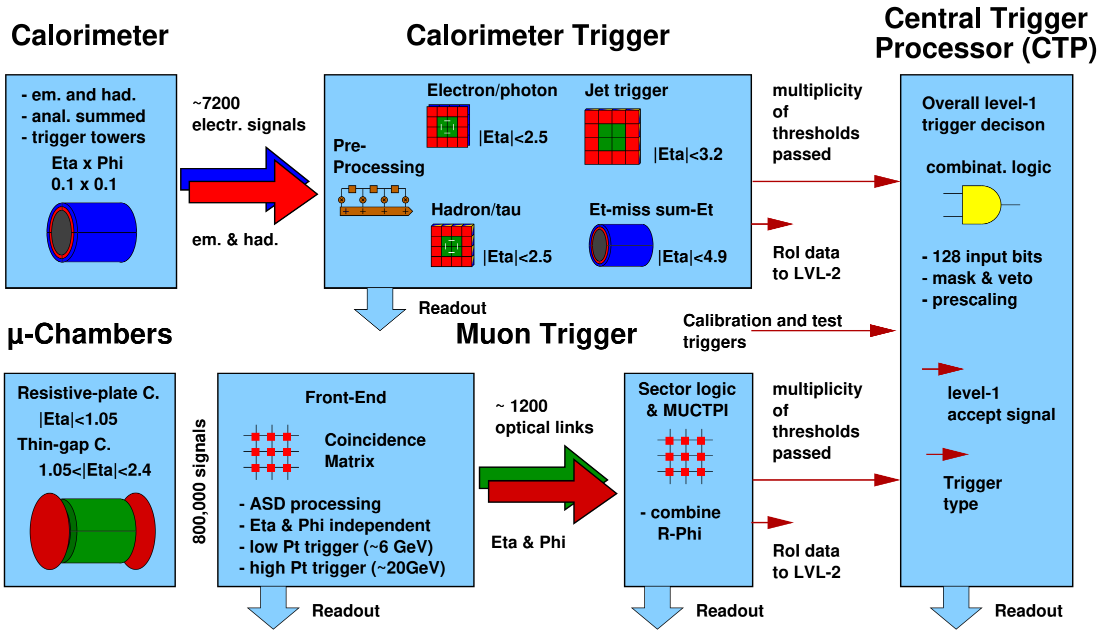

Figure 2.2: Overall Level-1 Calorimeter Trigger Architecture.

The Calorimeter Trigger gets analogue input signals from the electromagnetic and the hadronic calorimeters. These signals are summed separately in order to form trigger tower signals with a granularity of 0.1 x 0.1 in the η and ϕ directions. All in all the Calorimeter Trigger has about 7200 analogue input signals, which are transmitted electrically via twisted-pair cables from the detector to the Level-1 Trigger electronics, located in the trigger cavern. The maximum cable length for trigger input signals is 60 m. Based on these input signals the Calorimeter Trigger performs the following tasks:

- preprocessing of input signals;
- an electron/photon trigger algorithm;
- a hadron/tau trigger algorithm;
- a jet trigger algorithm;
- calculation of global trigger quantities (E_T-miss and sum-E_T).

For a detailed explanation of the individual algorithms and the associated processors see Section 2.2. The results from the Calorimeter Trigger are the multiplicity of showers passing transverse-energy thresholds, and the coordinates of regions of interrest (RoI). The multiplicity information is sent to the Central Trigger Processor (CTP) and the RoI data is sent to the Level-2 Trigger.

The Muon Trigger gets input signals from resistive-plate trigger chambers for the barrel region with an $|\eta|$-coverage smaller than 1.05. For the endcap region ($1.05 < |\eta| < 2.4$) it gets signals from thin-gap trigger chambers. See Section 1.3.1 for a description of the muon chambers. All in all 800,000 signals from the μ-chambers make up the input to the Muon Trigger. The Muon Trigger electronics is split up into electronics located on the detector and electronics located in the trigger cavern. The number of optical links, which are required to transmit data from the detector to the trigger cavern, is about 1200. For a more detailed description of the Muon Trigger architecture see Section 2.2.2.

The role of the CTP is to combine the threshold multiplicities from the Calorimeter and the Muon Triggers to make the final Level-1 Trigger decision. This is done by first classifying the muon and calorimeter results in up to 96 trigger 'objects', and then to combine these into trigger menu 'items' by doing boolean logic on the 'objects'. The decision is distributed via the Timing, Trigger and Control system (TTC) to the front-end electronics of the detector systems, where it initiates the readout. See Section 2.2.3 for a description of the CTP architecture and the trigger objects which are classified.

Level-2 Trigger

Event processing at the Level-2 trigger can be broken down decomposed into a number of broad steps:

- *Feature extraction*: Provide physics parameters for each RoI of the individual sub-detectors (data collection and data preprocessing).
- *Object building*: Combine several subdetectors for one RoI.
- *Trigger type selection*: Set flags for each item of a physics menu for which a match is found.

In feature extraction, the data from one RoI is gathered and processed to give physics-like quantities. For the calorimeter, this process takes cell information and produces cluster parameters. For the tracker, the basic hit information is converted to track or track-segment parameters. Object building takes the features for one RoI from all relevant detectors and returns the particle parameters and if possible the particle type. In the trigger-type selection phase all the objects found in an event are combined and compared with the topologies for a menu of physics selection (trigger menu). Items in the menu express physics selections in terms of seven basic trigger objects: muons, electrons, photons, charged hadrons, taus, jets, and missing-E_T. Each of the items defines a set of trigger conditions, including the required multiplicity of each object and the requirements on properties, e.g. p_T thresholds and isolation criteria.

The architecture of the Level-2 trigger system can be described in terms of four main functional blocks: the ROB complex, the supervisor, the network and the processing farms. The ROB is a standard functional component common to both the Level-2 Trigger and the Event Filter data flow. Its main function is to store the event fragments during the level-2 latency. The supervisor has three main functions. First it receives the RoI information from level-1, second it assigns processors as needed and it has to route the RoI data to the appropriate processors. Finally it has to receive the Level-2 decisions and broadcast them to the ROB complex. For the network one can distinguish two types of logical networks: a 'data collection' network which has to transmit RoI data and a 'control network' which carries control data, trigger accept signals and RoI requests.

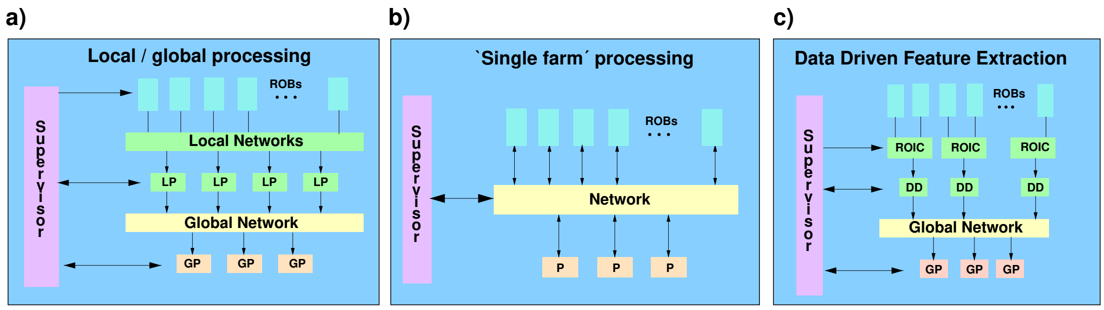

Figure 2.3: Level-2 Trigger Architecture.

The fundamental technical problem for the implementation of the Level-2 Trigger is to make the data available to the processors executing the trigger algorithms and to control the use of these processors. One should keep in mind that the full size of each event is about 1 Mbyte and the event data are spread over about 1700 ROBs of the readout system. The algorithms are expected to require a processing power of the order of 10^6 MIPS[1]. Figure 2.3 shows three conceptual designs proposed for the Level-2 Trigger architecture. These three designs have evolved because algorithms for object building differ and they are performed for each sub-detector individually. Hence, the object building algorithms could run on local processors assigned to the sub-detectors (see Figure 2.3 (a)) or as global processors (see Figure 2.3 (b)). The third option is to implement object building in 'hardware' processors based on FPGAs (see Figure 2.3 (c)).

The algorithms are derived from 'off-line' analysis code and are optimised for the individual target processors. The actual technology of the processors (RISC[2] or CISC[3]) is not yet chosen. The aim is to benefit from industrial developments in computing power and price, and hence to take a technology decision as late as possible. The real issue for the Level-2 architecture is the network technology

[1] MIPS: Mega Instructions Per Second
[2] RISC: Reduced Instruction Set
[3] CISC: Complex Instruction Set

decision. Various demonstrator projects for a reduced number of network nodes are currently in use: a 8-16 node Fast/Gigabit Ethernet, a 32 node ATM[4], and 6 node SCI[5] network.

Event Filter

In contrast to the Level-2 Trigger, which processes copies of relevant event fragments while the event data remains in the readout buffers, the Event Filter is functionally part of the main Data Acquisition (DAQ) data-flow. In fact, while being processed by the Event Filter program, events are stored in the event filter processor memories. The current ATLAS strategy for event filtering is to use the offline reconstruction and analysis programs to the maximum possible extent. This software will be the first element in the DAQ which has access to the full event and calibration data. It can also perform monitoring and calibration/alignment studies, e.g. trigger performance checks. The exact selection criteria to be used in the Event Filter are not yet chosen, since the final optimisation will most likely be performed using data taken during the initial operation.

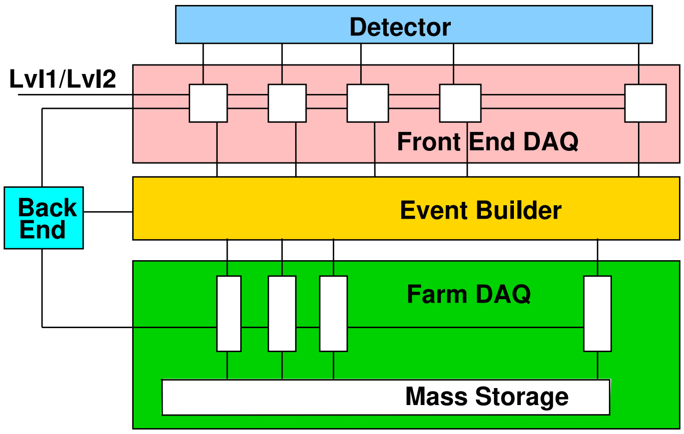

Figure 2.4: Event Filter Architecture.

The technology aspects which are valid for the Level-2 Trigger remain also for the Event Filter: delay the final technology choice as long as possible. Figure 2.4 shows a functional block diagram of an Event Filter prototype architecture. The figure shows the collection and buffering of data from the detector (the front-end DAQ), the merging of fragments into full events (the event builder) and the interaction with the Event Filter (the farm DAQ). The interface used for 'run' control is referred to as the 'Back-End' system. For the final Event Filter architecture a combination of Level-2 and the Event Filter is not yet excluded. They are both limited by

[4] ATM: Asynchronous Transfer Mode
[5] SCI: Scalable Coherent Interface

network technologies, they are hardly IO bound, and both are based on software programs running on computers. The Level-2 and the Event Filter are therefore often referred to as 'software' triggers.

2.2 The Level-1 Trigger system architecture

The Level-1 Trigger is divided into three subsystems: the Calorimeter Trigger, the Muon Trigger and the Central Trigger Processor. This has already been described in the trigger overview Section 2.1, and it is illustrated in Figure 2.2. Here, a more detailed description of the architecture of each of the individual subsystems is given. Emphasis is placed on the Calorimeter Trigger architecture with its algorithms and the associated subprocessor electronics modules: the Pre-Processor Module, the Cluster Processor Module, and the Jet/Energy-sum Module.

2.2.1 The Calorimeter Trigger

The ATLAS Level-1 Calorimeter Trigger requires analogue signals with a coarse granularity in pseudorapidity and azimuth (η, ϕ) as input from the calorimetry. For a total $\eta - \phi$ coverage of $|\eta| < 4.9$ and $0 < \phi < 2\pi$, calorimeter signals are summed together to form trigger-towers with a granularity of 0.1×0.1 in η and ϕ. This summation is done separately for signals from the electromagnetic and hadronic calorimeters. All in all about 7200 trigger-tower signals are transmitted with a cable length of up to 60 m to the trigger electronics, which are located near the detector. The chosen granularity is a balance between rejection of background and complexity of the trigger processor. The trigger is made up of: the Pre-Processor (PPr), the Cluster Processor (CP) and the Jet/Energy-Sum Processor (JEP).

The interconnections between the three processors and the input signals from the calorimeter are shown in Figure 2.5. The Pre-Processor performs several functions on each of the analogue input signals. It provides digital data for the CP and JEP processor, which work in parallel to each other. The digital data from the PPr contains the deposited transverse energy identified with a corresponding bunch-crossing in time (BCID). For the CP the preprocessing is done for all trigger-tower signals up to $|\eta| < 2.5$ and for the JEP up to $|\eta| < 4.9$. In case of the JEP the trigger-towers are summed further to form jet elements with a coarser size of 0.2×0.2. The summation is still treated separately for electromagnetic and hadronic signals. The CP performs the electron/photon and hadron/tau algorithms, whereas the JEP performs the jet algorithm and calculates the global sums for the E_T-miss and sum-E_T triggers. Physically, these processors are composed of printed-circuit modules mounted in crates.

In order to optimise the signal fan-out between crates and modules the trigger tower matrix is divided into quadrants in ϕ, see Figure 2.6. Each quadrant is then

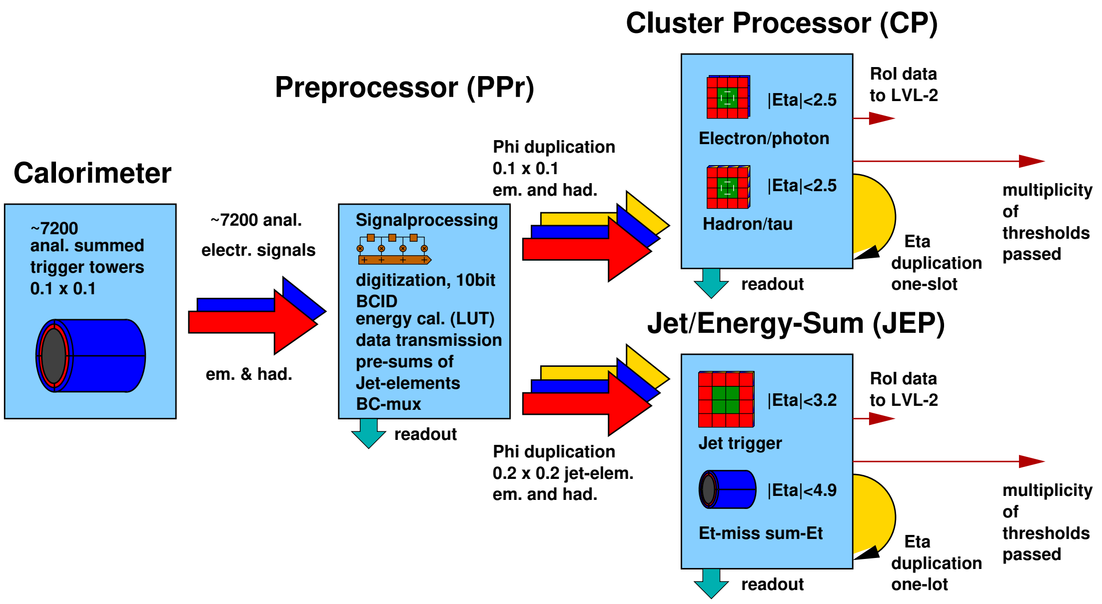

Figure 2.5: Calorimeter Trigger Architecture.

mapped into two PPr crates, one CP crate and one JEP crate. Signals at the boundary between each quadrant are duplicated at the Pre-Processor in such a way that shared signals, needed by 'stepped-window' algorithms from neighbouring quadrants, are received via the Pre-Processor rather than by inter-crate links. This has the advantage that signals shared between modules in the η-direction are transmitted via 'one-slot' backplane connections only. Inter-crate links have to be avoided in a pipelined system where a minimum of trigger latency is a crucial design aspect.

Calorimeter Trigger algorithms

The Level-1 Calorimeter Trigger performs four basic algorithms: an electron/photon, a hadron/tau, a jet and an E_T-miss/sum-E_T trigger. Each algorithm identifies objects according to their transverse energy. The number of objects passing a programmable E_T-threshold is counted as multiplicity. The multiplicity of objects which have passed is sent for eight different thresholds to the Central Trigger Processor (CTP). The regions-of-interest and the coordinates of the identified objects are sent to the Level-2 Trigger. The algorithms are based on 'stepped windows' in the $\eta - \phi$ trigger-tower space. Figure 2.7 illustrates the basic elements to construct the electron/photon and hadron/tau trigger algorithms implemented inside the Cluster Processor. The basic elements for the Jet/Energy-sum Processor are shown in Figure 2.8 for the jet and the E_T-miss/sum-E_T algorithms.

To cover the whole trigger region in parallel the algorithms are applied to trigger windows stepped by one trigger element at the same time. For each of these

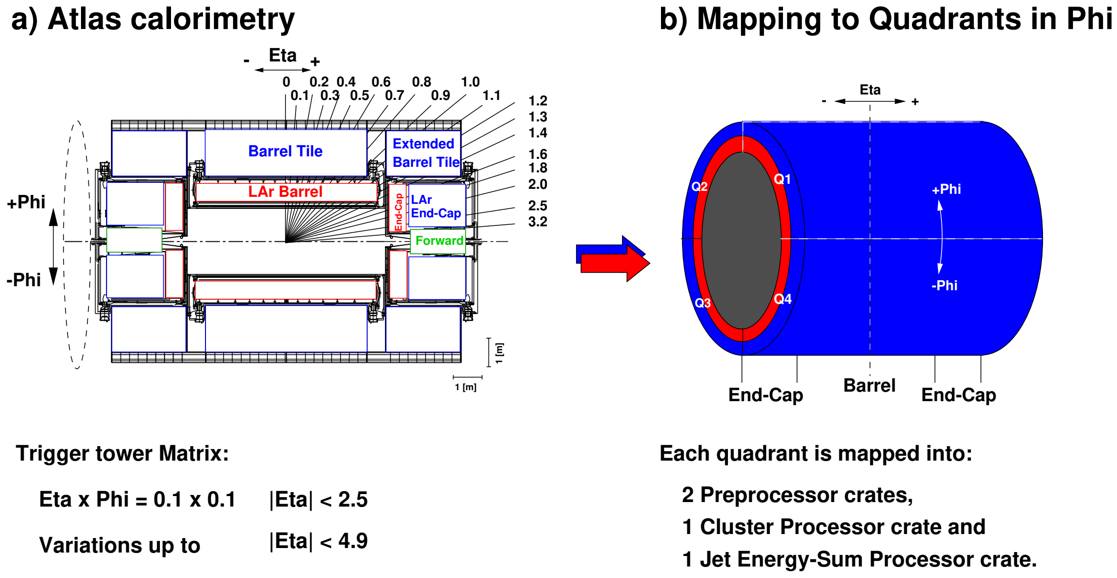

Figure 2.6: Calorimeter Trigger tower mapping. The ATLAS calorimetry is shown in (a) and the trigger tower mapping into quadrants in ϕ is illustrated in (b).

windows the algorithms decide whether an object passing the threshold conditions was found. Because of the stepping through the $\eta-\phi$ coverage, the algorithms must be resistant against double-counting of objects, which is referred to as declustering.

The first three algorithms search for different objects, but the basic computing steps are the same for all of them. The first two algorithms are performed within a window of $4 \times 4 \circ 2$ $(\Delta\eta \times \Delta\phi)$ trigger towers. The $\circ$ stands for separate processing of the electromagnetic and hadronic layers. In the case of the jet algorithm the window size is programmable $(4\times 4,\ 3\times 3,\ 2\times 2)\bullet 2$, where a $\bullet$ stands for summation of electromagnetic and hadronic layers. See Table 2.1 for a complete summary of the element sizes of the electron/photon and hadron/tau trigger and Table 2.2 for the jet and the E_T-miss/sum-E_T trigger. The sum-E_T trigger adds up all trigger elements up to $|\eta| = 4.9$ and the E_T-miss trigger does the same for the components E_x and E_y separately.

Except for the relevant distinctions given in Table 2.1 and Table 2.2 the processing of the first three algorithms can be summarised as follows:

- Each algorithm identifies a RoI cluster which has a local E_T maximum inside the trigger window (declustering). To pass this condition, the RoI cluster must be more energetic than the neighbouring RoI clusters to the right and to the one above, and at least as energetic as the nearest neighbours to the left and the one underneath.

- Inside the RoI cluster the energy of all possible E_T-measure clusters are compared with eight thresholds.

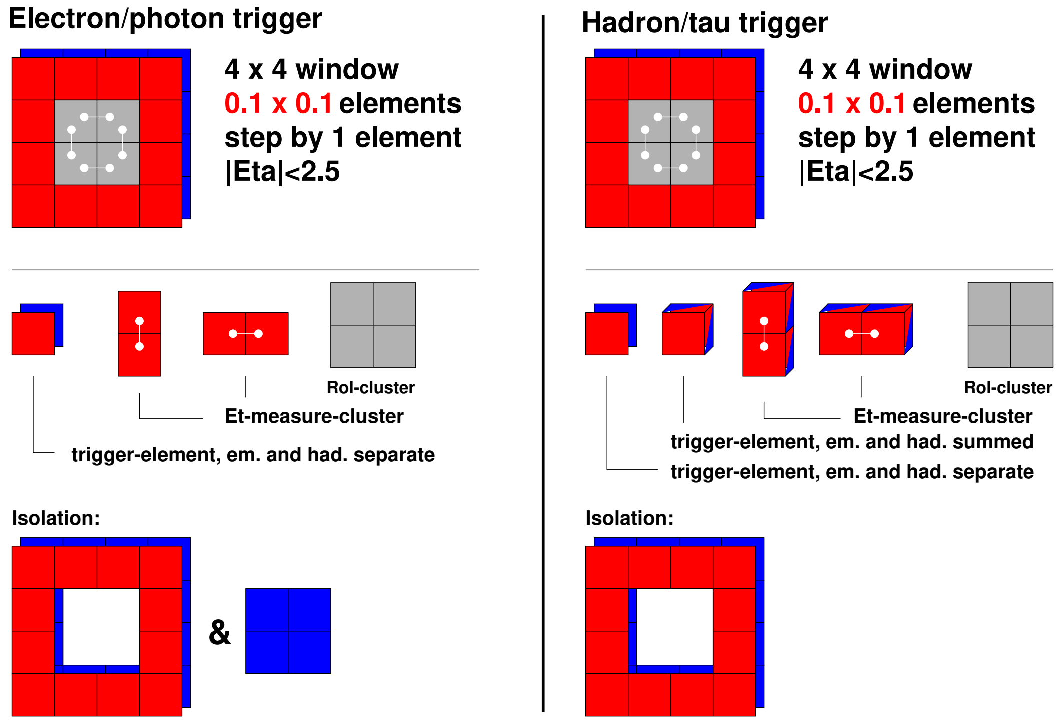

Figure 2.7: Cluster Processor algorithms.

- Trigger elements inside an isolation region are summed together and compared with an isolation threshold. This is not used in case of the jet algorithm.

- An object is identified if it has passed all threshold conditions. Then the multiplicity for the corresponding threshold is increased by one. The position of the RoI cluster is sent as a region-of-interest to the Level-2 Trigger.

Module mapping

As mentioned before, one quadrant in ϕ is mapped into crates of processors. Each processor crate contains electronics modules and components, which perform preprocessing or calculating of the actual stepped algorithms in parallel for a local area in $\eta - \phi$. Figure 2.9 illustrates the mapping of the $\eta - \phi$ trigger-tower matrix into the processor modules. As an example, it shows the mapping of one slice out of quadrant 3. The duplication of signals from the Pre-Processor and the duplication from adjacent modules is shown around a local area. A local area is fully processed on each module and does not require any fanout.

Table 2.3 summarises the occupied local areas for each of the electronics modules (PPM, CPM, JPM). In the following sections the focus is placed on the components and technologies which are used inside these processor modules.

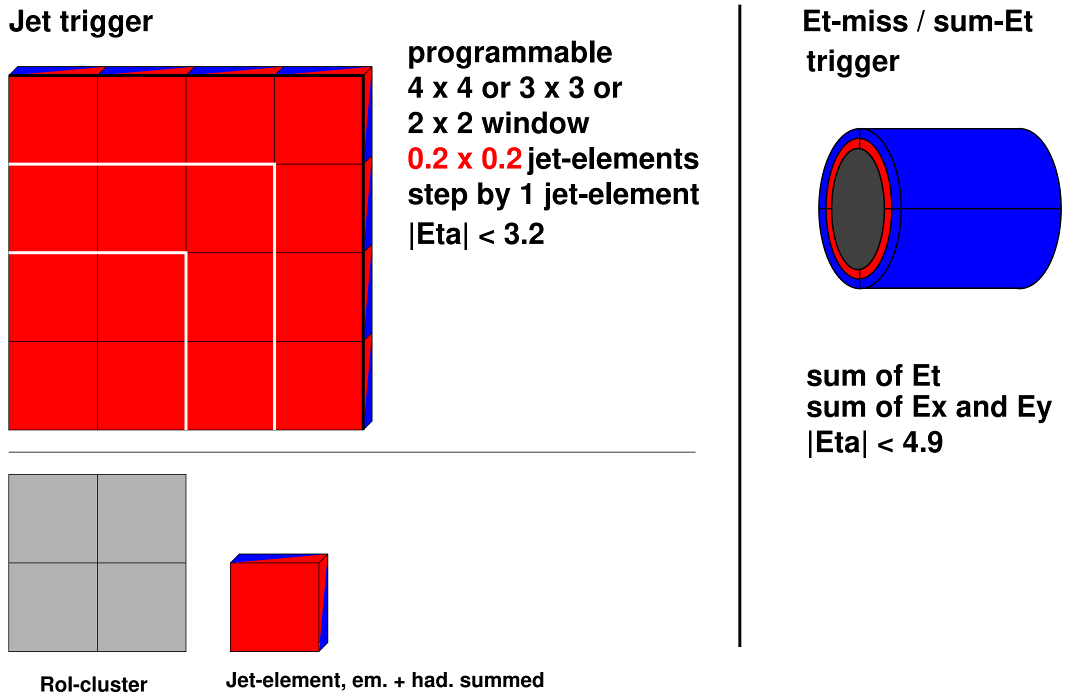

Figure 2.8: Jet/Energy-sum Processor algorithms.

Pre-Processor Module

The Pre-Processor Module (PPM) (see Figure 2.10) has 64 analogue receivers to accept the trigger tower signals from the calorimeter. The signals are processed by commercial integrated circuits (ICs) and application specific ICs (ASICs), most of which are located on 16 Multi-Chip Modules (PPr-MCMs). A MCM combines different IC technologies, e.g. analogue and digital components to increase the die to package-area ratio. Inside each PPr-MCM, four ADCs digitise the analogue signals to 10-bits. Two Pre-Processor ASICs (PPrAsics) perform the BCID algorithm, the transverse-energy calibration to 8-bits in a look-up table, and the pre-summing of jet elements. It also contains pipeline memories to store trigger data up to 2 μs, until a level-1 accept signal arrives from the CTP initiating the readout. Codes for the detection of transmission errors are generated before trigger tower data are serialised to 800 Mbit/s using gigabit links manufactured by Hewlett Packard (G-links). LVDS links from National Semiconductor may be used instead. A bunch-crossing multiplexing scheme, which doubles the effective bandwidth of the high-speed serial link by a factor of two is used for the transmission of preprocessed trigger towers to the CP. Pre-summed jet elements are serialised on the PPr-MCM and linked to the JEP with 9-bit resolution. The ϕ-duplication of links at quadrant boundaries is fanned out to cable drivers. The readout data from all PPrAsics are collected by one Readout Merger ASIC (RemAsic). This ASIC interfaces to a custom ring-like bus on the backplane (PipelineBus), which

size of	El./photon	Had./tau
Trigger element	$0.1 \times 0.1 \circ 2$	$0.1 \times 0.1 \circ 2$
Trigger window	$4 \times 4 \circ 2$	$4 \times 4 \circ 2$
Window step	0.1	0.1
RoI cluster	2×2	2×2
Isolation (em./had.)	(ring/square)	(ring/ring)
E_T-measure cluster	2×1 or 1×2	$1\times2\bullet2$ or $2\times1\bullet2$
η-coverage	$\lvert\eta\rvert$ <2.5	$\lvert\eta\rvert$ <2.5

Table 2.1: Comparison of trigger algorithms implemented inside the Cluster Processor. The × means $\Delta\eta \times \Delta\phi$ of the $\eta - \phi$ coverage or the number of elements in each direction. A $\bullet$ stands for summation of electromagnetic and hadronic layers and $\circ$ for separate processing.

size of	Jet	E_T-miss/sum
Trigger element	$0.2 \times 0.2 \bullet 2$	$0.2 \times 0.2 \bullet 2$
Trigger window	$(4 \times 4,\ 3 \times 3,\ 2 \times 2) \bullet 2$	–
Window step	0.2	–
RoI cluster	2×2	–
Isolation (em./had.)	–	–
E_T-measure cluster	trigger window	–
η-coverage	$\lvert\eta\rvert$ <3.2	$\lvert\eta\rvert$ <4.9

Table 2.2: Comparison of Trigger algorithms implemented in the Jet/Energy-sum processor. The × means $\Delta\eta \times \Delta\phi$ of the $\eta - \phi$ coverage or the number of elements in each direction. A $\bullet$ stands for summation of electromagnetic and hadronic layers.

shifts readout data via a readout driver board to the readout buffers. Slow control is used to set up the configuration and to load test data.

Cluster Processor Module

The Cluster Processor Module (CPM) receives a local area of 4×16 trigger towers plus ϕ-duplicated towers from the PPr (see Figure 2.11). A total number of 40 gigabit links per module is required to receive these signals. Two of these gigabit chips and two Serialising ASICs (SAsic) are placed inside a Multi-Chip Module (CPMCM). This requires 20 CPMCMs on one CPM. In the case of LVDS links, the number of receiver chips would be twice as much, but the reduced power density and the small package size of LVDS receiver chips may allow the use of a small daughter board instead of an MCM. The SAsic reduces the number of lines needed to fan-out trigger-tower signals to adjacent slots on the backplane by serialising the data from 40 MHz to 160 MHz. In order to avoid a trigger on faulty data the SAsic checks the data for transmission errors. The SAsic also provides a

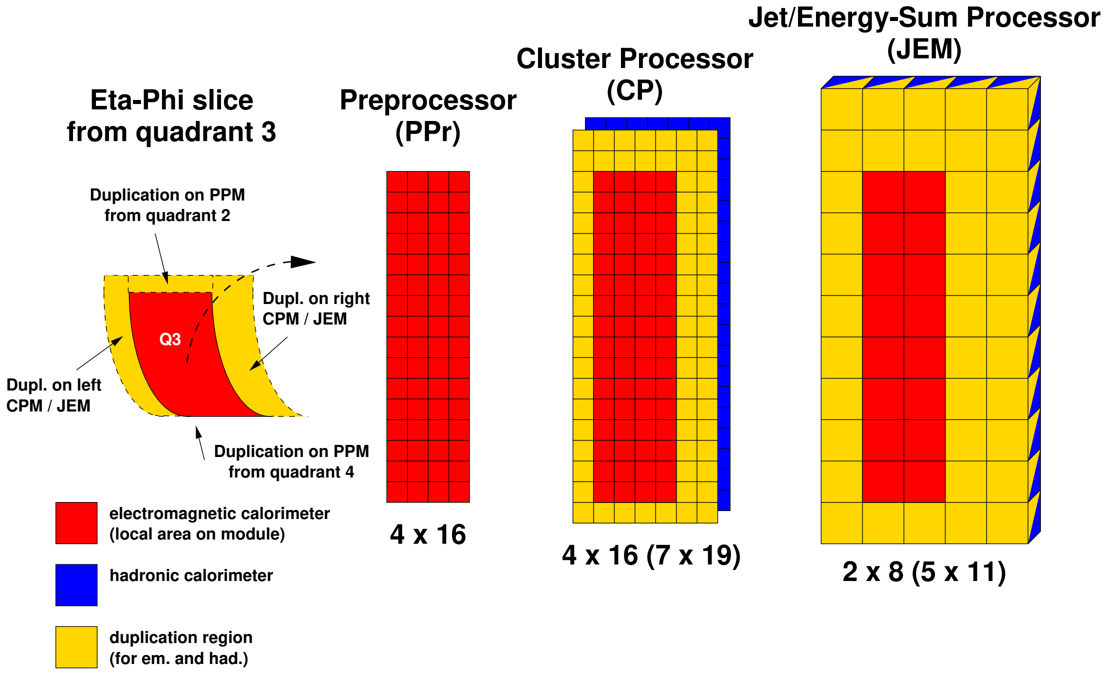

Figure 2.9: Mapping of $\eta - \phi$ trigger tower matrix into processor modules.

size of	PPM	CPM	JEM
element	0.1×0.1	$0.1 \times 0.1 \circ 2$	$0.2 \times 0.2 \bullet 2$
local area	4×16	$4 \times 16 \circ 2$	$2 \times 8 \bullet 2$
total area	4×16	$7 \times 19 \circ 2$	$5 \times 11 \bullet 2$
from neighbour	–	$3 \times 19 \circ 2$	$3 \times 11 \bullet 2$
to neighbour	–	$3 \times 19 \circ 2$	$3 \times 11 \bullet 2$

Table 2.3: Mapping of calorimeter regions into processor modules. The meaning of $\times$, $\circ$ and $\bullet$ corresponds to Table 2.1.

synchronised clock for the Cluster Processor ASIC (CPAsic). If it turns out that available FPGA devices can run at a speed of 160 MHz, it may be possible that the SAsic will be replaced by an FPGA. Eight CPAsics on a module perform the actual CP algorithms. Besides local signals, which are private on each CPM, the CP algorithms require duplicated signals from two sources: One source are the SAsics on the same CPM board providing ϕ-duplicated signals from the PPr. The other source are neighbouring CPMs which provide η-duplicated towers via one-slot backplane connections at 160 MHz. The multiplicity of objects which have passed a configurable set of thresholds are merged by Cluster Merger Modules (CMMs). The merged multiplicity results are sent from the CMMs to the CTP. RoI data are sent in parallel to the level-2 trigger and to the readout buffers. RoI data are collected via gigabit serial links and then merged into S-links, which represent a CERN standard link for the ATLAS data acquisition [Sli]. Slow control is used to

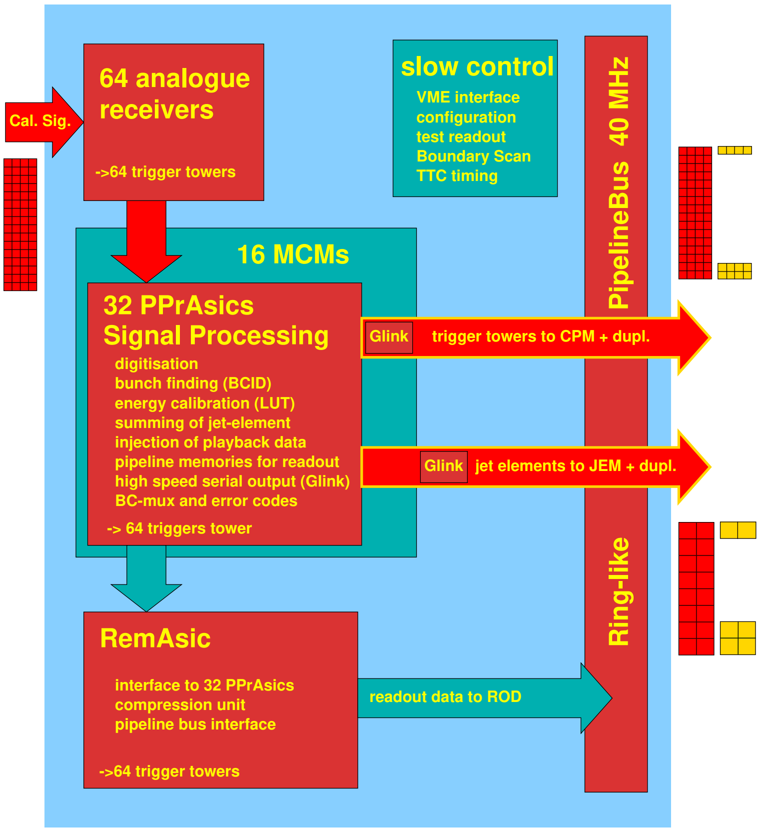

Figure 2.10: Pre-Processor Module (PPM).

set up configuration data, e.g. thresholds, timing parameters, and test data.

Jet/Energy-sum module

On a Jet/Energy-sum Module (JEM), 22 gigabit receivers convert the serial data from the PPr back to 40 MHz parallel data (see Figure 2.12). Compared to the CPM, the data volume is smaller because of the larger jet-element size used for the jet and E_T-miss/sum-E_T algorithms (see Table 2.3). Therefore, serialising the fanned-out backplane data to twice the clock speed (80 MHz) is sufficient. The serialising is performed by field programmable gate arrays (Sum-FPGAs). A FPGA device at 80 MHz is preferred to ASICs because of its programmability, an easier and faster design flow, and because FPGA prices are competitive with ASICs at low quantities. The Sum-FPGA sums electromagnetic and hadronic layers together to form the final jet elements. The energy sum and the components in E_x and E_y are computed for the local area. These sub-sums are merged inside a Sum Merger Module (SMM) to global values. Readout data from the Sum-FPGA are collected and transmitted to the readout-driver (ROD). The ROD makes use of the S-link standard to transmit the readout data to the readout buffers. Jet algorithms are

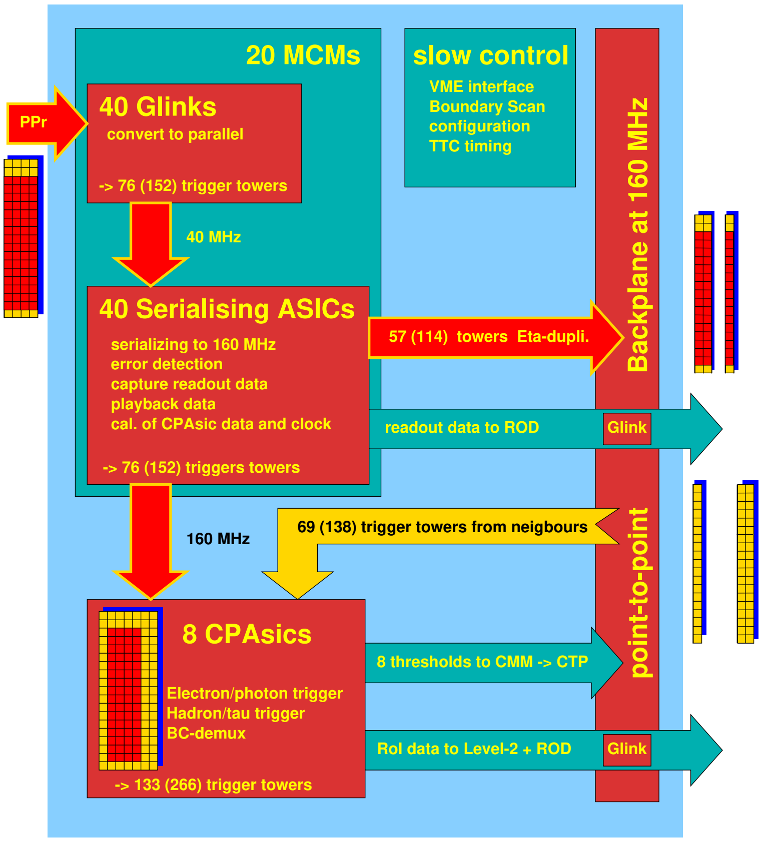

Figure 2.11: Cluster Processor Module (CPM).

also implemented inside an FPGA (Jet-FPGA). This offers the flexibility to change the algorithm depending on the size of the actual chosen trigger window. A total number of 22 jet elements from the Sum-FPGA and 33 jet elements from the neighbours are used to perform the jet algorithms. Finally, a Jet Merger Module (JMM) counts the multiplicity and RoI data from all modules and it finally sends the merged data results to the CTP.

2.2.2 The Muon Trigger

The Level-1 Muon Trigger is based on fast and finely-segmented muon trigger chambers, see Figure 2.13. In the barrel region Resistive-Plate Chambers (RPCs) are used to cover the range $|\eta| < 1.05$. The RPCs are wireless strip detectors in η and ϕ, which are easy to build to cover large areas. Three trigger stations (RPC1–RPC3) are used with a total number of 430,000 trigger channels. In the end-cap region Thin-Gap Chambers (TGCs) cover the region between $1.05 < |\eta| < 2.4$. At the end-caps a finer granularity is required because the main trigger stations are located outside the toroidal field, with smaller gaps between stations. The TGCs

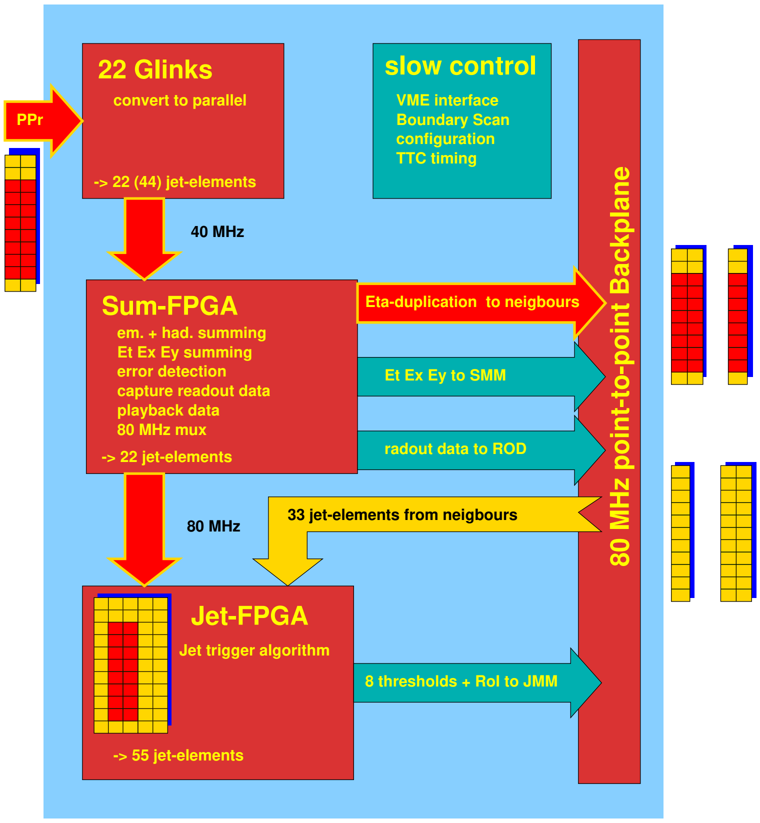

Figure 2.12: Jet/Energy-sum Module (JEM).

use strips in ϕ and wires in radial direction. All in all three stations (TGC1–TGC3) plus two smaller TGCs close to the hadronic end-caps are used. Because of the high background rate in forward direction, a high rate capability is needed. The total number of trigger channels for the end-cap region is about 370,000.

The Muon Trigger measures the muon trajectories in stations for two transverse momentum ranges (low-p_t and high-p_t). For the high-p_t range a coincidence in three stations is required. For the low-p_t range a 2-fold or a 3-fold coincidence can be used. What will be used depends on the actual low-p_t background. For the low-p_t range the thresholds are between 6–10 GeV and for the high-p_t range between 8–35 GeV.

Due to the high number of trigger channels, the trigger electronics is split up between electronics located on the detector and electronics in the trigger cavern. This reduces the number of optical links between detector and cavern to about 800 for the RPCs and to about 600 for the TGCs. The architecture of the Muon Trigger is shown in Figure 2.14. For the RPCs a Coincidence Matrix ASIC performs most of the processing. It finds the muon tracks separately for η and ϕ. In case of the TGC electronics the coincidence logic is partly performed on the detector using Slave Boards and in sector logic boards off the detector. In both cases the

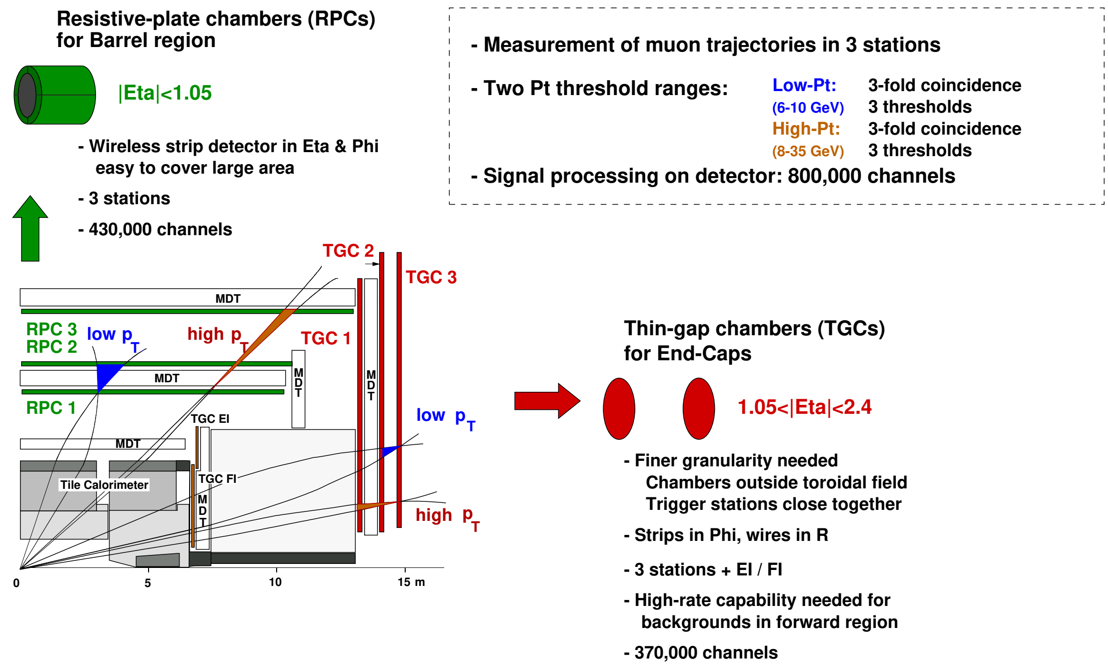

Figure 2.13: Muon trigger chambers.

sector logic reduces the granularity by keeping the highest two p_t candidates for each sector. The Muon to Central Trigger Processor Interface (MUCTPI) combines sector results for 6 thresholds and calculates the multiplicity for each threshold. It also performs declustering of overlap regions. Finally, it sends the multiplicity of muon candidates for 6 thresholds to the Central Trigger Processor and transmits the RoI data to the Level-2 Trigger.

2.2.3 The Central Trigger Processor

The Central Trigger Processor combines the multiplicity of thresholds in order to generate the final level-1 decision. It uses six types of trigger 'inputs'. A group of eight thresholds with 3-bits each are used to count the multiplicity of the electron/photon, hadron/tau, and the jet algorithm. The E_t-miss/sum-E_t trigger has eight thresholds for the missing transverse energy and four for the sum E_t. From the muon trigger six thresholds with three bits for the multiplicity are received. In addition 13 input bits are used for calibration and test triggers. All in all the total number of input bits to the CTP is 115 plus a contingency of 13-bits. A block diagram of the CTP is given in Figure 2.15.

In contrast to the other level-1 trigger electronics systems, the CTP is rather small. It fits on a single printed circuit board. It consists of synchronisation FIFOs used for timing alignment and synchronisation to the bunch-crossing clock. Inside look-up tables (LUTs) trigger 'objects' are formed from 128 input bits. Examples for

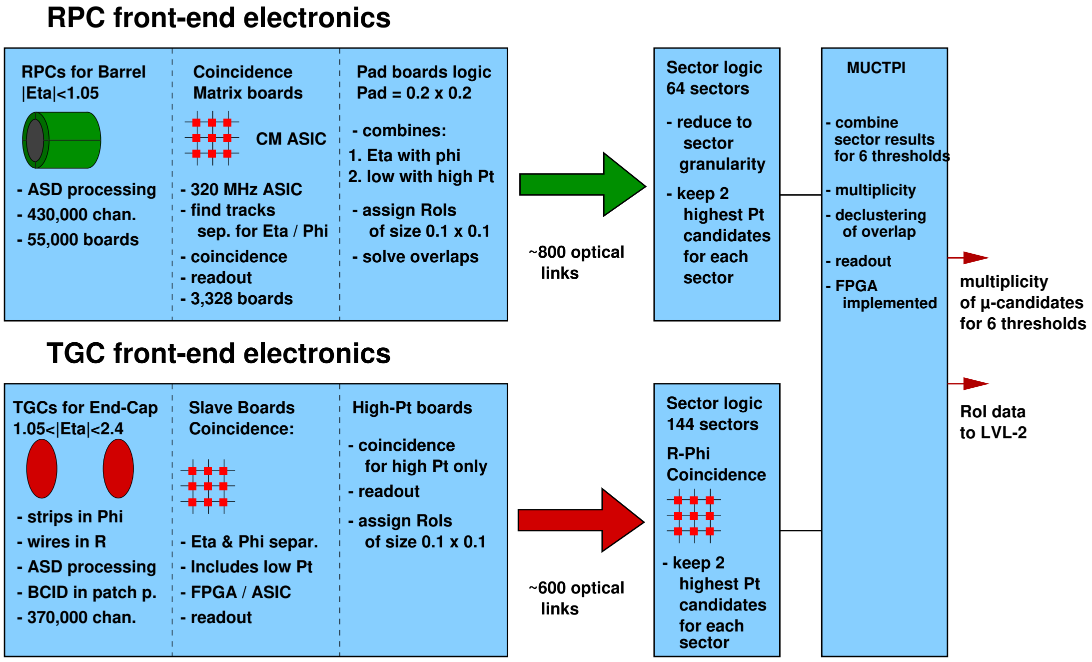

Figure 2.14: Muon Trigger Architecture.

trigger 'objects' are: two or more muons with an energy larger than 10 GeV or an electromagnetic cluster with more than 15 GeV. From these objects 96 trigger 'items' are formed by the use of a logical combination logic programmed into CPLDs[6]. Single items are given just by the objects themself whereas complex items are a combination of objects. The CTP provides the possibility to have triggers with low or high priority. Each trigger object can be gated and triggers can be prescaled. In addition to the Level-1 accept signal, the CTP provides trigger type information (8-bits), a BCID number (12-bits), and an event identification number (EVID, 24-bits). Those results are sent to the front-end electronics via the TTC system. The CTP is responsible for the dead-time handling of the Level-1 Trigger. It allows setting 0–16 'dead' bunch-crossings after the CTP has said 'yes'. For this time duration no Level-1 accept signal will be generated. In normal operation 4 BCs are used as default dead-time settings. The CTP also adjusts the average Level-1 rate by allowing only 1–32 triggers in 0-1.7 ms. The normal number would be to allow 8 triggers within 80 μs, which corresponds to 100 kHz Level-1 accept rate. The latency occupied by the CTP electronics is less than 100 ns (4 BCs).

2.2.4 Summary

The Level-1 Trigger system architecture is divided into three subsystems, the Calorimeter Trigger, the Muon Trigger and the Central Trigger Processor. The

[6]CPLD: Complex Programmable Logic Device

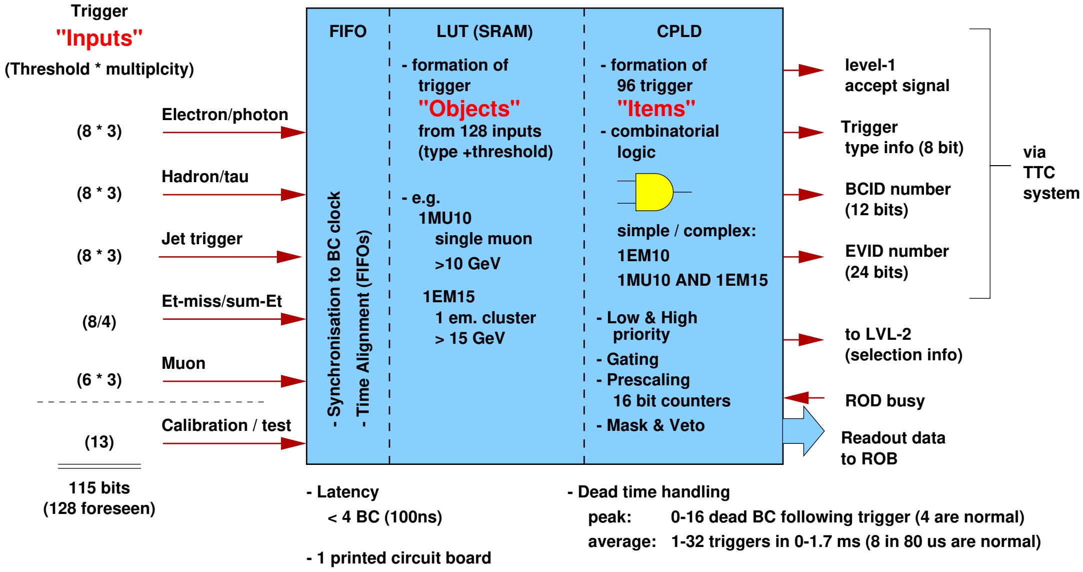

Figure 2.15: Central Trigger Processor Architecture.

parallel-processing calorimeter trigger algorithms are based on overlapping windows implemented in a ϕ-quadrant architecture. This architecture has been made possible because preprocessed trigger tower signals are duplicated at the Pre-Processor for quadrant boundaries, and jet elements are presummed for electromagnetic and hadronic signals (0.2×0.2). This minimises the data sharing between electronics modules and trigger processors. It simplifies the design of transmission-line backplanes and the layout of modules in terms of system size and costs. However, it requires the use of advanced technologies to process the large number of input signals within the total Level-1 Trigger latency of 2.0 μs. Figure 2.16 shows a summary of the overall Level-1 Trigger latency, which is an important constraint for the system design. The latency is shared between the subtriggers and the individual processors. The time measurement for the latency starts at the time of interaction, including the time-of-flight from the interaction point to the sub-detectors. The measurement includes all cable delays from the detector to the trigger electronics, and the cables delays required to distribute the level-1 accept signal back to the front-end electronics.

Table 2.4 gives a listing of some parameters which are of importance for the Calorimeter and Muon trigger system designs.

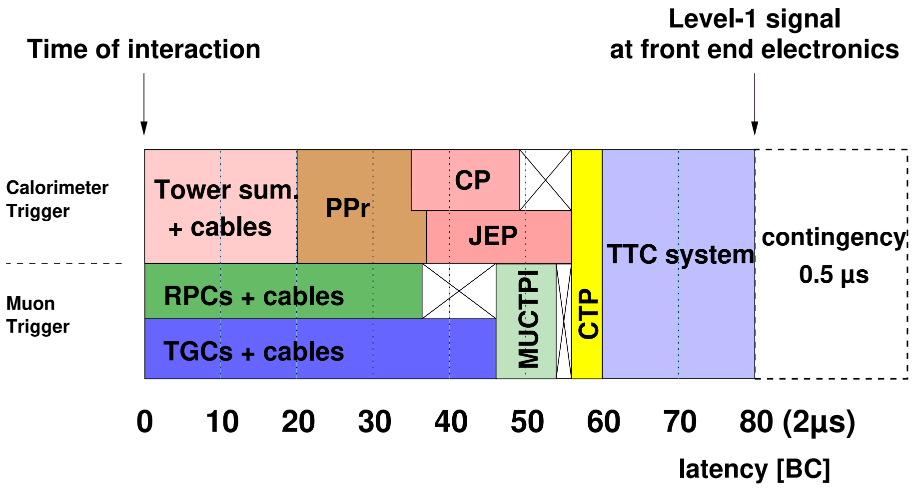

Figure 2.16: Attribution of the Level-1 Calorimeter Trigger latency to the individual sub-triggers and their associated processors.

Level-1 Trigger parameters	
Average trigger rate	75 kHz (100 kHz max.)
Max. peak trigger rate	8 MHz
Latency	2.0 μs (80 BCs)
Latency contingency	0.5 μs (20 BCs)
Trigger 'objects'	96
Calorimeter Trigger	
Trigger tower granularity	0.1×0.1
No. of trigger tower signals	~ 7200
η coverage	$\|\eta\| < 4.9$
Muon Trigger	
No. of optical detector links	~ 1400
No. of trigger chamber signals	$\sim 800,000$
η coverage	$\|\eta\| < 2.4$

Table 2.4: Level-1 Trigger parameters.

Chapter 3

The Level-1 Calorimeter Trigger Pre-Processor

- System overview
- Tasks of the Pre-Processor
- Key components
- Pre-Processor MCM

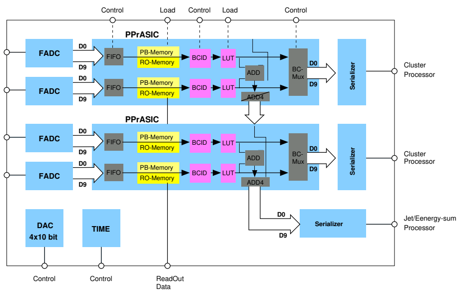

3.1 System overview

An overview of the Level-1 Calorimeter Trigger architecture was given in detail in Section 2.2.1. There, the arrangement and connection between the individual Calorimeter Trigger processors, the Pre-Processor, the Cluster Processor, and the Jet/Energy-sum Processor were described. The trigger space ($\eta \times \phi$) was economically mapped as ϕ-quadrants into electronics modules, in order to minimise the signal duplication at the Pre-Processor and to simplify fanout in the Cluster and Jet/Energy-sum Processors. If it turns out that the density target required for a ϕ-quadrant architecture can not be met in the Calorimeter Trigger, the system can be re-partitioned into ϕ-octants by dividing each quadrant in half. Beside additional electronics modules, crates, and racks, a consequence of a ϕ-octant architecture would be that the amount of duplicated signals at ϕ-edges of a module would be doubled. Hence, the ϕ-quadrant architecture for the Calorimeter Trigger is a major design target. A compact 64-channel Pre-Processor Module with a high degree of component integration is required to preprocess all trigger tower signals economically in four electronics crates.

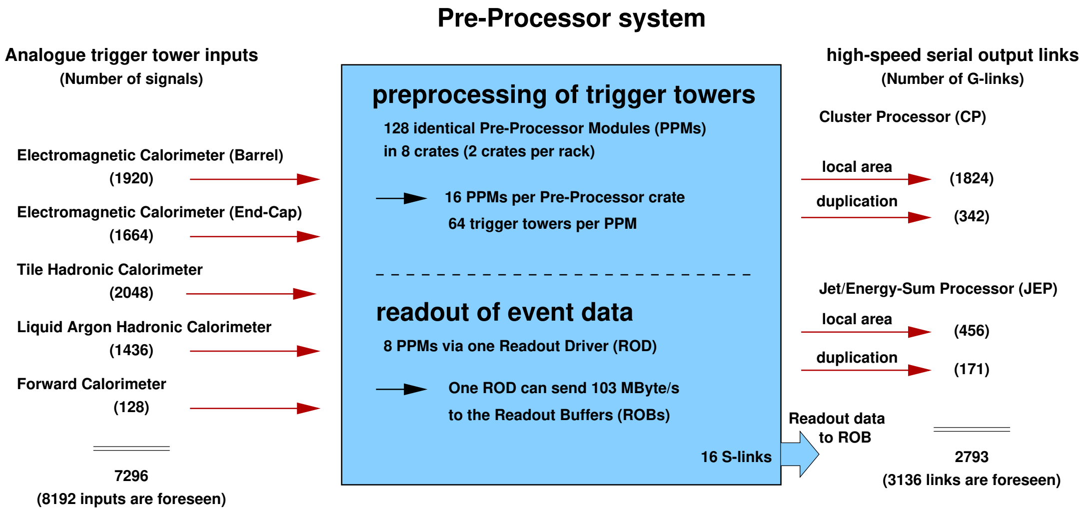

Figure 3.1: Pre-Processor system overview.

On the basis of a 64-channel Pre-Processor Module, the number of modules, the number of input signals, and the number of serial output links are shown in Figure 3.1. All in all the Pre-Processor receives 7296 analogue input signals from the ATLAS calorimetry. These signals are converted to transverse energy, within a $\pm 10\,\%$ band. For the liquid-argon calorimeters and similar for the Tile Calorimeter, this is done in 'Receiver Stations' in front of the Pre-Processor. The total number of serial output links is 2793, including the duplication at ϕ-boundaries. This number would be doubled if LVDS links would be used instead of G-links.

The tasks that the Pre-Processor system has to perform, based on its input signals can be summarised as follows:

- **Preprocessing:** Provide the trigger processors downstream with digital data containing the transverse energy deposited, identified with the corresponding bunch-crossing. For the Cluster Processor (CP) the granularity is 0.1×0.1 for $|\eta| < 2.5$ and for the Jet/Energy-Sum Processor (JEP) the granularity is 0.2×0.2 for $|\eta| < 4.9$. In both cases the input data are separate for the electromagnetic and the hadronic calorimeters.

- **Readout of event data:** Raw trigger data from the Pre-Processor are needed to be able to tell what has caused a trigger and to allow monitoring of the performance of the trigger system.

The following section explains each of these tasks in detail. Then the key components used for the implementation of the Pre-Processor system are described.

3.2 Tasks of the Pre-Processor

The Pre-Processor is of importance for the running of the ATLAS experiment, because all the Level-1 Calorimeter Trigger input data have to go through it. In case that the Pre-Processor needs to be repaired, a quick exchange of broken modules is required. Therefore, it is not desirable to have different types of modules implying a large set of different spare modules. Hence, all the Pre-Processor Modules will be identical, having the same functionality.

Since the preprocessing is mainly channel oriented, it is sufficient to explain the stages of preprocessing for one channel only.

3.2.1 Preprocessing tasks

The preprocessing tasks are illustrated in Figure 3.2 for the processing of one trigger tower signal. The preprocessing steps are marked as **A** to **H** and are described as follows:

A: **Reception of analogue trigger tower signals:** The differential analogue trigger tower signals are received by a differential line receiver circuit. The input voltage range is linearly mapped, with 0–2.5 V representing 0–250 GeV. A programmable DAC with 10-bit resolution is used to adjust the zero baseline for each input signal. For a description of the signal shape see Section 4.1.2.

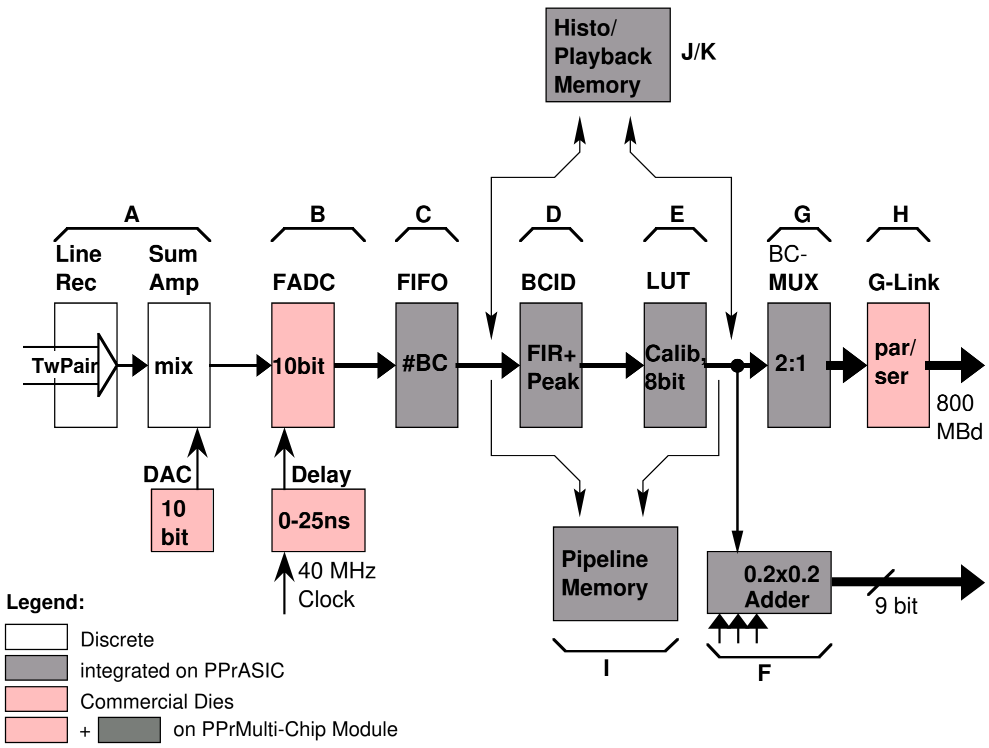

Figure 3.2: Preprocessing of one trigger tower signal by the Pre-Processor [TDR98].

B: **Digitisation and phase adjustment:** Each analogue input signal will be digitised by a flash analogue-to-digital converter (FADC) with 10-bit resolution. The time position of the sampling strobe with respect to the analogue input signal can be adjusted in steps of 1 ns within a range of 25 ns. This is required to perform the fine synchronisation of each trigger tower signal and to sample each pulse at its maximum. An ASIC (Phos4) developed by the CERN Microelectronics group will be used for this time adjustment [Phos4].

C: **Synchronisation**: The digitised data needs to be synchronised to the same bunch-crossing, because of different time-of-flight for particles from the interaction point to the calorimeter and the different cable lengths from the calorimeter to the trigger cavern (USA15). Synchronisation is done in steps of 25 ns by a FIFO with a programmable depth of 16 bunch-crossings. Assuming a cable propagation delay of 5 ns/m, this corresponds to the delay of an 80 m long cable. This includes enough contingency because the actual cables are not going to be longer than 60 m.

D: **Bunch-crossing identification (BCID):** This circuit consists of two algorithms to identify the transverse energy deposition represented by a trigger tower signal, and the corresponding bunch-crossing in time. One algorithm is applied to non-saturated signals and one is applied to saturated signals. For a detailed description and simulation results of these algorithms see Chapter 4.

E: **Lookup table:** A lookup table is used to fine-calibrate the digitised data to the deposited transverse energy E_T. It maps the 10-bit data after BCID to 8-bit, with a least significant bit (LSB) of 1 GeV. In addition, it can be used to subtract a pedestal and it can apply a minimum threshold to suppress noise.

F: **Formation of jet elements:** The Pre-Processor pre-sums four 8-bit trigger towers to coarser jet elements with a size of 0.2×0.2. The summing is done separately for the electromagnetic and the hadronic calorimetry, leaving the option to apply separate jet thresholds for electromagnetic and hadronic clusters in the Jet/Energy-Sum Processor. The summing tree in the Pre-Processor requires all four inputs for a jet element to be in adjacent trigger tower channels. The resolution for the jet elements is reduced to 9-bit accuracy, with a least count of 1 GeV, before transmission.

G: **Bunch-crossing multiplexing:** This transmission scheme (BC-mux) doubles the effective bandwidth of the serial links to the Cluster Processor. In case of a G-link transmitter/receiver chip-set the number of links is only 2166 instead of 4332, with each G-link carrying four trigger-tower signals. This BC-mux scheme can not be used for the transmission of jet elements to the Jet/Energy-Sum Processor because the presumming removes empty bunch-crossings. For a detailed description of the BC-mux scheme see Section 4.5.

H: **Serial data transmission:** Preprocessing results are sent to the downstream processors via high-speed serial links. A high-speed serial transmission is required to keep the number of data links to an acceptable value. The feasibility of a serial data rate of 800 MBd was demonstrated using the G-link chip-set, see Chapter 7. Because of the high power dissipation of G-links, LVDS links will probably be used for the final system.

3.2.2 Readout tasks

The Pre-Processor provides pipelined readout of raw trigger input data as well as E_T values after the lookup table in order to tell what has caused a trigger and to provide diagnostic information. It allows the monitoring of the performance of the trigger system and the injection of test data for trigger system tests. These tasks are marked as **I** to **K** in Figure 3.2 and are described as follows:

I: **Pipelined readout:** The role of the pipelined readout of event data was described in the trigger system overview Section 2.1. The function of the readout pipelines in the Pre-Processor is equivalent, but independent of those of the detector readout. The Level-1 Trigger captures its own event data as soon as it has triggered. Two sets of pipeline memories capture event data in the Pre-Processor. One records the raw FADC data at the Pre-Processor input and one records after the lookup table and BCID. The number of time slices around an accepted event can be preset to read up to 128 time slices. Without introducing deadtime to the readout, the identified bunch-crossing and two time slices around can be read out. See Section 3.3.2 for a detailed description of the implementation of this in an ASIC.

J: **Data playback:** The Pre-Processor, comprising 7296 input channels, can inject data for technical tests of the Level-1 Trigger system. This allows testing the functioning of the trigger processors and the relative timing of the processors and the input channels.

K: **Histogramming:** This is a useful feature for monitoring of the trigger performance. Two modes of 'online' trigger monitoring are foreseen for each trigger tower input at the Pre-Processor. The first mode is *rate monitoring*, where entries in a histogram above a programable threshold are counted for a given time duration. The second mode is used for monitoring of the *transverse energy spectrum* of either the raw FADC data or the energy calibrated output after the lookup table and BCID. The latter mode allows the monitoring of a 'bunch window' out of the 2961 LHC bunches in one turn.

3.3 Key components of the Pre-Processor

The Pre-Processor comprises 128 identical Pre-Processor Modules (PPMs), which are hosted in 8 electronic crates. One Readout Driver (ROD) collects the readout data of 8 PPMs via a 'ring-like' bus (PipelineBus) on the backplane and sends the readout data via a S-link [Sli] to the Readout Buffers (ROBs). Besides 16 Readout Drivers and 8 VMEbus Crate Controllers, 128 Pre-Processor Modules represent the main electronics of the system. The smallest exchangeable electronics component on a PPM will be the Pre-Processor Multi-Chip Module (PPr-MCM). The PPr-MCM contains all the preprocessing stages of four trigger tower signals, which were marked in Figure 3.2, except for the analogue line receiver circuit, marked as

A. The digital signal processing, which is shaded in dark grey, is contained in a Pre-Processor ASIC (PPrAsic).

A 64-channel Pre-Processor Module has to carry 16 PPr-MCMs. Input connectors for 64 differential analogue signals will be located at the front panel. On the backplane side of a module the PipelineBus connectors, the VMEbus connectors and the connectors for the high-speed serial links will be located.

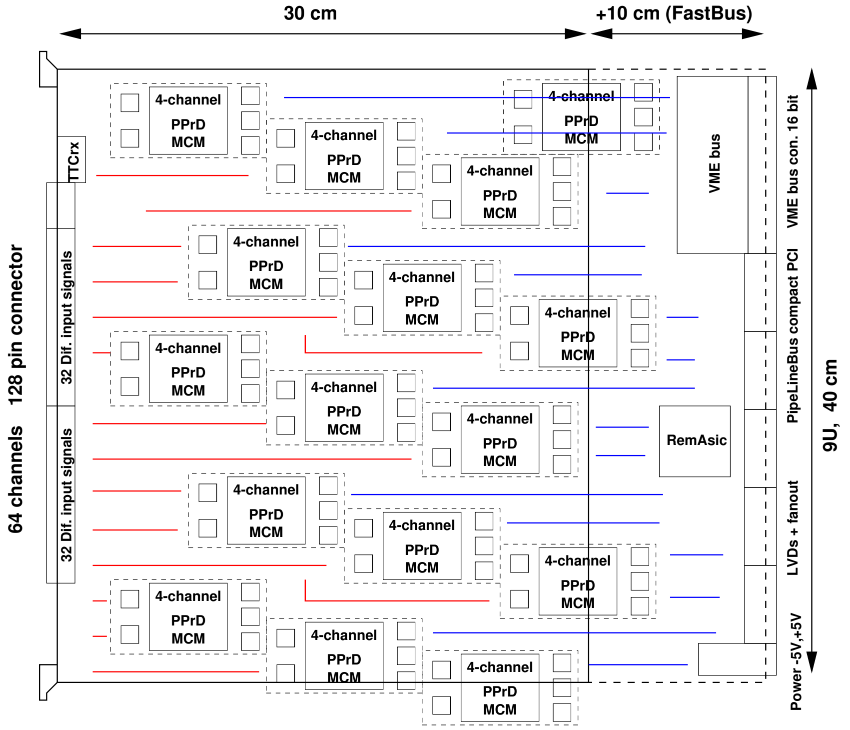

Figure 3.3: Component view of a Pre-Processor Module (on relative scale).

Figure 3.3 shows an estimate of the Pre-Processor Module board space. The module will have 9 units (9U = 40 cm) and 40 cm in depth. The form factor of the PPr-MCM is estimated to be the same as the demonstrator MCM. This is because the demonstrator MCM has similar functionality to the final PPr-MCM, see Chapter 5 for a detailed description. The position of the PPr-MCM on the Pre-Processor Module was optimised in terms of good heat exchange and a minimum of cross talk between the high-speed MCM output signals and the analogue inputs. The board space occupied by the differential line receivers, the clock buffers, and the serial line drivers needed for the high-speed signal fanout, is estimated around the PPr-MCM symbol. This figure makes it obvious that efficient signal routing and component placement will be of importance, and that the feasibility of 64 channels per module relies on the Pre-Processor Multi-Chip Module being compact in size.

On each Pre-Processor Module a Readout Merger ASIC (RemAsic) will collect the readout data from all the on-board PPr-MCMs. This ASIC has a 70 pin (2×35 bits) connection to the PipelineBus at the backplane. In the following section the three key components of a Pre-Processor Module will be described, which are the Pre-Processor MCM, the Pre-Processor ASIC, and the Readout Merger ASIC.

3.3.1 The Pre-Processor Multi-Chip Module

A block diagram of the Pre-Processor Multi-Chip Module is shown in Figure 3.4. It will consist of 11 dies, it will be manufactured in a laminated MCM-L process, the same as that investigated for the demonstrator Multi-Chip Module. For details of the MCM-L production process see Chapter 5.

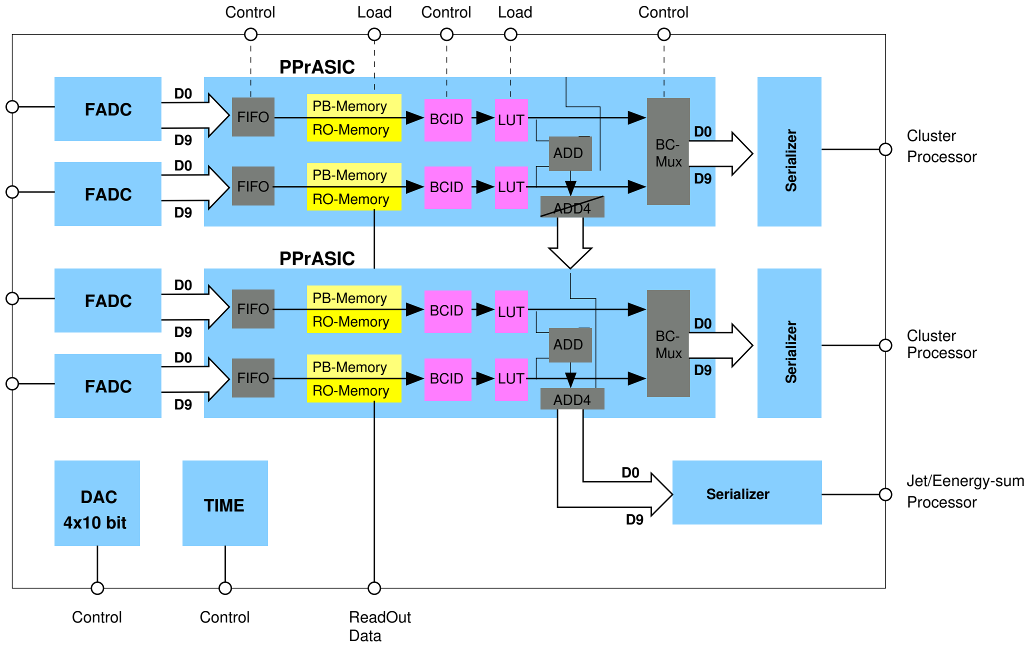

Figure 3.4: Block diagram of Pre-Processor Multi-Chip Module.

The bock diagram of the PPr-MCM contains four 10-bit FADCs, two dual channel Pre-Processor ASICs, and serialisers for the digital data transmission to the processors. Three LVDS chips may be used instead of one G-link per MCM. The jet-element data transmission is done externaly in case of G-links. A timer chip (Phos4) is needed for the phase adjustment of the FADC strobes with respect to the analogue input signals, and a four channel digital-to-analogue converter (DAC) is required for the analogue baseline adjustment of the four inputs. If it turns out to be a problem to use LVDS transmitter/receiver chip-sets for the serial data transmission, G-links can be used instead. This was investigated and was shown to be feasible on the demonstrator MCM. The summing of jet elements are shown as cross connections between the two dual-channel PPrAsics. It might be possible to fit all four channels into one PPrAsic, which would reduce the number of dies per PPr-MCM to ten.

The design of the final PPr-MCM will benefit from the design experience established by the demonstrator project of the PPrD-MCM. The demonstrator MCM has a similar partition into dies and includes most of the preprocessing of four trigger tower signals. The main difference is that the demonstrator MCM digitises only at 8-bit precision, it uses a former Pre-Processor ASIC prototype (FeAsic), and it requires multiplexing of G-link input data and level-conversion to transmit four preprocessed trigger tower signals via one G-link device. The maximum serial data rate of a G-link device on the PPrD-MCM is 1600 MBd. The required serial data rate for the final PPr-MCM is reduced, due to the use of the bunch-crossing multiplexing scheme (BC-mux) described in Section 4.5, and the intention to use two LVDS serialisers, which share the data.

	PPrD-MCM	Final PPr-MCM	
Analogue dies	3	$\nearrow$ +2	5
Total number of dies	9	$\nearrow$ +2	11 (10)
MCM area	15.9 cm^2	$\rightarrow$ or $\searrow$	?
Serial data rate	800 or 1600 MBd	$\searrow$ -75 %	400 MBd
Power consumption	9 W	$\searrow$ -30 %	$\sim$6 W
Temperature rise	+20 C°	$\searrow$ -25 %	$\sim$+15 C°
Number of bonds	613	$\searrow$ -37 %	$\sim$386
SMD connector pins	120	$\searrow$ -30 %	$\sim$80
$MTBF_{MCM}$	100.000		111.111
Reliability per year	92.4 % (9 dies)	$\searrow$ -0.8 %	91.6 % (10 dies)

Table 3.1: Demanding parameters of the demonstrator MCM in comparison to the final PPr-MCM. The MCM temperature assumes air cooling. The reliability calculation and the $MTBF_{MCM}$ assumes a failure rate of one part per million hours.

Table 3.1 shows a comparison of the demonstrator MCM with the final PPr-MCM. A demanding design parameters is the number of dies per MCM, which will increase by one or two, depending on a dual or quad PPrAsic. The demonstrator MCM uses dual FADCs dies, which are not available for a 10-bit precision. The power dissipation will go down by roughly 30 %, due to the less power-hungry LVDS serialisers. The reduction in power gained for the MCM is less than one would expect because the power is taken over by the 10-bit FADCs, which are more power-hungry if a short conversion time is needed. For a conversion time less than three clock cycles (3 ticks latency), the power dissipation is roughly one Watt for a 10-bit ADC. Hence, the reduced power will reduce the MCM temperature only slightly. An improvement in terms of reliability is the reduced number of total bonds, even if the number of dies is increased. The reliability, calculated on the basis of die count, is slightly reduced by 0.8% for the inclusion of one more die on the MCM substrate. See Section 5.7.8 for the calculation of the reliability of a MCM.

All in all, the final MCM will be less demanding on MCM technology, which will be an advantage for the reliability of the Pre-Processor system.

3.3.2 The Pre-Processor ASIC

The Pre-Processor ASIC performs the digital processing tasks for two or if possible for four trigger towers. It will be completely described using Verilog as a Hardware Description Language (HDL) [Tho96]. A 0.6 μm CMOS[1] process offered by AMS (Austria Micro Systems) will most likely be used for manufacturing. A specification of the Pre-Processor ASIC created as part of the ATLAS Preliminary Design Review (PDR) can be found in [PPR99]. The following summarises the tasks which need to be performed by the PPrAsic. Most of them were already described in Section 3.2.1.

- Synchronise digital data from FADCs.
- Provide pipelined event data readout.
- Perform BCID for non-saturated and saturated calorimeter signals.
- Foresee an input bit from an optional external BCID algorithm.
- Provide a 'playback' mode for test data replay.
- Allow monitoring of rates for each trigger tower.
- Fine-calibration to the transverse energy in lookup tables.
- Double the effective transmission bandwidth of serial transmitters (BC-mux).
- Generate error bits for the serial transmission.
- Provide a read-back capability for set-up and control data.
- Monitor the MCM temperature.
- Provide a Test Access Port (TAP) for in-circuit testing and boundary scan (JTAG).

Constraints of the PPrAsic design are as follows:

- Reduce the occupied latency to a minimum.
- Keep the die size as small as possible.

[1]CMOS: Complementary Metal Oxide Semiconductor

- Keep power consumption as low as possible.
- Allow Flip-Chip bonding on an MCM-L substrate.

Figure 3.5 shows a block diagram of the PPrAsic for two trigger towers. For the MCM design, the PPrAsic layout should be optimised in terms of die size and pad layout. Flip-Chip mounting of the PPrAsic improves the reliability of the chip bonding and it reduces the occupied chip area on the MCM substrate. A Test Access Port (TAP) for in-circuit testing and boundary scan (JTAG) is essential for a chip which is intended to be embedded inside a Multi-Chip Module. The PPrAsic is located between the FADC and the serialisers, where it is not possible to test the connectivity of this ASIC. The ASIC is isolated from the MCM environment and a JTAG interface provides the only way to either preload defined input and output pad stages or to scan the input to this chip. The Pre-Processor system requires at least 2048 working MCMs, not including spares. Hence, an automated test procedure is needed to identify connection errors at an early stage of the MCM assembly.

3.3.3 The Readout Merger ASIC

The Readout Merger ASIC (RemAsic) plays a key role in the Pre-Processor. It has to perform two tasks for all 64 PPrAsics on a Pre-Processor Module: it has to collect all the event data, and it has to load all the configuration and control data.

The ASIC will have 16 ports to connect to all the serial interfaces of the PPrAsics on a module. In the case of a dual channel PPrAsic, two serial interfaces will be daisy-chained. The event data are stored as an event block in an internal buffer memory. On that data it performs a data compression algorithm (Huffman, Run-Length or a difference encoding), which achieves a typical data reduction factor of 2.5 [Nie98]. After compression the data are made available to the custom readout bus (PipelineBus). For configuration, the RemAsic receives data from the PipelineBus or from the local VMEbus interface and sends it to the PPrAsics via the serial interfaces.

A block diagram of a prototype RemAsic [Rem98] is shown in Figure 3.6. This ASIC was designed using Verilog HDL and manufactured using a 0.7 μm CMOS process from Atmel ES2. This prototype has a subset of four ports to connect to the serial interfaces of the PPrAsic prototypes (FeAsics). On each port four FeAsics can be used in a daisy-chained configuration. The final RemAsic will most likely be implemented in an FPGA because only 128 of these devices are required for the Pre-Processor system. In low quantities and at a speed of 40 MHz FPGAs are competitive with ASICs regarding price, even if a large design is required.

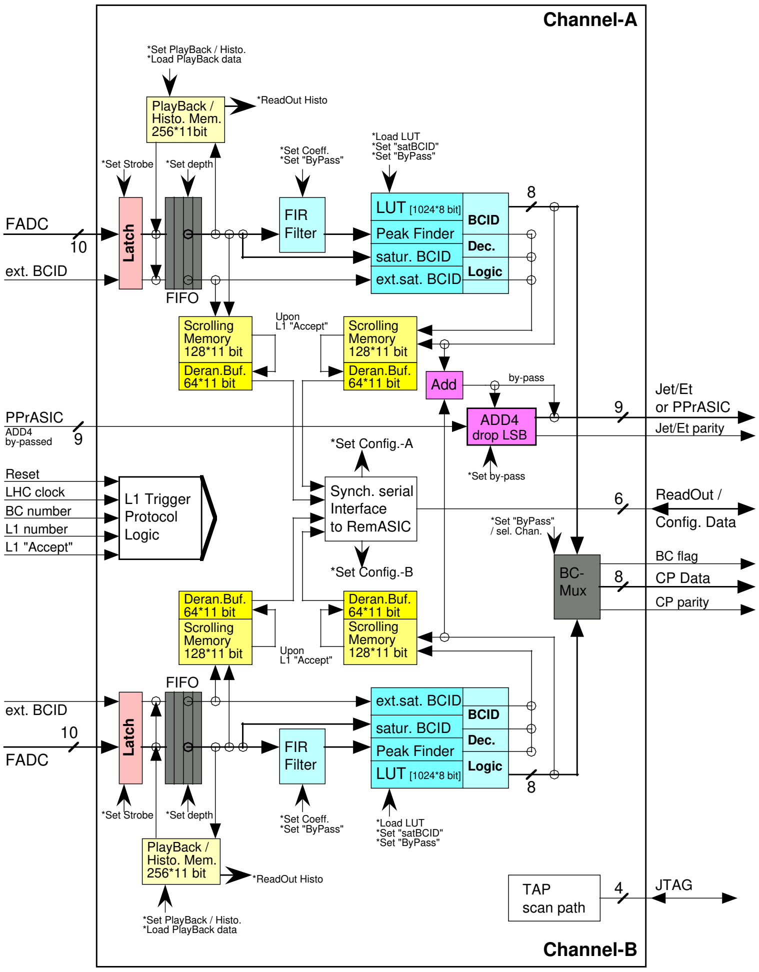

Figure 3.5: Block diagram of the Pre-Processor ASIC [PPR99].

The PipelineBus

The PipelineBus will be used to connect 8 Pre-Processor Modules to one Readout Driver. The PipelineBus consists of a 35-bit-wide parallel shift register (pipeline) of scalable depth. In principle, the bus pipeline is bent in a ring-like fashion to connect modules, where each pipeline stage corresponds to a bus node.

Four different types of nodes exist: a *master node* which controls and reacts to the response of other nodes by injecting commands into the pipeline, a *readout node* which is the source of readout data or the sink for control data, an *S-link node* which transfers readout data to the DAQ via the S-link standard readout link, and a *monitor node*. The RemAsic is an implementation of a readout node and all the other nodes are implemented in FPGAs located on the Readout Driver.

The 35 bits on the bus include a 32-bit data word. Two bits are used as control bits identifying the user data type, and one bit is used as parity bit. The maximum

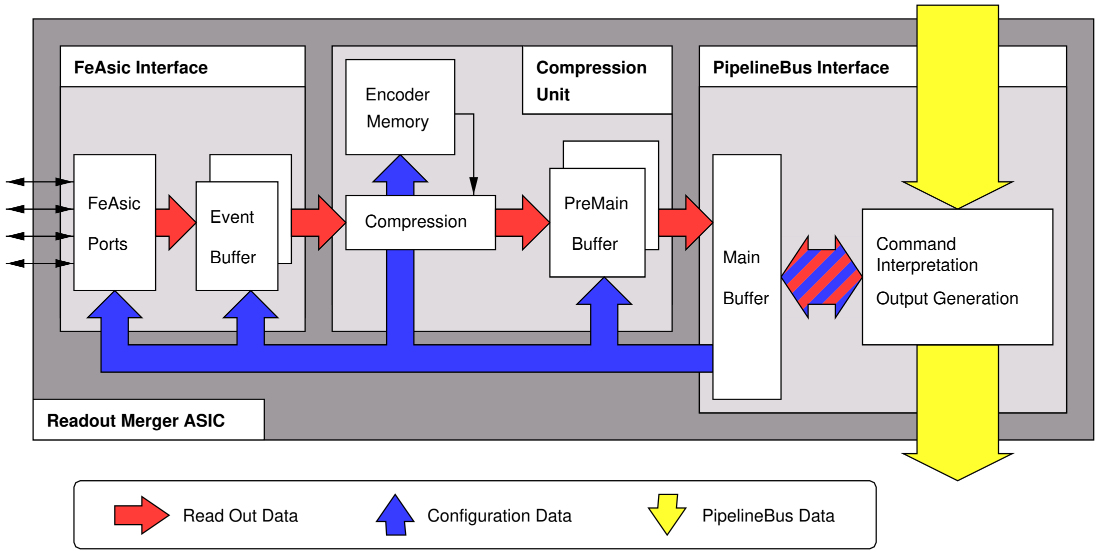

Figure 3.6: Readout Merger ASIC block diagram [Rem98].

bandwidth depends on the frequency of the bus clock. At a frequency of 40 MHz the bandwidth is 40 MHz$\times$32 bit=152 Mbyte/s. This is the theoretical limit, including protocol and idle times.

The Readout Driver

A Readout Driver Module transmits the readout data of 8 PPMs to the ATLAS data acquisition system (DAQ). It consists of three PipelineBus nodes: the master node, the S-link node, and a monitoring and control node independent from the main ATLAS DAQ system. All nodes will be implemented using FPGA technology.

Chapter 4

Bunch-crossing identification

- Non-saturated BCID
- Saturated BCID
- Simulation results
- Bunch-crossing multiplexing

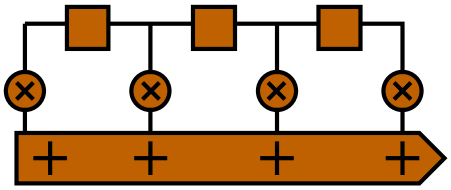

4.1 Introduction

The analogue input signals to the Pre-Processor are called trigger tower signals. They have bipolar shape in the case of the Liquid Argon Calorimeter and unipolar shape in the case of the Hadronic Tile Calorimeter. All trigger tower signals extend over a number of LHC bunch-crossings. The reason for this is that the effect of electronics and pile-up noise is then minimised. The pulse height of a trigger tower signal corresponds to the deposited transverse energy and the position of the pulse maximum identifies the corresponding bunch-crossing. For each of these input signals the Pre-Processor has to extract the following information:

- the amount of transverse energy deposited in a trigger tower, due to an interaction;
- and the corresponding unique bunch-crossing in time to which the energy deposition is related.

This task is referred to as bunch-crossing identification (BCID). The actual BCID decision is performed by two algorithms: a *non-saturated* algorithm responsible for an energy range up to 250 GeV, where the trigger energy scale saturates, and a *saturated* algorithm responsible for energies above that range. The BCID efficiency of the algorithms is of prime importance for the ATLAS experiment. Data from physics events are spread over the whole ATLAS detector. Hence, event data fragments must be aligned in time to the same bunch-crossing. It is not desirable for events to be chopped and distributed between more than one bunch-crossing. Event building and the event reconstruction, which assume that the event data are contained in a single bunch-crossing, would then not be able to extract the physics quantities. At low transverse energy deposits ($<$5 GeV) the efficiency of the energy measurement will affect the performance of the trigger algorithms. A poor BCID for small pulses will degrade the sharpness of the electron/photon, hadron/tau, jet, E_T-miss and sum-E_T triggers and weaken the effectiveness of isolation criteria.

This chapter describes both algorithms which will be implemented in the Pre-Processor ASIC of the ATLAS Level-1 Calorimeter trigger. The chapter starts with an explanation of the requirements for BCID and the difference between saturated and non-saturated calorimeter signals. Next, a short description of the non-saturated BCID algorithm is given. Details of this algorithm and its performance can be found in [Bra96] and [Ree97].

The emphasis in this chapter is placed on the contributions made for saturated BCID. A saturated algorithm and its simulated performance is described in Section 4.3. Simulation results are derived from a PSPICE schematic of the Liquid Argon trigger tower electronics and a PTOLEMY schematic of the digital signal processing inside the Pre-Processor. The PSPICE schematic allows simulation of the saturation in the analogue trigger-tower electronics. The PTOLEMY simulation allows simulation of the digital processing, noise effects, and shifts of the digitisation strobe of the FADC against the pulse position. All simulation results presented here are also described in an ATLAS note, see [Pfe99].

4.1.1 Requirements for BCID

A major requirement for a BCID algorithm is to provide bunch-crossing identification with good efficiency for both low and high E_T pulses. The crucial factor in determining the efficiency for BCID at low E_T is the immunity to noise. There are three sources of noise affecting the signals observed by the BCID logic:

- **Electronics noise:** This consists mainly of thermal noise, where the thermal movement of charge carriers causes statistical fluctuations measured at both ends of a conductor. There are other noise effects applying to semiconductors and active electronics devices. Electronics noise is a criterion of quality for the design of an electronics system.
- **Pile-up noise:** This occurs when two calorimeter signals overlap. It is a function of the luminosity of the accelerator machine, the granularity of the detector and the length of the pulse produced by the calorimeter shaping circuit.
- **Quantisation noise:** This is a source of noise for digital signal processing. It represents the rounding error which occurs in an analogue-to-digital conversion process.

Another requirement for the BCID algorithm is that the implementation of the algorithm should be simple and that it should introduce a minimum additional latency to the Level-1 Trigger. To achieve these goals, the BCID algorithm for non-saturated trigger tower signals requires an input pulse shape which is stable and does not vary with the pulse amplitude (energy). This requirement allows a non-saturated BCID algorithm to successfully extract the timing of a low-E_T pulse peak and therefore the corresponding bunch-crossing in time. In contrast to that, saturated pulses do change their shape. The shape depends on the energy and on the saturation point inside the summing stages of a trigger tower builder. This fact has made it necessary to have an extra algorithm optimised for saturated pulses.

The advantage of the saturated algorithm, which is proposed in Section 4.3, is that it does not rely on a stable pulse shape. It relies on some parameters of the analogue trigger tower electronics. The rising edge of the pulse must be unaffected by saturation at least for two samples before the pulse maximum. In addition, the output voltage slope and bandwidth of the operational amplifiers used in the

BCID requirements and parameters	
Liquid Argon Calorimeter pile-up noise	400 MeV
Liquid Argon Calorimeter thermal noise	217-410 MeV $\eta = 2 - 0$
Liquid Argon Calorimeter total noise	450-570 MeV $\eta = 2 - 0$
Hadronic Tile Calorimeter pile-up noise	90 MeV
Hadronic Tile Calorimeter thermal noise	35 MeV
Hadronic Tile Calorimeter total noise	97 MeV
FADC quantisation noise	75 MeV
Digital saturation level	250 GeV (2,5 V;10-bit)
LAr. Analogue saturation voltage level	300 GeV (3.0 V)
LAr. Shaper time constant	15 ns ± 0.5 ns
LAr. Peaking time (pole-zero)	51-52 ns
LAr. Peaking time at receiver (due to 70 m long cable)	63 ns
Trigger tower timing jitter	always within ± 3 ns
Calorimeter cell timing for trigger tower summing	± 2.5 ns

Table 4.1: Requirements and characteristic parameters for trigger tower signals from the ATLAS calorimeters (at high luminosity).

trigger tower electronics should only be as high as necessary for linear transmission of trigger tower signals. Table 4.1 summarises some requirements and parameters as documented in [TDR98] and [URD98]. The parameters were used as input to the simulations described in Section 4.3.

4.1.2 Comparison between non-saturated and saturated trigger tower signals

The simulated performance of previous BCID options for saturated pulses, e.g. a lookup table BCID described in [TDR98] and [Ree97], were based on an 'ideal' analytic shaper function. A description of this function can be found in [Cha95]. This function allows calculation of the shaper waveform as a response to a triangular detector drift current. Figure 4.1 shows an example for three calculated pulse shapes. Signals simulated with a PSPICE model of the trigger tower electronics [Cle98] are shown in Figure 4.2. The PSPICE pulses shown were simulated to show the saturation inside the layer sum board.

Apart from obvious saturation effects, an important difference between the 'ideal' calculated and the PSPICE simulated pulses is the rise time of a pulse. This time is finite for PSPICE pulses because of the limited output voltage slope of the operational amplifiers.

A problem for previously investigated saturated algorithms is that they rely on a stable leading and falling edge shape of the pulse. Therefore, they might not cope

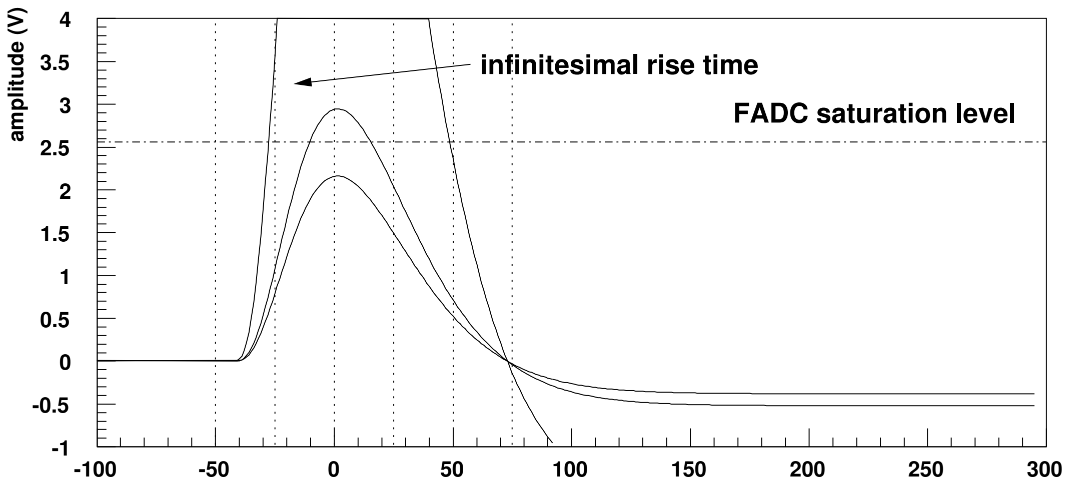

Figure 4.1: Example of trigger tower signals of about 200 GeV, 280 GeV and 1 TeV. The pulse shapes are calculated using an analytical shaper function for a $CRRC^2$ monolithic shaper for the Liquid Argon electromagnetic calorimeter [Cha95].

with pulses saturated at different stages of the trigger tower electronics, because saturation affects mainly the falling edge of the pulse as can be seen in Section 4.3.1.

4.2 BCID for non-saturated calorimeter signals

The BCID logic for non-saturated calorimeter signals consists of a Finite Impulse Response (FIR) filter followed by a peak-finding algorithm. A FIR filter was chosen to achieve good performance at small pulse heights, where the noise influence is significant because of the reduced signal-to-noise ratio. Note that without any noise and without any pulse distortion a peak-finder would be the best one could do to perform BCID. The real world in fact is quite different. The BCID logic in the Pre-Processor must be seen as a receiver in a comunication system. The transmission system involves the transmission of an ideal analogue signal that is affected by noise before being received and digitised in the Pre-Processor.

In general, the behaviour of a comunication system is described by its transfer function:

$$H(j\omega) = \frac{O(j\omega)}{I(j\omega)},$$

where $I(j\omega)$ is the Fourier transformation of the input signal $F[i(t)]$ and $O(j\omega)$ is the Fourier transformation of the output signal $F[o(t)]$.

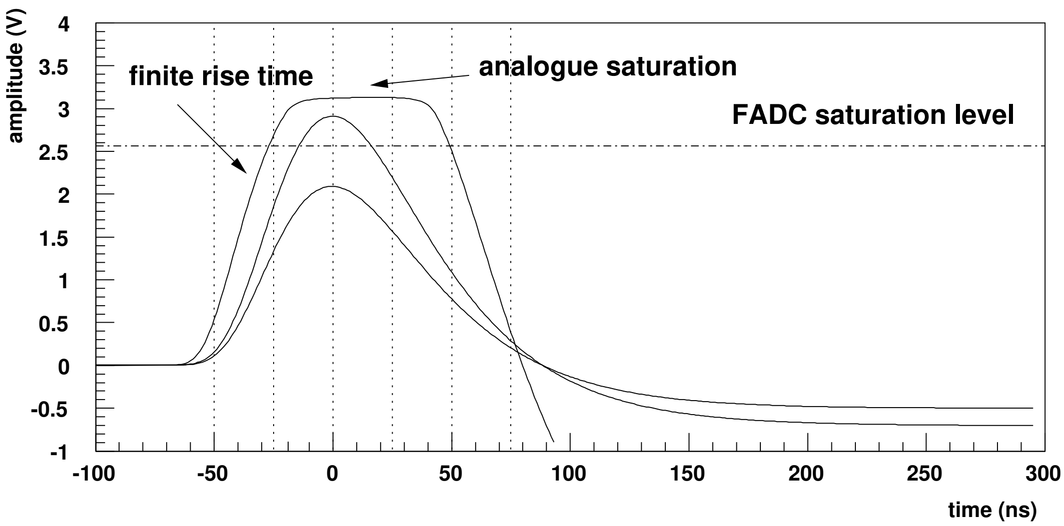

Figure 4.2: Example of simulated trigger tower signals at about 200 GeV, 280 GeV and 1 TeV. The simulation is based on a PSPICE model of the Liquid Argon trigger tower chain including the linear mixer, the layer sum board, the tower builder, the receiver station, and cables in between [Cle98].

The impulse response of the system $h(t)$ is defined to be the output produced by the comunication system when the input consists of a unit impulse at time zero $i(t) = \delta(t)$. This fact can be verified as follows:

$$\begin{aligned} O(j\omega) &= H(j\omega) \cdot I(j\omega) \\ &= H(j\omega) \cdot \underbrace{F[i(t)]}_{=1} \quad \| \, i(t) = \delta(t) \\ &= H(j\omega) \\ &= F[h(t)] \quad \Rightarrow \quad o(t) := h(t) \end{aligned}$$

If the impulse respons of a linear and time-invariant system is known, each possible output can then be calculated by a convolution of the input data with the impulse response:

$$\begin{aligned} o(t) &= F^{-1}[O(j\omega)] \\ &= F^{-1}[H(j\omega) \cdot I(j\omega)] \\ &= F^{-1}[H(j\omega) \cdot F[i(t)]] \\ &= \int_{-\infty}^{+\infty} i(\tau) \cdot h(t-\tau) \, \mathrm{d}\tau \\ &= i(t) * h(t) \end{aligned}$$

A * stands for the convolution integral representation. In case of a discrete comunication system, which calculates the discrete output signal $o(nT)$ at discrete time points 0, T, $2T$,... from the discrete input signal $i(nT)$, the transfer function is an infinite sum of the form [Bes91]:

$$H(z) = \frac{a_0 + a_1 z^{-1} + ...}{1 + b_1 z^{-1} + ...}.$$

The integer number n is the sample number, T is the discrete time separation, and z is a variable transformation of the form $z = e^{sT}$, where s is a complex variable including a phase offset ($s = \sigma + j\omega$). For most of the filter applications it is sufficient to set the denominator polynomial to 1. The transfer function is then an infinite sum of the form:

$$H(z) = \sum_{j=0}^{\infty} a_j z^{-j},$$

where the coefficients a_j define a specific realisation of a transfer function. The problem of designing a digital filter is to calculate the required coefficients for an implemention of a given transfer function. If the sum and therefore the number of coefficients is infinite the filter type is referred to as an Infinite Impulse Response filter (IIR filter). If the sum is finite the filter type is referred to as a Finite Impulse Response filter (FIR filter). In the time domain the discrete output is calculated as:

$$o(n) = \sum_{j=0}^{\infty} a_j \cdot i(n-j). \tag{4.1}$$

4.2.1 The finite impulse response filter

As mentioned above, the coefficients for a FIR filter define the transfer function of the system. This is independent from the actual implementation of a FIR filter. A FIR filter implementation looks always the same. Samples are presented to the filter sequentially and the most recent samples are kept in a 'pipeline' of registers. The value in each register is multiplied by the filter coefficient a_j and summed to form a filter output. Figure 4.3 illustrates the implementation of a FIR filter.

The quality of an implemented transfer function depends on the depth of the pipeline register and the precision of the coefficients. The transfer functions of the following filters have been considered for the Pre-Processor and can be implemented in a FIR filter with a considerable amount of digital logic [Bra96]:

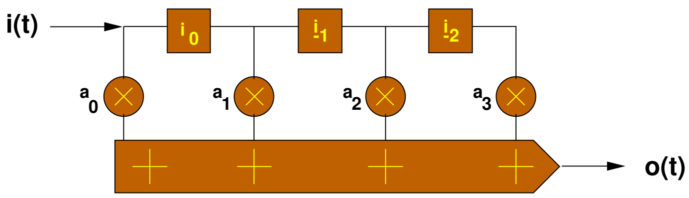

Figure 4.3: FIR filter implementation.

- **Deconvolution filter:** The aim is to deconvolute the shaped trigger tower signal back into an impulse signal.
- **Zero-crossing filter:** This filter extracts the timing of the zero-crossing of a bipolar pulse only.
- **Constant-fraction discriminator:** This filter sets a threshold to a constant fraction of the total pulse height and extracts the timing information when this threshold is crossed.
- **Comparison of the pulse with its integral:** This filter fires when the pulse height is equal to a constant fraction of its own integral.
- **Peak-sharpening algorithm:** This filter tries to sharpen the input pulse.
- **Differentiation filter:** This filter calculates the gradient of an input pulse.
- **Matched filter:** A matched filter calculates the maximum likelihood for the pulse samples. This filter has shown the best performance and has been chosen for the Pre-Processor.

Any of the filter types above can be used for the Pre-Processor because the filter coefficients can be loaded into the Pre-Processor ASIC. Limitations to the quality of the filters exist, because multipliers represent a considerable amount of digital logic. Hence, the pipeline depth and coefficient precision has to be optimised. Therefore, the amount of logic required for the FIR filter implementation is optimised for a matched filter, because this filter has shown the best BCID performance.

Matched filter

For an explanation of a matched filter it is convenient to represent the originally transmitted signal and the received signal as vectors. The vector $\overrightarrow{p}$ represents the original transmitted signal, where each column element (x_i) represents a sampled value of the pulse shape. The vector $\overrightarrow{x}$ represents the received signal after being affected by noise. The number of vector elements is defined by the pulse width divided by the time separation of the sample points. For each bunch-crossing a matched filter calculates the probability that the samples contained in the FIR filter pipeline 'match' those of the original vector $\vec{p}$. For noise which follows a normal distribution, the conditional probability density has to be evaluated for each bunch-crossing [Gar96]:

$$f(\vec{x}|\vec{p}) = Ae^{-\frac{1}{2}(\vec{x}-\vec{p})^T\mathbf{C}^{-1}(\vec{x}-\vec{p})},$$

where A is a normalisation constant and $\mathbf{C}$ is a noise covariance matrix that describes the variance of each sample and the covariance between samples. This equation has to be calculated for each bunch-crossing. A maximum of this probability is then equivalent to the hypothesis that the received vector $\vec{x}$ is contained in the same bunch-crossing as the vector $\vec{p}$. To determine the maximum of that equation (maximum likelihood) it is sufficient to calculate the logarithm and neglect the constant factors. This results in the following equation:

$$\begin{aligned}\ln f(\vec{x}|\vec{p}) &= \vec{p}^T\mathbf{C}^{-1}\vec{x} \\ &= \vec{a}\cdot\vec{x} \qquad \| \; \vec{a} := \mathbf{C}^{-1}\vec{p}\end{aligned}$$

This is just a simple sum over each of the vector elements (x_i) weighted by a coefficient (a_i). It corresponds to what was defined for the calculation of a FIR filter output, see Equation 4.1. Hence, the values of the matched filter coefficients depend on the pulse shape and the characteristic of the noise. If no noise were present the coefficients are just the samples of the expected pulse shape. For 'white' noise the covariance matrix $\mathbf{C}$ is diagonal and therefore it is its inverse. In this case the vector $\vec{a}$ is proportional to $\vec{p}$. For non-white noise, which indeed is what we will expect, the matrix and its inverse are not diagonal and the pipeline length needs to be enlarged to achieve optimal performance.

The following is a listing of the FIR filter parameters which have implications for the precision and quality of a matched FIR filter implementation. The parameters are chosen to provide good BCID efficency for a matched filter with a minimum of required digital logic [PPR99], [TDR98]:

- **Number of filter stages:** The depth of the FIR filter is defined to be five.
- **Coefficient precision:** Four-bit wide coefficients are accurate enough for good noise reduction.
- **Digitisation precision:** Digitisation to 10-bit precision provides a good resolution for BCID.
- **Sampling rate:** A sampling rate of 40 MHz is sufficient if the digitisation is done at 10-bit precision.

4.2.2 The peak finder

The peak finder takes as input three consecutive samples P. It recognises a peak within these three samples if the following condition is fulfilled:

$$P_{i+1} \leq P_i > P_{i-1}.$$

A peak finder can be easily implemented in hardware by the use of a set of comparators. For a set of samples which do not fulfil this condition, the output of the peak finder is set to zero.

4.3 BCID for saturated calorimeter signals

Each trigger tower signal will saturate for energies beyond the linear transverse energy range (0–250 GeV). Two saturation effects can be distinguished: a digitisation effect and an analogue saturation effect. The first effect is due to the upper digitisation bound of the FADC, which is 250 GeV (3FF hex) at 10-bit precision. The second effect is due to a limited output voltage of the analogue trigger tower electronics. Hence, a bunch-crossing identification algorithm for saturated trigger tower signals must cope with both saturation effects in a flexible way. This section describes an algorithm which extracts the corresponding bunch-crossing for saturated trigger tower signals after digitisation. This algorithm combines BCID efficency with enormous simplicity in design and implementation. It takes only two samples from the leading pulse edge and compares their values against programmable thresholds.

4.3.1 Simulation of analogue calorimeter signals

The simulations described here, were performed with Liquid Argon Calorimeter signals only. For the hadronic Tile Calorimeter, where big pulses will be rare, saturation will be less of a problem. In addition, the trigger-tower signals are formed by a much reduced number of calorimeter cells and hence the variation of saturation along the electronics chain will be reduced.

For the Liquid Argon calorimeter, saturation can first occur inside the analogue summing chain of trigger tower electronics. In addition to that the signal is clipped digitally. The digitisation range at the Pre-Processor is limited to 2.5 V at 10-bit precision. All signals which reach 2.5 V or more at the Pre-Processor are represented by a maximum FADC digitisation value of 3FF (hex). Analogue saturation can occur at three stages of the trigger tower processing:

- at the linear mixer, including the pre-amplifier and the shaper;

- at the layer sum board;
- at the tower builder board, which is the stage that is most likely to saturate.

The analogue saturation voltage of the simulated trigger tower electronics is 3.0 V. The distorted pulse shape received at the Pre-Processor differs depending on the stage where saturation occurs. The presence of pile-up noise and thermal (electronics) noise makes it necessary to optimise the Liquid Argon shaper time constant individually for different calorimeter regions. Because of this difference, the peaking time[1] must be readjusted at the tower builder board before signals are summed together. This adjustment is done with a pole-zero circuit, which differs for different layers and each value of pseudo-rapidity (η). The uniform peaking time after the pole-zero circuit is about 51–52 ns at its output. Cable integration[2] of a 70 m twisted-pair cable slows down the signal further to about 63 ns. Because of the pole-zero circuit and the cable integration the expected peaking time at the Pre-Processor input is between 55–63 ns.

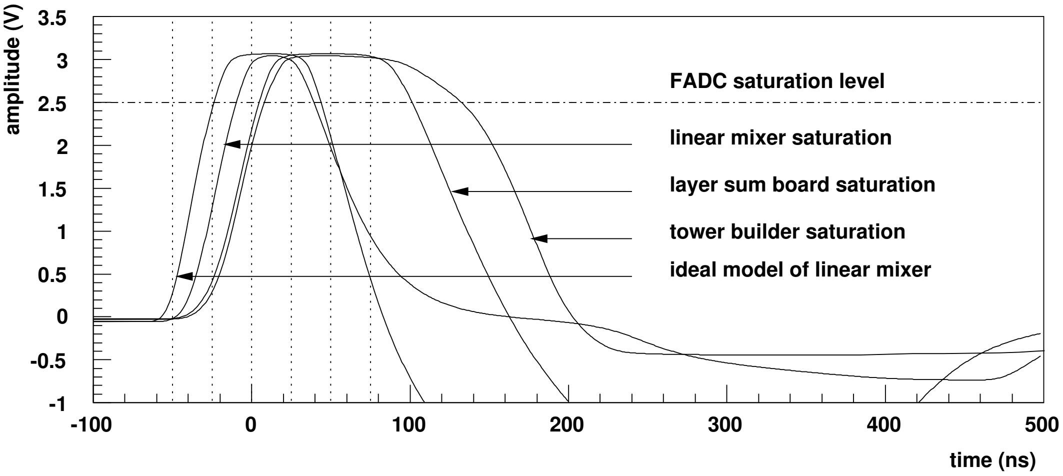

Figure 4.4: Pulse shape of saturated trigger tower signals. Saturation was simulated at different stages of the analogue trigger tower electronics. The vertical dashed lines are spaced 25 ns apart to indicate the LHC bunch-crossing. The horizontal dashed line at 2.5 V is the FADC saturation level (3FF hex).

Figure 4.4 shows a sample of simulated pulse shapes. All signal shapes were captured at the receiver station output, representing an energy of about 1 TeV. The saturation was simulated at different stages of the trigger tower electronics: at the linear mixer, the layer sum board, and the tower builder board, using a PSPICE model [Cle98]. The model for the analogue chain allows simulation of the expected

[1] The peaking time is measured from 5 % to 100 % of the peak maximum.

[2] A more conservative length of 70 m was used. The actual cables are not going to be longer than 60 m.

pulse shape for the middle layer at $\eta = 0$. To adjust the model to a different layer and/or η-segmentation the shaper circuit must have a different shaper time constant, the pole-zero circuit and the cable length have to be changed.

Saturation of the linear mixer has little effect on the pulse maximum. The peak is still well located and it does not get broader with higher energies. The integrated area is no longer zero because of a delayed start of the undershoot. The starting point of the undershoot moves to the right with higher energies. This shape corresponds with measurements carried out with the module-0 Liquid Argon Tri-gain shaper [Col98]. Saturation inside the layer sum board results in a large flat top at the maximum, which gets broader for higher energies. Due to the symmetrical power supply, the undershoot does not saturate at 1/5 of the positive saturation voltage. It is amplified up to –3.0 V and therefore the integrated area is no longer zero. The tower builder circuit affects the signal shape by giving the saturated signal a hump on the falling edge and it has by far the widest extension to later times. The distortion of the signal shape, as a consequence of saturation, affects mainly the falling edge of the shape and has little effect on the rising edge. The signal slope for the rising edge increases first for higher energies, but later it is limited by the maximum output voltage slope of the operational amplifiers. This fact is used by the digital saturated algorithm to efficiently identify the corresponding bunch-crossing even when the signal shape changes due to the saturation stage in the analogue chain. Variations between the origin of the pulses occur because of different propagation delays between an ideal PSPICE model for the linear mixer, which was recently replaced by a more realistic one. The pulse shape for the ideal model is included here, since all simulations, which follow, were done with it. Because of the ideal linear mixer circuit, its simulated saturation is equivalent to saturation in the following stage, the layer sum board.

4.3.2 Saturated BCID algorithm

The motivation to use a discriminator on the leading edge of the pulse shape arises from the fact that the pulse origin is the only fixed attribute of the pulse defining the interaction time. Once the origin of a pulse is located, it is just a matter of adding the peaking time to identify the correct bunch-crossing in time. In this case the signal shape can be distorted in various ways without affecting the identified bunch-crossing. An analogue discriminator BCID, described in [TDR98], is one possible realisation of that idea. The drawback of this solution is the additional analogue electronics and the cost needed for its implementation in the Pre-Processor system. It would be much preferred, if a digital saturated algorithm could instead be included in the Pre-Processor ASIC. Looking at the digitised pulse shape, one can see that the second sample before the peak is close to the origin of the pulse (the peaking time is about twice the LHC bunch-crossing of 25 ns). Setting a digital threshold for the digitised data is therefore quite similar to an analogue threshold before digitisation.

A digitisation at the maximum of the pulse is preferred to optimise the accuracy of the measured transverse energy. This is not a requirement for the Level-1 trigger, but the Pre-Processor can do so because it is able to adjust the relative timing of analogue trigger-tower signals to the FADC strobe. The required step size for the FADC strobe is only 2.5 ns [URD98], whereas the proposed one for the Pre-Processor is 1 ns [TDR98]. Starting at the peak maximum and going 50 ns (about one peaking time) back in time gives two samples on the rising edge. The first sample sits close to the origin of the pulse and the second one is in between the baseline and the upper digitisation limit of the FADC. The following section describes the working of a digital saturated BCID algorithm which applies two thresholds to both samples before saturation.

Saturated pulse detection

The algorithm has a programmable saturation level, which triggers its execution. When a 10-bit value from the FADC reaches or goes beyond that level the algorithm is performed once. It is only re-armed with the next FADC sample which falls below this saturation level. To maximise the available threshold range for the algorithm, this value should be identical to the FADC saturation of 3FF. It may be useful to lower this value in order to increase the overlap region between the saturated and the non-saturated algorithm.

Bunch-crossing determination

Little digital logic is required to determine the correct bunch-crossing for saturated signals. The algorithm is fully defined by the logic given in Table 4.2, and its validity can be seen from pulse examples shown in Figures 4.5 and 4.6.

The algorithm is triggered with the first occurrence of saturation, at (t_{sat}). Two samples before saturation, which are held inside of a FIFO buffer, are then compared with programmable thresholds stored in registers. Depending on the crossing of thresholds, either the first occurrence of saturation at t_{sat} is identified as the correct bunch-crossing, or the next time slice which follows at t_{sat+1}. Figure 4.5 a) gives an example of a trigger tower signal which just triggers the saturated algorithm at t_{sat}. Sample points which are used as input to the algorithm are drawn as dots sampled at t_{sat-2}, t_{sat-1} and t_{sat}. This figure is also valid for larger signals where the sampled FADC values are increased further until the sample at t_{sat-1} reaches saturation, too. In this case the algorithm would trigger and take its input sample points 25 ns earlier (see Figure 4.5 b)). For very large pulses in Figure 4.6 a) the influence of the limited output voltage slope of the operational amplifier used is shown. The sample at t_{sat-1} cannot cross the s_{high} threshold any more because

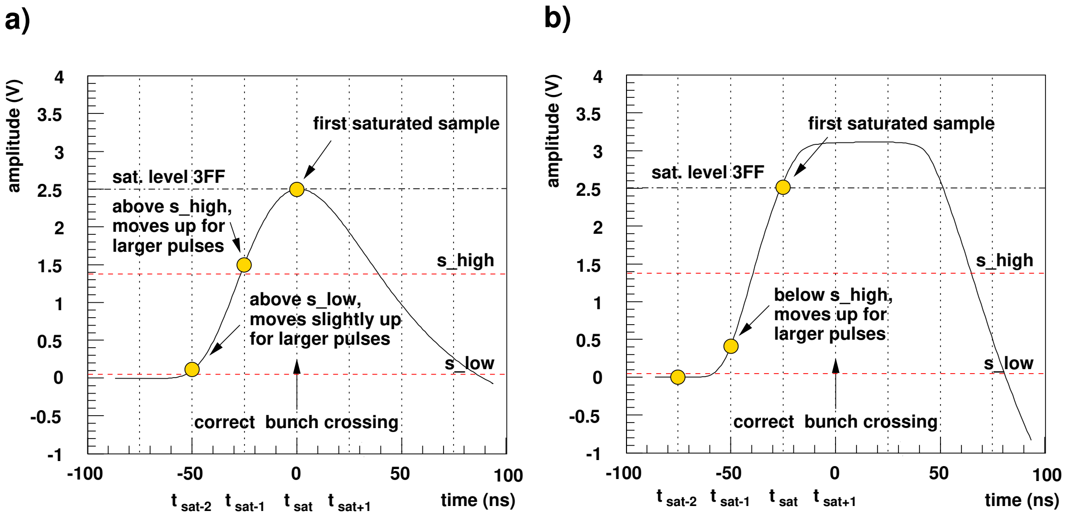

Figure 4.5: Pulse example (250 GeV) which just saturates the FADC is shown in a). Because of the sample at t_{sat-1} which is above the threshold s_{high}, the algorithm identifies the sample at t_{sat} as the correct bunch-crossing. A larger pulse which saturates 25 ns earlier is shown in b). In this case the sample at t_{sat-1} is below s_{high} and therefore the algorithm identifies the sample at t_{sat+1} as the correct one.

of the fact that the output voltage cannot reach saturation from 0 V to 2.5 V in infinitesimal time.

So far the threshold s_{low}, which gives further robustness against a misplaced FADC strobe, has not been needed. Its role can be seen from Figure 4.6 b), which one can get from Figure 4.6 a) just by introducing a phase offset between pulse and FADC digitisation strobe. Now the sample at t_{sat-1} crosses s_{high}, which would cause a misidentified bunch-crossing. In contrast to Figure 4.5 a) the sample at t_{sat-2} is now below s_{low} and therefore the s_{low} threshold can correct the decision.

Transverse energy measurement

The saturated digital BCID algorithm identifies the corresponding bunch-crossing only, it does not measure the energy. The setting of the measured transverse energy is up to the programming of the BCID decision logic, which collects the results for non-saturated and saturated BCID algorithms. Possible values are 3FF, the saturation level if it differs from 3FF, or the output result from the FIR filter lookup table for non-saturated pulses. The later case might differ from 3FF for signals which just reach the saturation level. For a description of the processing chain of the Pre-Processor see Chapter 3.

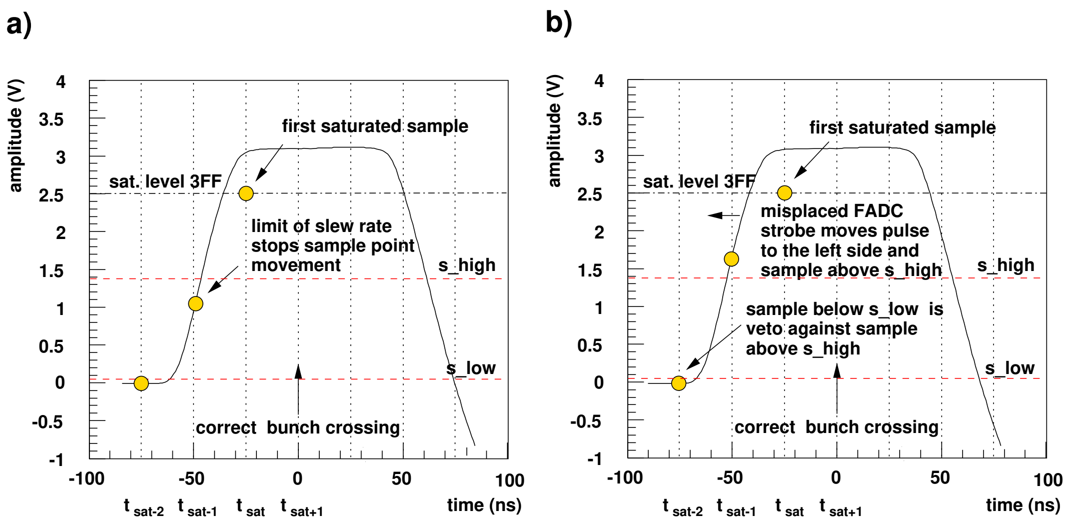

Figure 4.6: The effect of the limited output voltage slope of the operational amplifiers is shown in a). The sample at t_{sat-1} will not cross the s_{high} threshold for any further increased energy. For a misplaced FADC strobe, a crossing of s_{high} might be possible as shown in b). This case can be distinguished from the one shown in a) of Figure 4.5 by the s_{low} threshold.

Implementation

The algorithm can be described with Verilog HDL by implementing the logic shown in Table 4.2. A possible schematic implementation is shown in Figure 4.7. It consists of three programmable registers, a 10-bit FIFO with depth of three, and three 10-bit comparators applied to FADC samples before saturation. The result is sent to a programmable BCID decision logic which collects results for the non-saturated and saturated algorithms. The synchronisation of the result bits and the re-arming of the algorithm is not shown in this schematic. Another solution would be to compare each FADC sample in parallel with both thresholds and to pipeline the resulting bits only.

$t_{sat} \geq 3FF$	$t_{sat-2} > s_{low}$	$t_{sat-1} > s_{high}$	t_{sat}	t_{sat+1}
0	X	X	0	0
1	1	1	1	0
1	0	1	0	1
1	X	0	0	1

Table 4.2: Logic table to be implemented in the Pre-Processor ASIC. Fulfilled threshold conditions are indicated by a logic ‘1’; the identified bunch-crossing has a logic ‘1’ at t_{sat} or t_{sat+1}.

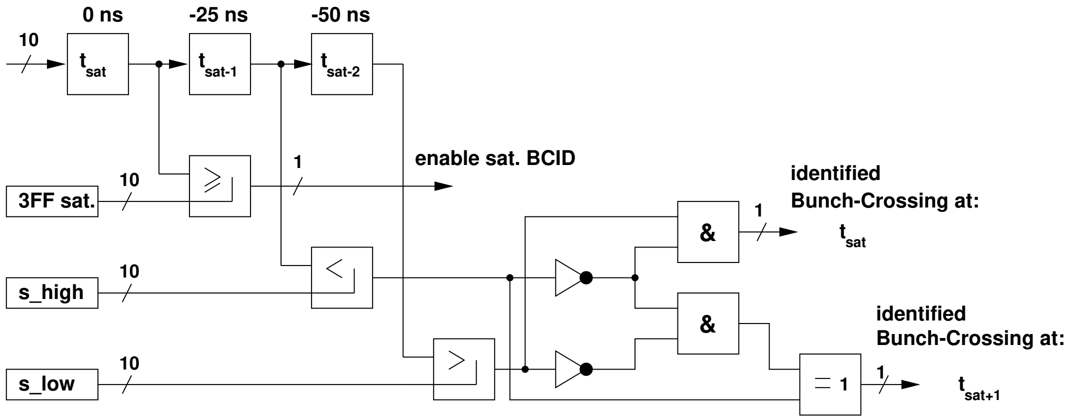

Figure 4.7: Possible schematic implementation of the saturated BCID logic. The logic sets thresholds on raw FADC data stored inside a FIFO-buffer. The identified time slice is indicated to the BCID decision logic. Timing synchronisation of the result bits and the re-arming of the algorithms are not shown.

FIR filter discussion

The inclusion of a finite impulse response filter (FIR filter) in the algorithm was investigated. The coefficients were chosen to perform different algorithms: a matched filter, which calculates the maximum likelihood for the samples inside the FIR filter, a differentiation, and a peak-sharpening algorithm. A matched filter, which uses the convolution of the pulse samples with a noise covariance matrix as coefficients, relies on a fixed saturated-pulse shape (see Section 4.2.1). For very large pulses the distortion of the falling edge precipitates mis-identified bunch-crossings. A peak-sharpening algorithm seems to work better. It extends the timing stability, but again for very large pulses the falling edge affects the efficiency, which is not acceptable. Only if one could avoid more than two saturated samples is it an attractive option. Differentiation of the pulse has a large uncertainty for the threshold settings. This is the reason why the simulation results presented here are limited to thresholds applied to the raw FADC data only.

Latency summary

Without a FIR filter the latency of the digital saturated algorithm is one bunch-crossing (25 ns). The inclusion of a FIR filter in the circuit would increase the latency to 5 bunch-crossings (125 ns).

4.3.3 Simulation environment

This section describes the simulation environment for the digital saturated BCID algorithm. The simulation is a combined simulation of the analogue trigger tower electronics and a simulation of the digital signal processing in the Pre-Processor.

Simulation chain

Figure 4.8 shows the simulation chain. It is divided into two simulators: a MicroSim PSPICE-V8 simulator for the analogue part and PTOLEMY 0.7.1 from the University of California at Berkeley, which combines a synchronous data flow simulation (SDF) and a discrete event-driven (DE) simulation for the digital signal processing. The PSPICE schematic [Cle98] includes a circuit for the linear mixer, the layer sum board, and the tower builder. It has an ideal model for a 70 m twisted pair cable included as well as a schematic of the receiver station. The input for that circuit is a triangular Liquid Argon drift current, with a drift time of $t_{dr} = 400$ ns and a peak current of 3.2 μA/GeV (middle layer at $\eta = 0$). The pole-zero circuit is chosen for the middle layer at $\eta = 0$. The circuit is supposed to have a peaking time of 51–52 ns after the tower builder. At the receiver station output the peaking time is about 63 ns.

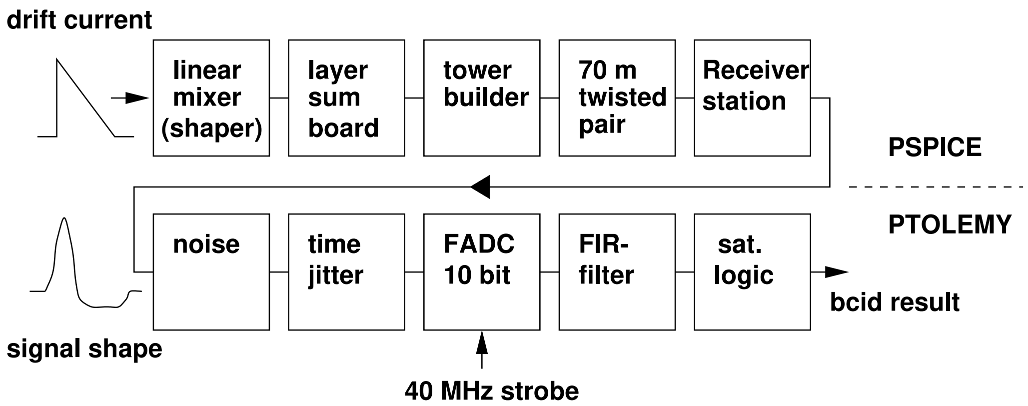

Figure 4.8: Simulation chain used to simulate analogue trigger tower signals and the digital signal processing inside the Pre-Processor system.

The output of that PSPICE simulation is interfaced to a PTOLEMY schematic. The PTOLEMY simulation adds first electronic and then pile-up noise to the PSPICE waveform (RMS = 440 MeV). The whole signal is then delayed to simulate time jitter, with a Gaussian distribution of $\sigma = 1$ ns. An ideal model of an FADC digitises the signal at 10-bit precision, where the FADC digitisation strobe can be moved in steps of 1 ns against the peak maximum. The FIR filter that follows is included here, but it is used in bypass mode. The saturation logic can be described in C++ within PTOLEMY. As a result of the simulation, the exact position of the peak maximum is compared to the result from the saturated BCID logic.

PSPICE pulse generation

The PSPICE pulses which were used as input to the PTOLEMY simulation were generated for discrete drift currents equivalent to energies ranging from 100 GeV up to 10 TeV. The saturation took place inside the tower builder board. In order to simulate the effect of different peaking times at the Pre-Processor input, the shaper time constant was varied. Pulses were generated for shaper time constants of 5 ns, 10 ns and 15 ns (t_{pk}=43 ns, 50 ns, 63 ns). Figure 4.9 shows examples for non-saturated 100 GeV pulses. Due to the longer integration time invoked by a longer shaper time constant, the peak height changes as well as the peak position. The peak energy at this point was not readjusted for the simulation.

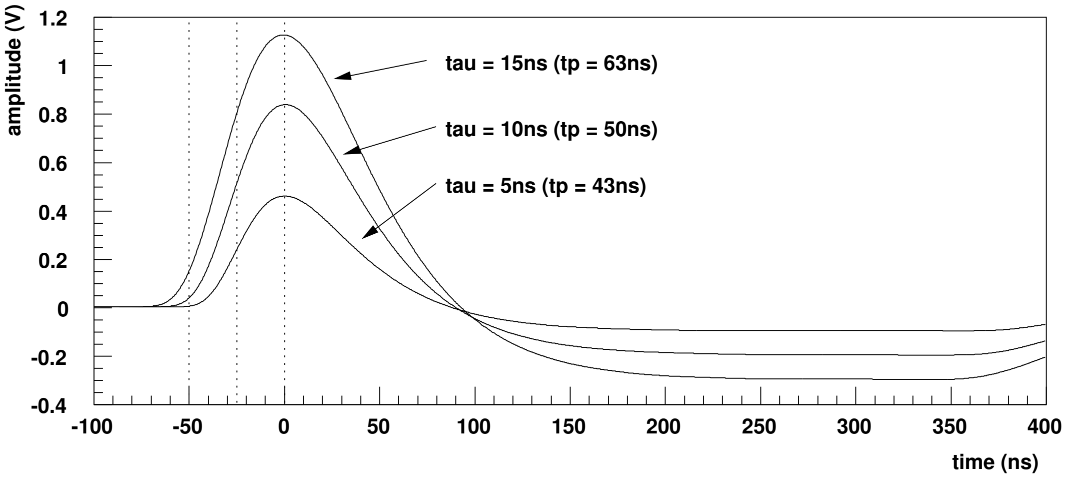

Figure 4.9: Simulated pulse shape (100 GeV) for different shaper time constants. The corresponding peaking time is given in brackets. The dashed lines represent the FADC digitisation strobe of 40 MHz.

Because of an expected[3] peaking time at the Pre-Processor input of about 42 ns to 50 ns, these sets of peaking times should be sufficient to test the flexibility of the algorithm against changes of the peaking time. Once an exact measurement of the cable characteristics is done, the expected variation of peaking times can be estimated more precisely.

Pulse shape for different detector capacitances

Figure 4.10 shows how the pulse shape is affected by changes made to the detector capacitance. Changes of the detector capacitances between 800 pF up to 1200 pF

[3]The expected shaper time constant has just recently changed to 15 ns. This is the reason why longer peaking times were not investigated. The expected peaking time at the Preprocessor input is now in between 55–63 ns depending on the cable length. Longer peaking times should be explored to give confidence that the algorithm is also valid for extreme values.

have only a small effect on the pulse shape and on the peaking time, as can be seen from the figure.

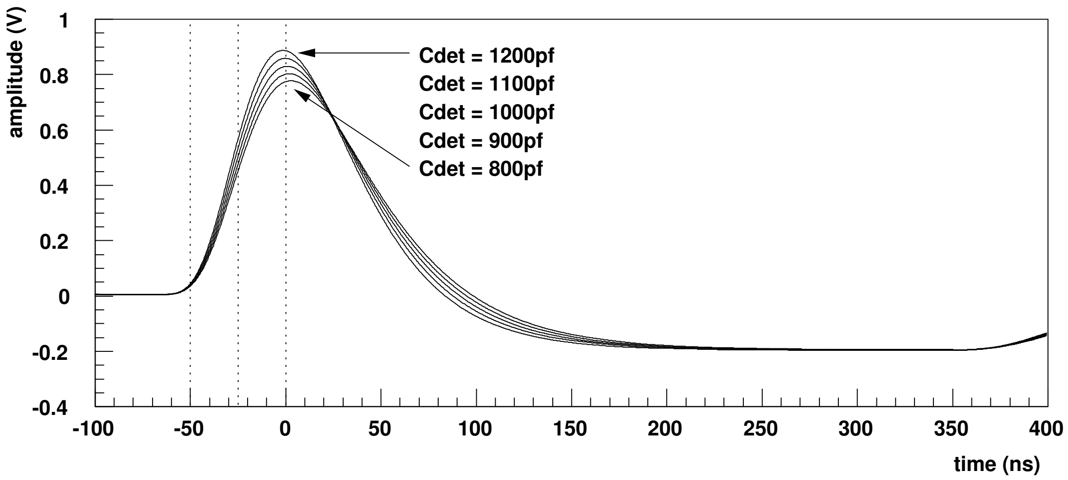

Figure 4.10: Simulated pulse shape (100 GeV) for different detector capacitances (800pF - 1200pF). The dashed lines represent the FADC digitisation strobe of 40 MHz.

Effects on pulse shape due to radiation effects

Preliminary irradiation results for the backup version of the Liquid Argon TRI-gain shaper [Col98] have shown that the peaking time is stable within ± 0.5 ns for 10 times the expected dose at LHC. It shifts to a faster peaking time for long term irradiation. Instant changes of the peaking time or a distortion of the shape were not examined.

Simulated noise contribution

The noise which was added to the signal waveform was generated as described in [TDR98]. It contains the pile-up spectrum taken from ATRIG[4] at high luminosity and the expected noise spectrum for the Liquid Argon electronics, followed by a fourth-order low-pass Butterworth filter with a cut-off frequency of 20 MHz. Figure 4.11 shows the amplitude distribution of that noise. The RMS value taken from that distribution is equivalent to 440 MeV[5].

Simulated time jitter

All PSPICE simulated signals were delayed by a value taken from a Gaussian distribution with a standard deviation of 1 ns. The maximum values from that

[4] ATRIG: ATLAS software package for trigger simulation

[5] 10 mV corresponds to 1 GeV transverse energy

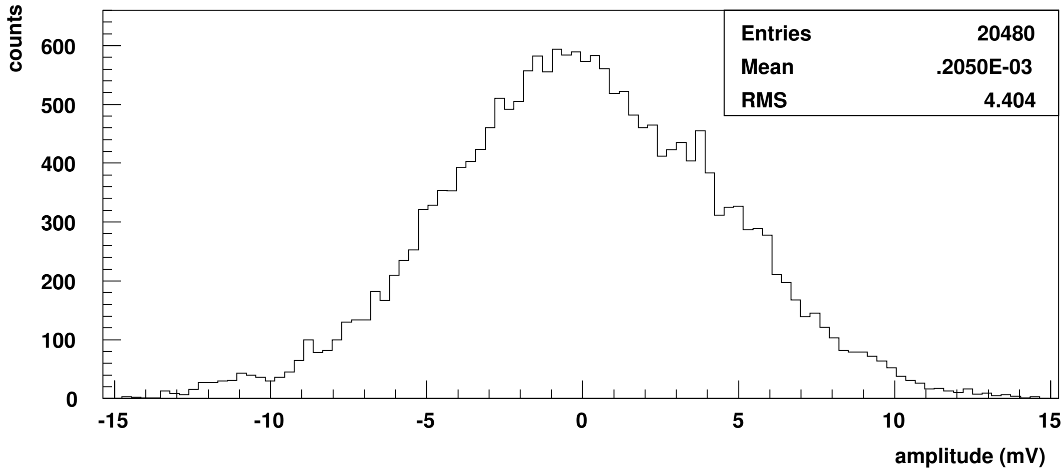

Figure 4.11: Expected amplitude distribution of the Liquid Argon calorimeter noise, including electronics noise and pile-up noise taken from ATRIG at high luminosity [TDR98]. The RMS noise is 4.4 mV (440 MeV).

distribution were limited to ± 2 ns. This should simulate unexpected movements of the signal origin caused by short-term effects which cannot be calibrated for.

Synchronisation of digitisation phase to pulse maximum

The threshold settings of the algorithm are programmable and should be optimised for each position of the FADC digitisation strobe relative to the peak maximum. However, the simulation described here changes the timing of the FADC digitisation strobe without loading new thresholds. This is done to test the timing stability and robustness of the algorithm.

4.3.4 Simulation results

For each simulation two plots were generated. The plot labeled a) contains a result bit for each simulated pulse amplitude. The entry is '1' for those events which were identified with the correct bunch-crossing. For amplitudes which are below saturation (2.5 V), and for those which were mis-identified, the entry is '0'. The plot labeled b) shows the value of the FADC sample before the first occurrence of saturation. This is the sample at t_{sat-1} which is compared to the threshold s_{high}. A dashed line indicates the threshold setting which is a parameter for that simulation. Entries below the threshold have identified the bunch-crossing to the pulse sample at t_{sat+1}, entries above to t_{sat}. The sample at t_{sat} is the one which first triggers the saturated algorithm, and the sample at t_{sat+1} follows one time slice (25 ns) after the first saturation.

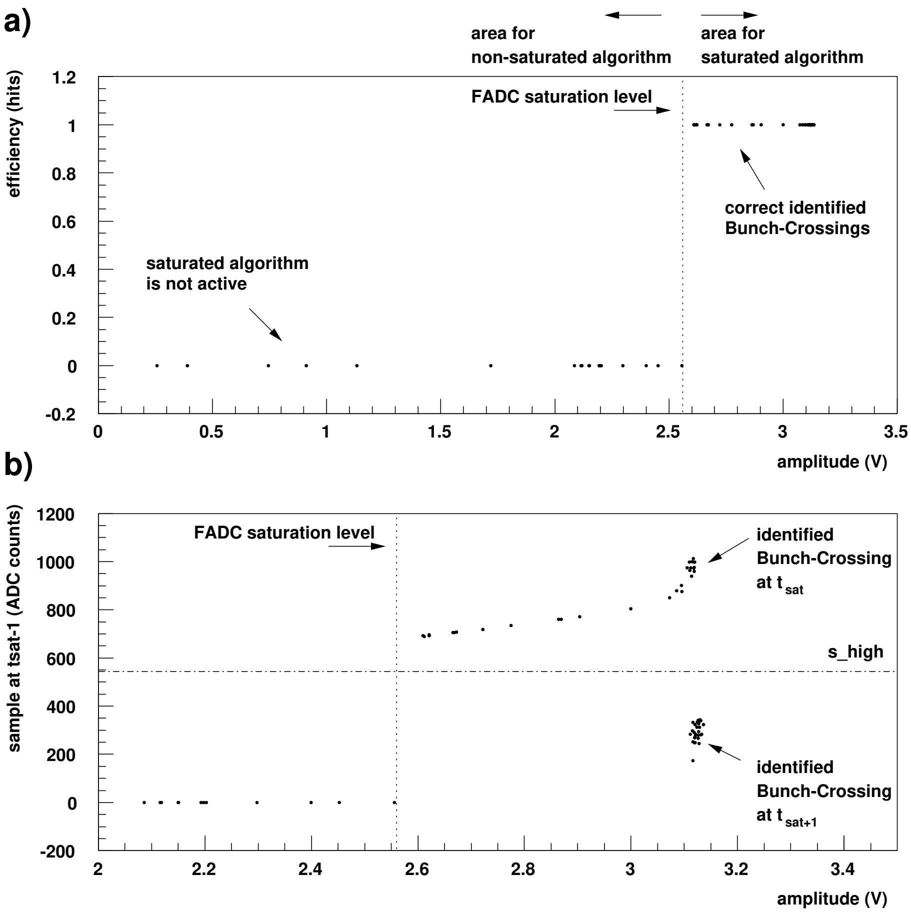

Figure 4.12: Efficiency plot a), and sampled FADC value before first occurrence of saturation at t_{sat-1} b). Shaper time constant is 10 ns ($t_{pk} = 50ns$), the digital discriminator threshold is set to 550 and the RMS noise is 4.4 mV. The strobe of the FADC is exactly synchronised to the pulse maximum and no time jitter was included.

Figure 4.12 shows the effect of noise applied to the signal. The pulse maximum was exactly synchronised to the FADC digitisation strobe and no time jitter was included. The top 'efficiency' plot has a sharp transition region at the saturation level of the FADC (2.5 V). All events were identified correctly up to simulated pulse energies of 10 TeV. The plot underneath, which shows the sampled FADC value before saturation, is smeared out because of the noise effect on the signal. From this plot one can see that the separation for the samples which must be discriminated at the threshold (s_{high}) is about 350 FADC counts. Within that region the threshold can be set to any desired value, but an optimised position would be the centre of that region. Theoretically, that region would be about 400 FADC counts (about 1 V) without any noise.

Figure 4.13 includes the simulated time jitter, as described in section 4.3.3. The time jitter has no effect on the efficiency of the algorithm. A threshold of 550

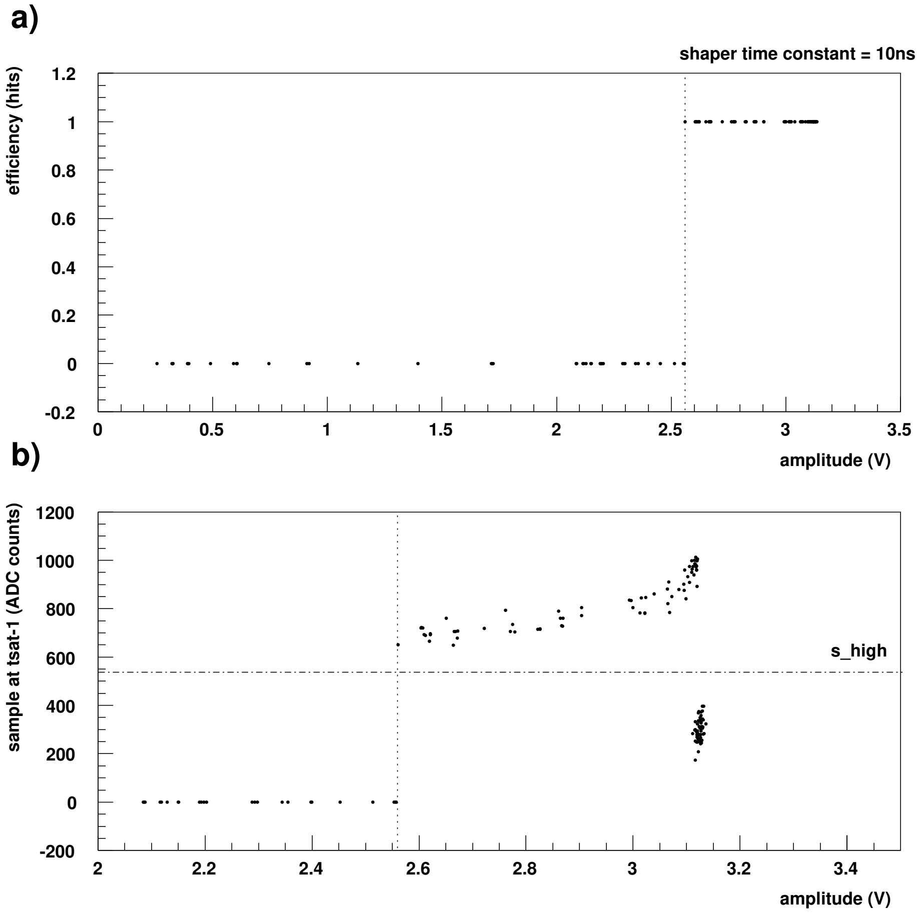

Figure 4.13: Efficiency plot a) and sampled FADC value before first occurrence of saturation at t_{sat-1} b). Shaper time constant is 10 ns (t_{pk} = 50 ns), the digital discriminator threshold is set to 550, the RMS noise is 4.4 mV, and time jitter of the pulse is ± 2 ns. The strobe of the FADC is synchronised to the pulse maximum.

is still well located to distinguish the sampled values. If the FADC digitisation strobe were now to be moved away from the peak maximum position, the samples shown inside the bottom diagram would move. They would move up for an FADC sampling point shifted towards the falling edge, and they would move down for an FADC sampling point shifted towards the rising edge of the pulse. In such a case the threshold is not centred anymore and should be re-optimised with respect to the FADC strobe. A worst case simulation, which includes noise, time jitter and a misplaced FADC digitisation strobe of ± 3ns without changing the threshold setting, is shown in Figure 4.14. The algorithm still identifies the correct bunch-crossing up to simulated saturated signals of 10 TeV, but the safety margin to set the threshold is reduced. The simulation in Figure 4.15 was done as described before, but the noise was increased by a factor of three. One can see that the samples are smeared out further, but all events are still identified correctly.

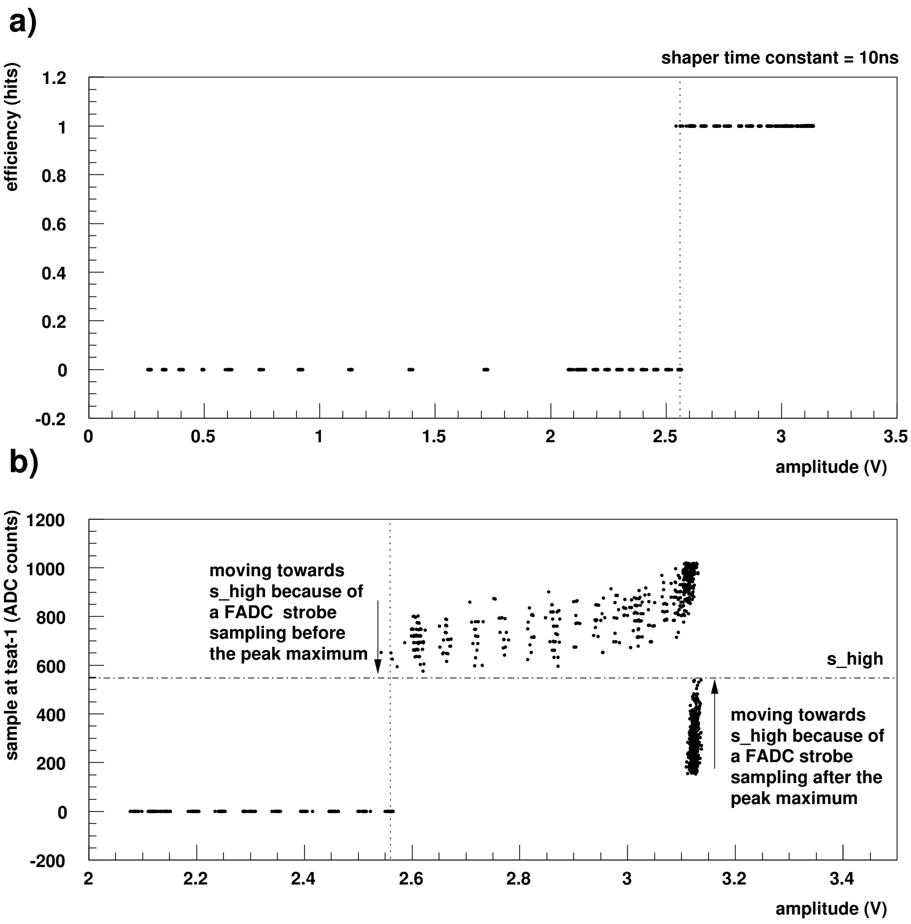

Figure 4.14: Efficiency plot a) and sampled FADC value before first occurrence of saturation at t_{sat-1} b). Shaper time constant is 10 ns ($t_{pk} = 50$ ns), the digital discriminator threshold is set to 550, the RMS noise is 4.4 mV, the time jitter is ± 2 ns, and the strobe of the FADC is moved within ± 3 ns.

The algorithm identifies the wrong bunch-crossing as soon as sample points move across the threshold (see Figure 4.14 b)). That would first happen at the FADC saturation level for samples coming from above the threshold and at the higher energy limit for those samples coming from beneath the threshold. If the FADC strobe is moved further away from the peak without changing the threshold level, then we expect samples which cross the threshold. A mis-identified bunch-crossing at the FADC saturation level would lower the overlap region for non-saturated and saturated BCID, whereas mis-identified samples at very high energies are of course not acceptable.

To prove the flexibility of the algorithm, it was tested with pulses having a different pulse peaking time. For a shaper time constants of 5 ns the FADC digitisation strobe timing can be changed within ± 4 ns without any misidentified event. For a larger peaking time of 63 ns the range for the FADC strobe is reduced. In such

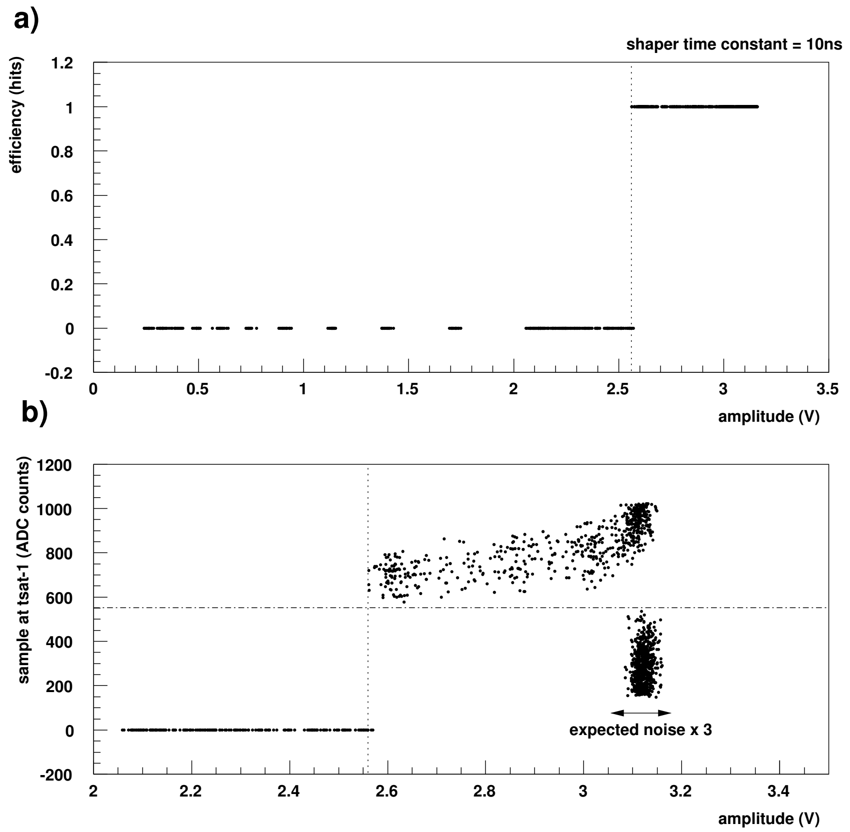

Figure 4.15: Efficiency plot a) and sampled FADC value before first occurrence of saturation at t_{sat-1} b). Shaper time constant is 10 ns (t_{pk} = 50 ns), the digital discriminator threshold is set to 550, the RMS noise is 13.2 mV, the time jitter is ±2 ns, and the strobe of the FADC is moved within ±3 ns.

a case it can only be changed without a wrongly identified event in a region from –2 ns up to 1 ns. An extension of this FADC strobe range can be achieved by including the threshold s_{low} applied to t_{sat-2}. The simulations done so far have only used the threshold s_{high} applied to t_{sat-1}. A second threshold s_{low}, set to 20 FADC counts, extends the FADC strobe range to ±3 ns. For further simulation results see [Pfe99].

Cable integration effect

Since signal integration by the cable is the only effect that causes a variation in the peaking time of the Pre-Processor input signal, one needs confidence that the cable simulation is close to reality. Therefore a laboratory test setup was used to measure the influence of a twisted-pair cable (type: Kerpen MegaLine 627 flex-4p). This cable has a bandwidth of 300 MHz and a propagation delay of 4.26 ns/m. Figure 4.16 shows the peaking time versus pulse amplitude of a trigger tower signal generated by an arbitrary function generator (AFG). The AFG waveform is used as input to a cable driver which contains clipping operational amplifiers (CLC 502/CL) in the same way as the PSPICE schematic. The input peaking time (45 ns) of the AFG signal is independent of its amplitude. A short cable of 0.86 m and the transmitter and receiver electronics increase the peaking time by 10–15 ns. A long 71 m cable increases the peaking time further by 8–12 ns for signals below the FADC saturation level. This corresponds to the length independent value (12 ns) introduced by the simulation model of a 70 m long cable.

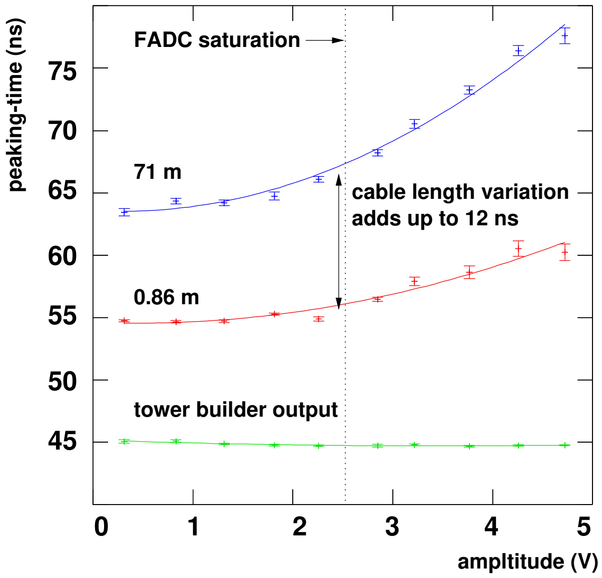

Figure 4.16: Peaking time versus tower builder output voltage [Hoe99]. The curves are drawn for different cable lengths (71 m, 0.86 m) to give confidence for the cable integration effect.

4.4 BCID summary

This chapter has described the simulation of a digital algorithm used for bunch-crossing identification of saturated trigger tower signals. Table 4.3 summarises the parameters investigated, for which the digital saturated algorithm has shown 100% efficiency. Table 4.4 gives a comparison between an analogue threshold BCID as described in [TDR98] and the digital algorithm described here. Measurements of the integration effect of a 71 m long twisted pair cable have shown the effect on the cable integration to the pulse peaking-time.

Parameter	Range
Shaper time constant	5 ns, 10 ns, 15 ns
Peaking time	43 ns, 50 ns, 63 ns
Noise	4.4 mV, 13.2 mV
Time jitter	$\sigma = 1$ ns of a Gaussian distribution (± 2 ns)
Mis-aligned FADC dig. strobe	6 ns (± 3 ns)
Energy range	256 GeV – 10 GeV
Saturation stages	lin. mixer, tower builder, layer sum

Table 4.3: Investigated parameters for which the algorithm has shown 100% efficiency.

	Saturated analogue BCID	Saturated digital BCID
Latency	no additional latency	no additional latency parallel to non-sat. BCID
Chips on PPr-MCM	additional analogue chip	included inside PPrAsic
Bonds per PPr-MCM	increases	does not increase
Efficiency overlap	overlap (20 GeV - 256 GeV)	no
Timing stability	17.5 ns	8 ns
Energy range	>20 GeV	>256 GeV
X-talk	sensitive to giga bit signals	no
Noise immunity	insensitive	insensitive
Time independent	independent of 40 MHz clock	tied to the 40 MHz clock
Costs	extra chip	no additional costs

Table 4.4: Comparison between analogue and digital saturated algorithm.

BCID timing strategy

For each channel an automatic process is needed to set the position of the FADC digitisation to the pulse maximum and to estimate the optimised threshold values for the saturated algorithm. One possible strategy is to use the Liquid Argon calibration system to generate a non-saturated pulse which is digitised with a default setting of the FADC strobe. A multi-pole fit function of the bipolar shaping is then applied to digitised data of the calibration pulse. From that fit one can calculate the correct FADC strobe setting and the optimised thresholds. Thresholds are calculated from two scaled versions of the fit function. One is scaled to a pulse that just has reached saturation, and one is large enough to give an estimation of a saturated pulse. From those functions one can calculate the thresholds by looking at the fictitious sampling points at t_{sat-1} and t_{sat-2}.

4.5 Bunch-crossing multiplexing

This section describes a data transmission scheme used to double the effective high-speed serial link bandwidth from the Pre-Processor to the Cluster Processor. The scheme makes use of the nature of the non-saturated BCID algorithm, which produces empty time slices after it has identified a bunch-crossing. The empty time slice can be used to transmit a value of a neighbouring trigger tower. Hence, the scheme is referred to as bunch-crossing Multiplexing (BC-mux). The scheme is independent of the actual link implementation. Either a G-link from Hewlett-Packard (HDMP-1012/14) or an LVDS[6] link from National Semiconductor (DS92LV1021/10) could be used to act as a transmitter/receiver chip-set.

Because of the large number of trigger tower signals (~7200) which must be transmitted to the Cluster Processor, a doubled link bandwidth has large implication for the architecture of the Level-1 Calorimeter Trigger. Transmitter/receiver chipsets are expensive in terms of power and cost. Heat dissipation is a limiting factor for the number of transmitters or receivers which can be placed on an electronics board. Hence, half the number of transmitters and receivers can reduce the system size and/or can contribute to producing a much more reliable system. Besides the required number of transmitter and receiver chips, connectors and cables represent a considerable part of the system cost. The attraction of the BC-mux scheme is, that it halves the number of transmitter and receiver chips, connectors, and cables. It has also been adapted to the η-duplicated one-slot connections between Cluster Processor Modules, where it has relaxed the track density of the backplane design in a Cluster Processor crate. The BC-mux scheme has been described in the Technical Design Report of the ATLAS Level-1 Calorimeter Trigger [TDR98].

4.5.1 BC-mux implementation

The BCID algorithm for non-saturated trigger tower signals, described in Section 4.2, consists of a peak finder with the condition $P_{t+1} \leq P_t > P_{t-1}$, where P_t is the digitised pulse amplitude of a trigger tower signal at time slice t. If that condition is not fulfilled, the peakfinder will produce empty bunch-crossings represented by a zero value. This leads to the fact that the bunch-crossing (25 ns) after a pulse maximum is always empty. Therefore, the maximum possible peak density is the alternation of a peak followed by an empty bunch-crossing. The BC-mux scheme makes use of the empty slices to transmit the bunch-crossings of a neighbouring trigger tower.

The bunch-crossing multiplexing will be performed by the Pre-Processor ASIC. It can be described with Verilog HDL. The digital input data to the BC-mux scheme are two 8-bit trigger tower channels, labled A and B. An additional flag bit must

[6]LVDS: Low-Voltage Differential Signalling

be generated to identify the transmitted trigger tower and the corresponding time slice. This flag bit is required at the receiver end to demultiplex the trigger tower channels. The total data to be transmitted are 9 bits, 8-bit data plus one flag bit for 2 trigger towers. Using the single frame mode of a G-link device, one can transmit 20+1 bits at 40 MHz, with a serial data rate of 960 MBd. This means that one G-link is able to transmit 4 channels of preprocessed trigger data. Using the double frame mode a G-link device, one could transmit 16+1 bits at 80 MHz, with a serial data rate of 1600 MBb (8 trigger tower channels). In this case the four flag bits must then be transmitted by an additional G-link per Pre-Processor Module. This would add further complexity to the board design and latency to the trigger system. Anyway, because of bit error measurements, the double frame mode of a G-link device is no longer considered as a reliable option. With LVDS links as an alternative, two preprocessed trigger towers could be transmitted over one chip-set.

	tower A (8-bit)		tower B (8-bit)		BC-mux output (8+1 bit)	
Case	BC_i	BC_{i+1}	BC_i	BC_{i+1}	BC_i,flag$_i$	BC_{i+1},flag$_{i+1}$
1	0	-	0	-	0,0	-
2	X	0	0	0	X,0	0,1
3	X	0	0	Y	X,0	Y,1
4	X	0	Y	0	X,0	Y,0
5	0	X	Y	0	Y,1	X,1
6	0	0	Y	0	Y,1	0,1

Table 4.5: Logic table for bunch-crossing multiplexing. The X stands for a 8-bit value of the trigger tower A. The Y stands for a 8-bit value of the trigger tower B.

Table 4.5 is a listing of all possible logic states of the BC-mux scheme. The first non-zero trigger tower data out of the pair A and B are sent out immediately, with a flag bit indicating whether it was tower A or B. On the second bunch-crossing (BC_{i+1}) the other tower is sent, with a flag bit indicating whether it was the same bunch-crossing (BC_i) as the first or one later. Note that the flag bit has a different meaning in the second bunch-crossing as in the first. The BC-mux of course needs time to demultiplex, but only on the receiving end. The receiver has to rearrange the data words. It has to buffer the first time slice until the second time slice arrives. This introduces 25 ns to the total trigger latency.

The BC-mux scheme is illustrated in Figure 4.17 for logic state 4 in Table 4.5. The figure shows the multiplexing of two trigger tower channels A and B into 8-bit + 1 flag bit. The buffering of the first bunch-crossing is shown as well as the de-multiplexing at the receiver.

Figure 4.18 shows the content of the 20+1 bit data word of the G-link device running at 960 MBd. Three error bits are possible to indicate errors during the data transmission. The G-link device provides a convenient intrinsic error check,

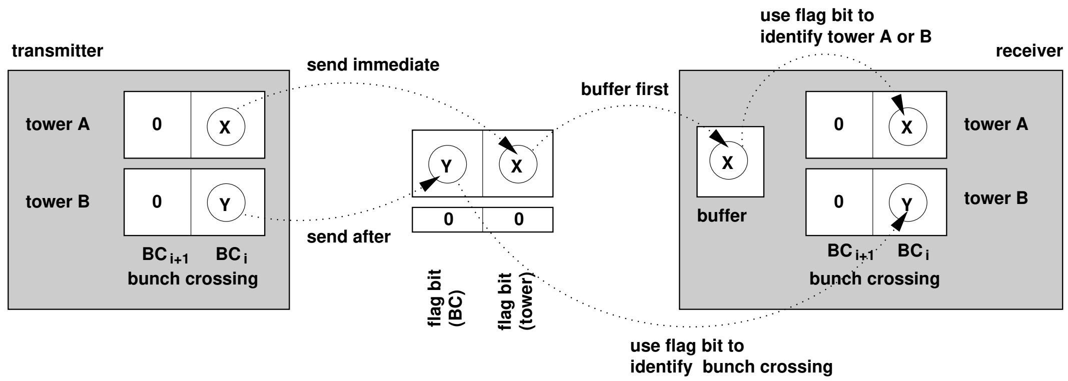

Figure 4.17: Illustration of the BC-mux scheme for logic state 4 of Table 4.5.

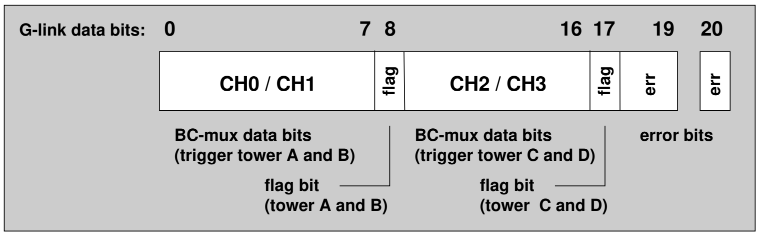

Figure 4.18: Data word format to transmit four channels with one G-link running at a serial data rate of 960 MBd.

but only on its protocol bits and not on the data bits. Therefore, the additional three data bits can be used as parity bits. Faulty event data can then be set to zero in order to avoid a false trigger due to an error.

Chapter 5

A compact Pre-Processor Multi-Chip Module

- MCM technology overview
- Aim and functionality
- Production technique
- Design and layout
- Thermal management
- Signal integrity
- System reliability

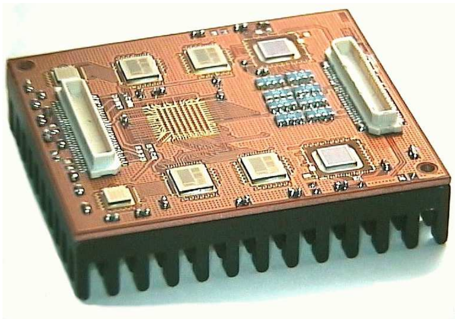

5.1 Multi-Chip Module technology overview

Multi-Chip Module packaging technology originally evolved out of hybrid microcircuit electronics. In hybrid circuits many electronic components, made from different semiconductor materials, are mounted on a single layer or multilayer substrate. In addition, a Multi-Chip Module is a packaging technology which encapsulates the entire system hermetically. The attraction of Multi-Chip Module technology, especially for large and complex systems, arises from increasing specialisation of semiconductor devices. A system often consists of a larger number of chips (dies), each optimised for its specific task. This makes it desirable to assemble many chips in a single package to reduce system size and external connections. Inside a Multi-Chip Module a very high internal pin-count can be implemented, with the package having a much reduced external pin-count. Input and output signals used as one-to-one connections between chips, e.g. wide digital buses, can be kept internal. Only signals like control, power supply and module input and output signals are used as package pins to interface to the outside printed circuit board. With Multi-Chip Module technology, a ratio of silicon-to-substrate area greater than 30 % can be achieved.

This section summarises the basic types of Multi-Chip Modules, classified by their substrate type and the production process used for the multilayer structure. The types are named: MCM-L, MCM-C, and MCM-D, where the characters refer to *laminated*, *ceramic*, and *depositing*.

- **MCM-L:** The substrate is based on *laminated* multilayer printed circuit board technology (PCB). It is an extension of chip-on-board technology (COB), where the bare dies are mounted directly on a substrate. Multilayer structures are formed by etching patterns in copper foils laminated to both sides of a resin based organic core. Layers are laminated together with one or more layers of basic resin in between to act as an insulator. Typical insulator materials are Polyimid or bismaleide-triazine (BT-epoxy). Core or substrate materials are often epoxy/fiberglass (FR4), aluminium or copper. Interconnections between layers may be formed by 'drilled' vias, which extend all the way through the board, 'blind' vias, which only extend from the surface part-way trough the board, or 'buried' vias, which connect only adjacent inner layers and do not extend to the surface in either direction. The smallest feasible layout structure in a design is referred to as feature size. In MCM-L technology this size can go down to about 60 μm.

- **MCM-C:** The substrate is based on confired *ceramic* or glass-ceramic materials. This technology evolved from thick-film hybrid technology. The process begins with thin sheets of unfired material which are referred to as *green tape*. The individual layers are printed with thick-film paste to create the metallisation patterns, aligned with the other layers, and are laminated at elevated temperature and pressure. Resistors and capacitors, compatible with the *green tape* process, may also be fabricated. The feature size depends on the thick film process, which has a typical conductor line width of 200 μm down to 100 μm.

- **MCM-D:** Interconnection patterns are created by *depositing* dielectrics and conductors on a base substrate of ceramic or silicon. The dielectric materials are used in the liquid state and are applied by spinning. Vias may be opened in the dielectric film by applying photoresistive material and etching, or by using photosensitive materials. The typical process which is used to deposit thin layers of dielectric materials is a plasma-enhanced chemical vapour deposition process (PECVD) to deposit layers of silicon dioxide for example. The fine-line lithographic technique used in MCM-D technology can produce very high-density interconnections using feature sizes of a few microns.

5.2 The aim of the demonstrator project

The aim of the demonstrator Multi-Chip Module (PPrD-MCM) was to establish MCM design techniques and experience for the final Pre-Processor. Therefore it had to include the preprocessing and the readout of the ATLAS Level-1 Calorimeter Trigger, for four trigger tower signals, in a single electronics module. The compactness and complexity of the final PPr-MCM is crucial to the architecture of the ATLAS Level-1 Calorimeter Trigger. In addition, its reliability is of importance for the running of the ATLAS experiment, since all data of the Calorimeter Trigger system have to go through this module before any event can be accepted. This section lists the tasks which were addressed by the demonstrator Multi-Chip Module project to give confidence for the final PPr-MCM design.

The objectives are as follows:

- The PPrD-MCM has to have a similar functionality as the final PPr-MCM. See Section 5.3 for a description.
- A suitable MCM packaging technology needs to be investigated.
- A compact layout, using a small feature size, is crucial to the system architecture.
- The mixed-signal MCM design needs to combine analogue and digital components.
- The MCM has to cope with high-speed serial signals in a frequency range similar to microwaves.

- An advanced bonding technique needs to be investigated. Flip-Chip bonding in combination with the MCM technology has to improve bonding yield and MCM compactness.
- The MCM reliability needs to be investigated in terms of yield and performance.
- The testability of the MCM must be investigated.
- A MCM design flow is required which can analyse the performance prior to manufacturing.
- Analysis results for signal integrity, maximum temperature, and electromagnetic influence are required, to prove the MCM performance.

5.3 Functional description of the PPrD-MCM

The boundaries of the PPrD-MCM were chosen at points of the processing chain where only few signals come in and out of the MCM package. The MCM has analogue and digital signals as inputs and outputs, respectively, and it includes different semiconductor devices such as mixed-signal and pure digital chips. Some of them are commercially available and others are application specific. The demonstrator Multi-Chip Module has a similar partitioning into dies as the final PPr-MCM. It processes four trigger tower signals using a prototype Pre-Processor ASIC (FeAsic[1]) developed at the ASIC laboratory of the University of Heidelberg [Fea96].

The tasks of the PPrD-MCM are:

- to digitise four analogue trigger tower signals at 40 MHz with 8-bit resolution;
- to preprocess each trigger tower data in terms of energy calibration and bunch-crossing identification;
- to serialise preprocessed trigger tower data using high-speed gigabit chip-sets.
- to provide deadtime-free readout of four trigger towers.

In order to achieve these, the MCM consists of:

- two dual flash ADCs from Analog Devices (AD9058, [AD]);
- four FeAsics, providing readout and preprocessing;

[1]FeAsic (Front-End ASIC) stems from a former name for the Pre-Processor.

- one Flip-chip INterCOnnection ASIC (Finco) for data multiplexing, signal-level conversion, digital-to-analogue signal conversion, in-circuit testing and temperature monitoring;
- either two Hewlett Packard gigabit transmitters (HDMP-1012, [HP]) serialising 16-bit at 40 MHz (640 Mbit/s) or, if the Finco chip performs data pre-multiplexing, 16-bit at 80 MHz (1280 Mbit/s).

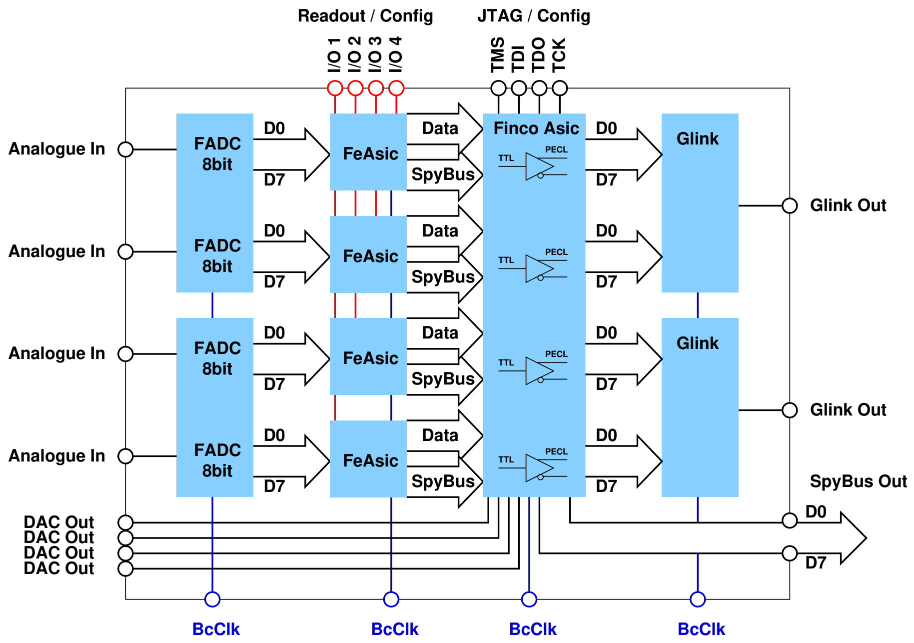

Figure 5.1: PPrD-MCM block diagram.

Figure 5.1 shows a block diagram of the PPrD-MCM. From that figure, one can see the scope of the Multi-Chip Module, where wide parallel data buses are kept internal and only analogue input, control, and the high-speed serial data signal come out of the MCM package. Each component has its own clock input, which allows an adjustment of the clock phase to ensure valid data at the following stage.

The real time signal processing flows from the left to the right. First, the FADC digitises the analogue trigger tower signals at 40 MHz to 8-bit precision in a range from 0 to 2 V. Each FADC die generates its own reference voltage, shared between the two internal FADCs. A dual FADC die has the advantage that the number of analogue components is reduced and a higher integration factor on the ASIC level is achieved, e.g. common reference voltage circuits. Next, the four 8-bit data buses are each digitally preprocessed inside a FeAsic. The FeAsic output is interfaced to the G-link dies via the Finco ASIC. This ASIC is a mixed-signal design, specific for the demonstrator Multi-Chip Module. It contains sixteen 2×1 multiplexers to double the normal operating rate of the G-links to 1600 MBd.

In this configuration one G-link transmits four trigger tower channels instead of two. Both G-links are then operating parallel to each other, producing a fanout of the serial data stream. The Finco ASIC is also required for level conversion of the FeAsic driver technology from TTL[2] to PECL[3], since the available[4] G-link dies have no TTL inputs. The Finco ASIC provides four digital-to-analogue converters (DACs) with a resolution of 10-bits. They are required to adjust the baseline of the analogue line receivers, which are placed outside the MCM. A JTAG interface provides boundary-scan I/O, to preload defined pad states and to scan the values received from the FeAsic. This is a very useful feature for testing wire connections between Multi-Chip Module dies. A temperature sensor is included on the Finco substrate to measure the temperature inside the PPrD-MCM. For a detailed description of the Finco ASIC see Section 6.

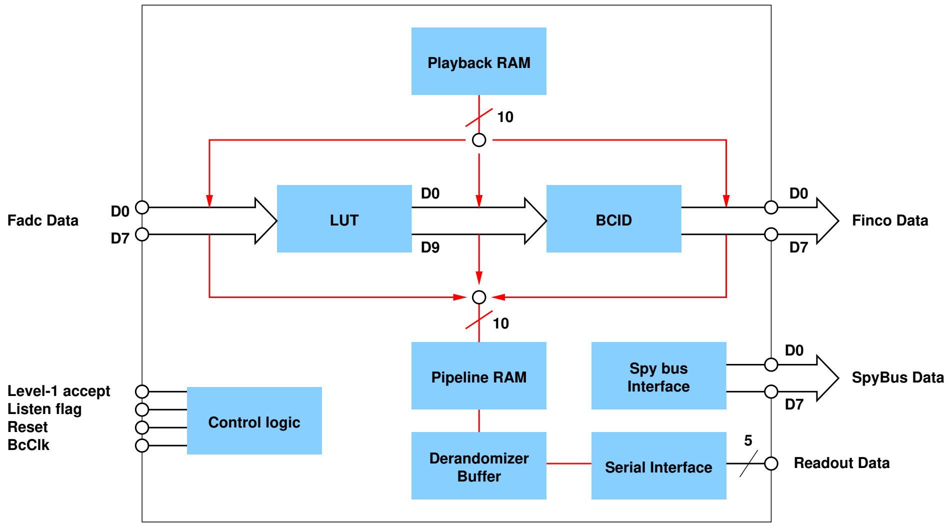

Figure 5.2: FeAsic block diagram.

An FeAsic block diagram is shown in Figure 5.2. The direct data path contains a lookup table (LUT) for data calibration and a Bunch-Crossing Identification unit (BCID) for unsaturated trigger tower signals. The BCID unit consists of a finite impulse response filter (FIR filter) with fixed coefficients (–1, 0, 4, –1) and a peak finder. The readout data path can record data in two conceptually different ways. Either it can record 128 time-slices of 25 ns in the so-called 'stop' mode or it can record data, deadtime free, using a 'fast' readout mode. The 'fast' readout mode does not stop the scrolling of pipeline memories for readout. It supports data readout of 1, 4, or 8 time-slices, where no deadtime was introduced up to a level-1 accept rate of 128 kHz for four time-slices and up to 64 kHz for eight time-slices.

[2] Transistor Transistor Logic
[3] Positive Emitter Coupled Logic
[4] Hewlett Packard has now changed its gigabit transmitter/receiver chip-set (HDMP-1012/14) to one which has TTL compatible input and output levels (HDMP-1022/24).

A detailed test of the performance of this readout concept was presented at the Third Workshop on Electronics for LHC Experiments (see [LEB97]). The FeAsic has the capability to 'spy' on internal registers at the full system clock speed of 40 MHz. These 'spybus' data, 8-bits wide from each FeAsic, are multiplexed by the Finco into a single 8-bit bus connected to the external MCM package pins.

5.4 The PPrD-MCM production technique

For the processing of the large number of channels in each Pre-Processor Module, a laminated MCM-L technique was chosen to combine small feature sizes with low prices. The design process of the laminated multi-layer structure is based on an industrially-available production technique for high-density printed circuit boards. The process, which is offered by Würth Elektronik [Wue] is called DYCOstrate®. It is characterised by its use of plasma etched micro-vias, where plasma is used for 'dry' etching of organic insulating materials such as Polyimide. Radicals and ions from a plasma react chemically with organic molecules accessible through open areas in a layer mask. Reaction products are removed instantaneously during the etching process, because of the reduced pressure and density of the plasma phase. Plasma etching enables precise via contacts between layers with a diameter of 100 μm down to 50 μm. The number of vias which may be etched at once is not limited and does not effect the complexity of a design and its price.

A reduced via diameter is an important parameter when designing Multi-Chip Modules with a large number of interconnections. Figure 5.3 shows a comparison of a micro-via with conventional vias used for printed circuit boards. It illustrates the increased routing density which stems from a smaller via size.

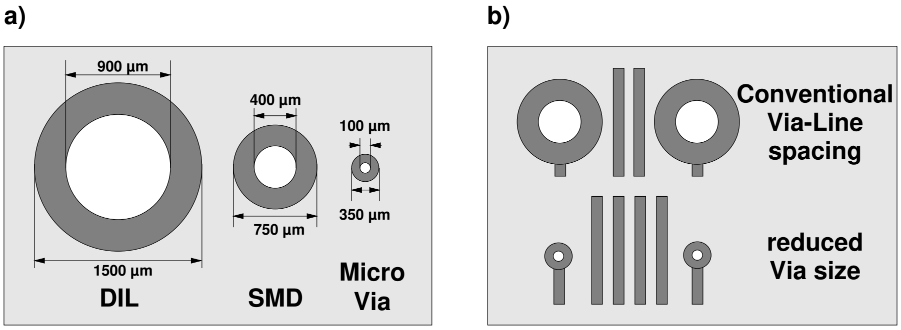

Figure 5.3: A comparison of micro-vias with conventional drilled vias is given in (a). A reduced via size increases the wiring density even if the line width stays constant (b).

The process of plasma etching can be used either within a surface-mount pad, or even in a pad suitable for Flip-Chip bonding, as investigated for the Finco ASIC of the demonstrator Multi-Chip Module. The reduction gained in system size is illustrated in Figure 5.4.

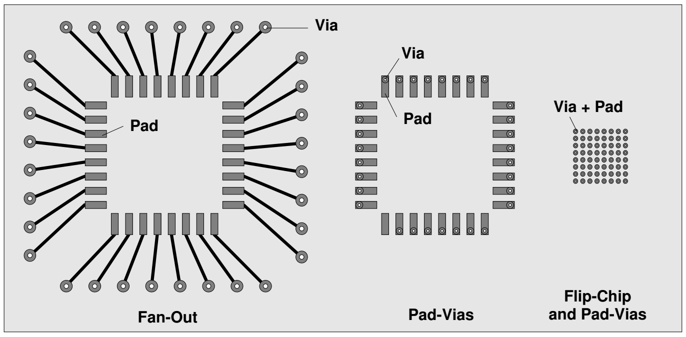

Figure 5.4: Reduction of packaging size due to vias placed in SMD pads. The approach of vias in Flip-Chip pads, as used for the Finco chip of the demonstrator Multi-Chip Module, is illustrated as an example of very high-density packaging.

5.4.1 MCM layer structure

The body of the demonstrator Multi-Chip Module is a combination of three flexible Polyimid foils laminated on a rigid copper substrate to form four routing layers. Here a description of the layer cross-section is given in the way they are manufactured: an inner-foil of 50 μm thickness, also referred to as *core* foil, carries 18 μm copper plates on either sides. Plasma etching is used for 'buried' via connections to adjacent layers and routing structures are formed in copper using conventional etching techniques. The core foil is surrounded by outer foils of 25 μm Polyimid, which are copper plated only on one side. Figure 5.5 shows a side view of each flexible foil prior to laminating. The actual contact through the core foil is accomplished with electroplated copper and after that, the routing structure are formed (see Figure 5.6). The electroplating process increases the track thickness from 18 μm to 25 μm.

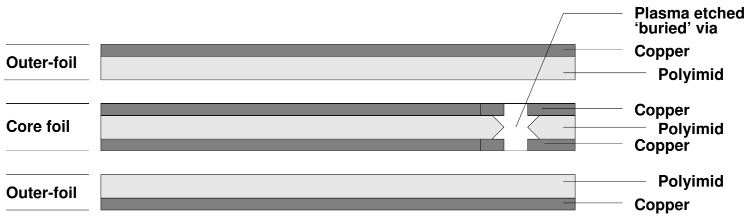

Figure 5.5: Flexible part of a four layer DYCOstrate® MCM-L, showing each Polyimid foil individual prior to laminating.

The apply of adhesive accomplishes laminating. Then insulating Polyimid is removed by plasma etching above a 'target' pad to form a surface 'blind' micro-via contact for the 'outer' foils. After electroplating of copper, routing structures for the top and bottom layer are etched to form final connections. Figure 5.7 shows the layer cross-section before electroplating and Figure 5.8 shows the final laminated

Figure 5.6: Cross-sectional view of the core layer after electroplating of 'buried' vias and etching of routing structures .

and flexible part of the MCM after etching of outer routing structures.

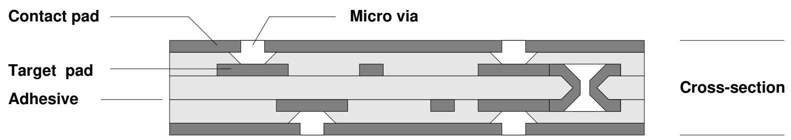

Figure 5.7: Laminated cross-section shows plasma-etched micro-vias for the top and bottom layer.

As shown in Figure 5.8, a combination of three vias is needed to accomplish a contact from the top to the bottom layer. A 'blind' surface micro-via connects layer 1 to layer 2, a 'buried' via connects layer 2 to layer 3, and finally a surface micro-via connects layer 3 to its destination on layer 4. Drilled vias are avoided, because they would need a more rigid and therefore thicker core foil, which would unnecessarily increase the total thermal resistance of the design.

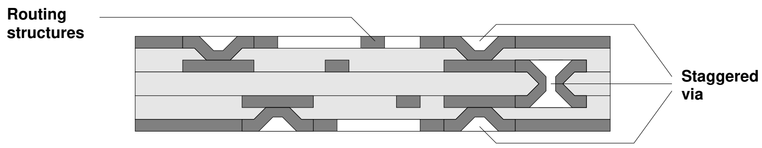

Figure 5.8: Final cross-section, showing staggered vias which are used instead of drilled vias for connections through all layers.

The flexible multilayer is further processed by milling of predefined cutout regions. Finally it is glued onto a copper substrate (see Figure 5.9). A cutout region is necessary for the the high-power G-link die, to ensure optimal thermal contact with a low junction temperature and low failure rate. See Section 5.7.2 for a detailed calculation of the improved thermal conductivity, which arises from cutouts.

The G-link die can be glued directly onto the copper substrate through its cutout. It is connected to bond pads on the top layer of the cross-section using a standard ultrasonic wire-bond technique.

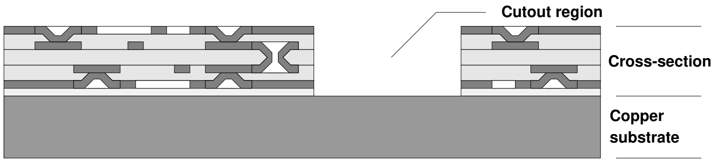

Figure 5.9: Laminated layer cross-section on top of a copper substrate. A cutout region is shown as it is used for the high-power G-link die.

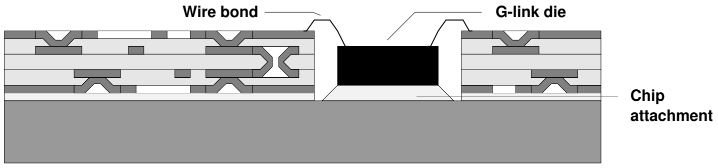

Figure 5.10: Improved thermal contact, achieved by a cutout region, for the high-power G-link die. The chip is glued right onto the copper substrate and bonded using a standard ultrasonic wire-bond technique.

Components such as capacitors and resistors are connected to the multi-layer structure using surface-mount technology. Advanced Flip-Chip mounting was investigated for a further reduction of the bonding space occupied. Staggered vias are grouped as close as possible to form thermal vias, which improve heat conduction to the substrate. Figure 5.11 illustrates the final build-up before the Multi-Chip Module is hermetically sealed by encapsulation materials.

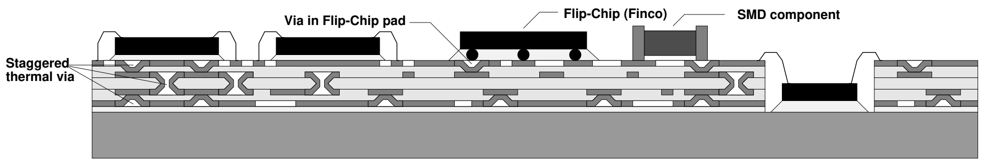

Figure 5.11: Illustration of different bonding technologies investigated for the demonstrator Multi-Chip Module. Thermal vias, SMD components, vias-in-pads and Flip-Chip pads are shown.

Due to the large number of semiconductor devices and Multi-Chip Modules at the Pre-Processor board level, a high-density SMD connector was investigated for the demonstrator Multi-Chip Module to permit a quick replacement of a broken MCM. If built-in tests have identified a broken MCM, it can be replaced without any soldering. Figure 5.12 shows the final build-up of the demonstrator Multi-Chip Module. Each chip is encapsulated individually, and finally an elastic encapsulation material is used to absorb stress arising from the use of different materials and to hermetically protect all components from their environment.

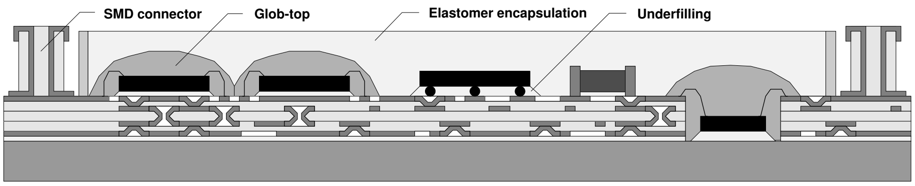

Figure 5.12: Side view of the hermetically sealed demonstrator Multi-Chip Module. High-density SMD connectors were used to allow quick replacement upon component failure.

5.4.2 MCM feature size and design constraints

The DYCOstrate® MCM-L technology is offered by its manufacturer with no relation to any specific physical layout tool. The company does not support a software product with libraries containing information on materials, cross-section, or design constraints. The customer is free to choose the layout tool he desires, and he has to ensure that all process parameters and constraints given by the manufacturer are set up properly. The layout tool is then able to check the layout against given constraints automatically, which is referred to as Design Rule Check (DRC).

For the PPrD-MCM the physical layout tool APD (Advanced Package Designer, [Apd97]) was used. It is a layout system for creating and optimising microelectronic packages, such as Multi-Chip Modules or Single-Chip Modules (SCM). APD is part of the integrated software that the software company Cadence [Cad] sells for accomplishing the major physical layout tasks of package design, which are: library development, placing parts, routing, and output generation for manufacturing. APD can also be used together with other Cadence products, namely: Concept [Con96] for schematic capture, DF/Thermax™ [The97] for thermal analysis, DF/SigNoise™ [Sig97] for signal noise analysis, and DF/EMControl™ [Emc97] for checking of electromagnetic interference rules. Following this product division, one can define constraints in each of these products, e.g. for the geometrical layout, for the maximum allowed module and component temperature, for the electrical signal integrity, and the electromagnetic interference.

Table 5.1 describes the applied layout constraints as they were arranged in conjunction with Würth Elektronik for the PPrD-MCM. A characteristic number of 100 μm was chosen as the minimum allowed feature size. No line or shape of the layout is allowed to be smaller than this number. The etching technique can produce smaller routing structures of 80 μm, but 100 μm was used to lower the MCM price and to stay within a more conservative range for which the DYCOstrate® process has shown best results. The minimum feature size applies to the line space and width of copper tracks as well as to the solder-mask space and width. For the distance of lines and shapes to cutout regions and for the milling tool radius, a size of 500 μm was defined.

Layout constraints	
Line to line space	100 μm
Line width	100 μm
Line to shape space	100 μm
Line to cutout space	500 μm
Solder-mask width	100 μm
Solder-mask space	100 μm
Cutout radius	500 μm

Table 5.1: Geometrical layout constraints as they were arranged in conjunction with Würth Elektronik for the PPrD-MCM.

Via type	Via hole ø	Via pad ø	Via target-pad ø
drilled via (through all)	600 μm	900 μm	900 μm
'buried' via (drilled)	350 μm	600 μm	600 μm
'buried' micro-via	100 μm	300 μm	300 μm
'blind' micro-via	100 μm	350 μm	350 μm

Table 5.2: Constraints of each via type of the PPrD-MCM.

Via constraints are defined by a hole diameter, to form the contact to an adjacent layer, and the diameter of a surrounding contact and target pad. Table 5.2 lists their values for drilled vias, 'buried' vias, 'buried' micro-vias, and 'blind' micro-vias. The main via type which was used for the demonstrator Multi-Chip Module is a 'buried' micro-via and a 'blind' micro-via. Both can be combined to form a staggered contact through all four signal layers.

Layer thickness	
Surface metallurgy (Au)	100 nm
Surface metallurgy (Ni)	5 μm
Copper track	25 μm
Outer Polyimid insulator	25 μm
Epoxy glue between layers	10 μm
Core Polyimid insulator	50 μm
Epoxy glue to substrate	50 μm
Copper substrate	800 μm
Aluminium heatsink	8000 μm

Table 5.3: Layer thicknesses, as used for thermal and electrical simulations of the PPrD-MCM.

Apart from the definition of alternating materials, their thickness is important for both thermal and electrical characteristics. Table 5.3 lists the thickness of each cross-section layer.

5.5 MCM design flow

On each 64-channel Pre-Processor Module, the smallest exchangeable part is a four-channel Multi-Chip Module. The number 4 defines the scope of the MCM: the number of dies, input and output signals, total power dissipation, and cooling technique. These definitions form the system requirements on which the Multi-Chip Module design flow is based. Figure 5.13 shows the design flow, as investigated for the design of the PPrD-MCM, and how it is envisaged for the final Pre-Processor Multi-Chip Module design.

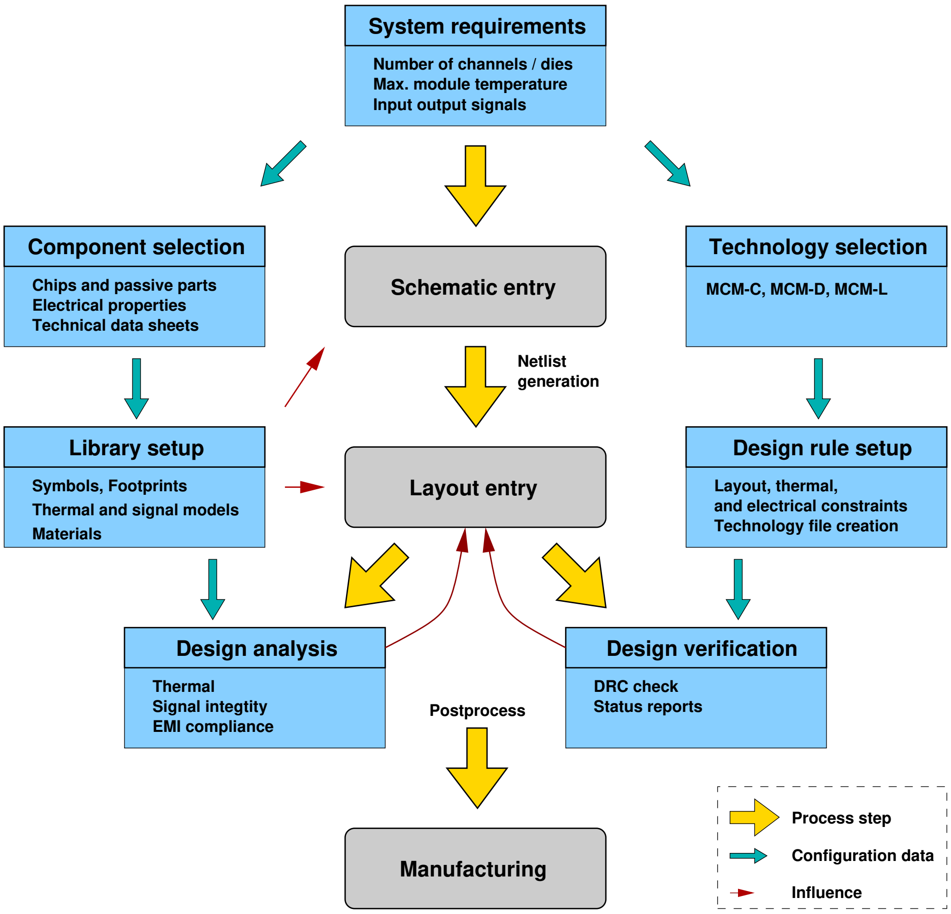

Figure 5.13: MCM design flow.

All tools which were used for the design are part of the Cadence product distribution. After the definition of the system requirements, the design flow continues with schematic entry. For the schematic entry Concept was used as the 'front-end' tool. Connectivity and component descriptions, derived from the schematic database, are forwarded via a design netlist to the Advanced Package Designer (APD). The software tool APD is used as physical layout tool, from which other

Cadence products can be started for design analysis, design verification, and for the final manufacturing output generation (postprocess).

Because of the lack of foundry-specified libraries and technology files, great emphasis has to be placed on the library development. At an early stage of the design, schematic symbol libraries are required. These libraries contain pin definitions for each schematic component, and electrical or physical layout properties, such as chip power dissipation or the maximum junction temperature. All properties which were defined for this 'Front-End' library will propagate to the physical layout tool. For the layout tool, package symbols containing footprints of the die bonding pads are assigned to each component. In addition to that, libraries used by the integrated analysis tools are required. Signal models to perform crosstalk and reflection simulations are needed from chip foundries or they must be measured or derived from ASIC simulations, as was done in case of the Finco chip. Thermal libraries, defining die attachment and packaging materials, are the basis for thermal analysis. Finally, design constraints specific for the selected MCM technology must be included (see Section 5.4.2 for details). Once all libraries are set up properly, design analysis can optimise the layout during die placement and routing. DRC checks and status reports of the design ensure that the manufacturing output does not violate the Multi-Chip Module specification.

5.6 MCM layout

This section describes the physical layout of the PPrD-MCM . It summarises various aspects of the design process, which can be used as a design guide for the final Pre-Processor layout.

Figure 5.14 shows an overview photo taken of the MCM top layer prior to any further processing. One can see the copper shapes for the die attachment, SMD pads, bonding pads, and the top ground shielding. The chip location is, from left to the right: FADCs, FeAsics, Finco, and G-links. Two trigger tower channels are processed by the top chip row and the other two by the bottom row. The Finco die sits between the two SMD connectors. Each connector has 60 pins. The left one carries only analogue, power, and slow control signals whereas the right one carries digital, clock, and high-speed serial output signals. As a guide line, the following points have to be considered:

- **Analogue and digital parts:** Separate analogue chips from digital chips, and keep power and ground signals separate. This applies also to the signal routing.
- **Component placement:** Keep optimum distance to other components to achieve uniform heat distribution. Remined space needed for different com-

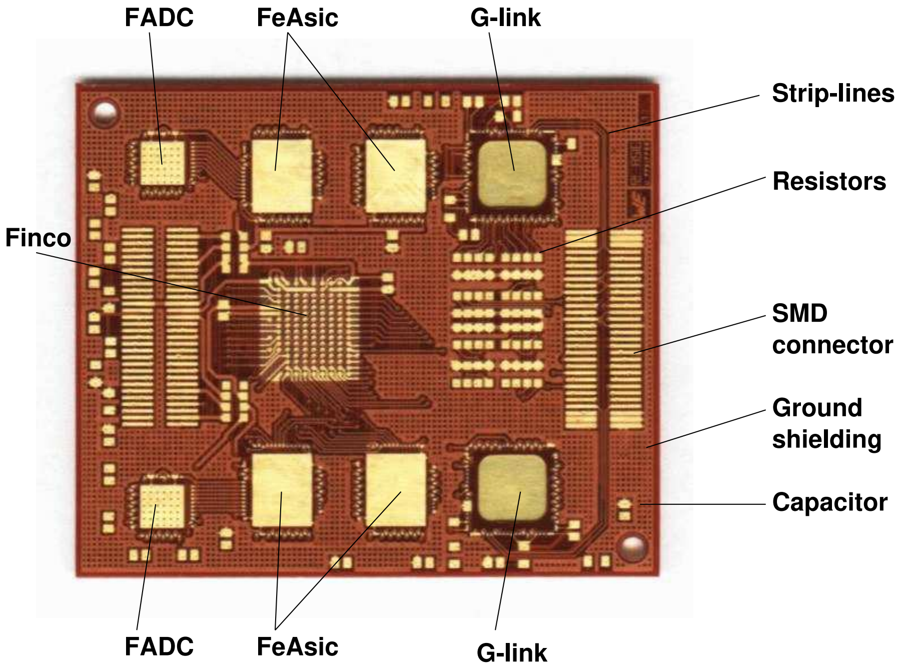

Figure 5.14: Top view of the MCM layout prior to any further processing.

ponent mounting technologies, e.g. wire-bonding, Flip-Chip bonding, and soldering.

- **Cutout region:** Figure 5.15 shows a picture of the G-link die attachment inside a cutout region prior to wire-bonding. When designing cutout regions, one has to provide bond pads on the surface to connect the copper base on which the die will be glued to its defined substrate potential. In case of the G-link die, a few wire-bonds for each chip were placed down through the cutout region onto the copper base to conduct the ground substrate potential.

- **Signal routing:** Because of the high signal density, no auto routing tool can be used. The layout is fully hand crafted, with two signal layers surrounded by a bottom ground plane and a top layer, consisting of shielding shapes, bonding pads, and SMD pads. One signal layer was mainly used for routing in x-direction and the other one for routing in y-direction.

- **Track width:** Use wider power tracks, especially for high power chips, to limit the voltage drop of power rails. For a copper foil having a resistivity of $1.724 \cdot 10^{-6}\,\Omega$m and a thickness of 25 μm, the resistance per square $R_\Box$ is 0.069 $\Omega/\Box$. The resistance of any track with constant width w is equal to the number of rectangular squares with size $w \times w$ multiplied by $R_\Box$. For example, a typical power track is 350 μm wide, 1 cm long, and it carries a third of the G-link current (133 mA). Such a track has resistance 1.97 Ω and a voltage drop of 0.26 V.

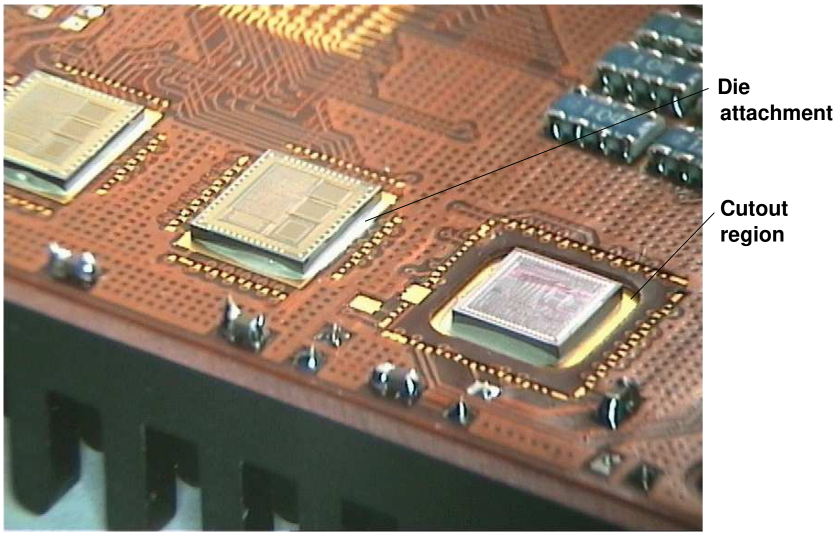

Figure 5.15: Die attachment on surface and inside a cutout region.

- **Chip power supply:** Use separate power tracks for each chip type and connect those tracks separately to the external power pins of the MCM package. For all high power chips, e.g. FADCs and G-links, single external power connections are preferred to lower a voltage drop caused by ohmic contacts.
- **Clock distribution:** Ensure a uniform propagation delay for clock signals.
- **Via placement:** A via can carry the same current as a copper track of 100 μm width. Hence, in wider power rails the number of vias must be increased.
- **Via count:** Reduce the number of vias for high-speed signals. A via, which connects adjacent layers of different impedance, can cause reflections affecting high-speed signals. For a simulation of the via influence see Section 5.8 for details.
- **Bond-pad size:** For the demonstrator MCM, a bond-pad size of the smallest feature size of 100 μm $\times$ 100 μm was used. This size is large enough for wire-bonding, but if one needs to probe at bonding pads during the MCM test the length should be at least twice as long.
- **Flip-Chip pads:** The via-in-pad technology was used for Flip-Chip pads of the Finco footprint. This pad layout has reached the maximum possible routing density for the DYCOstrate® MCM-L technology. The pads were arranged in a two-dimensional matrix as shown in Figure 5.16. Inner pads have contact through micro-vias to deeper routing layers, whereas the outer pad ring has contact only to the surface. The use of micro-vias for Flip-Chip pads has made it necessary to fill the via hole before reflow soldering of the

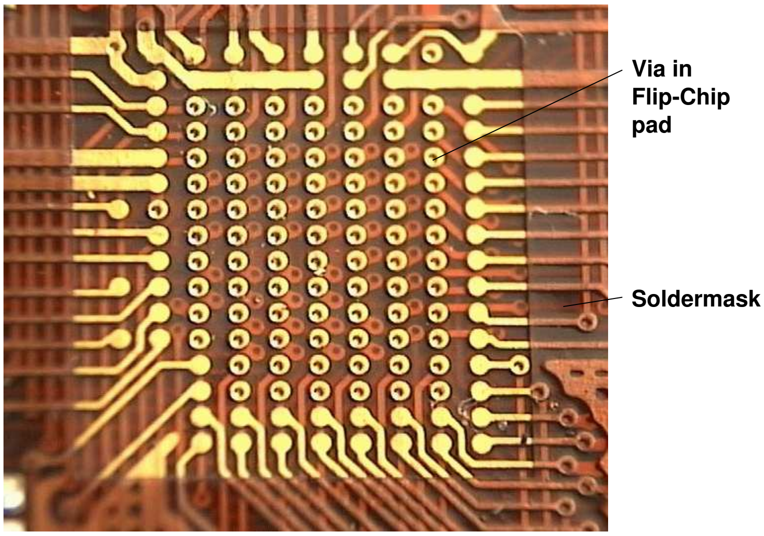

Figure 5.16: Finco footprint for Flip-Chip mounting.

Finco die, and to use a solder-mask for the outer pad contacts. See Section 6.1.6 for a detailed description of the Flip-Chip mount process.

- **Test pads:** Place test pads on the top layer. During the MCM test, needle probes can then spy on signals which are otherwise buried in the layer cross-section. This is important during the MCM test, where each chip needs to be tested.

- **Thermal vias:** A thermal via is a kind of staggered via, which connects through the cross-section to provide good thermal conduction. Figure 5.17 shows thermal vias placed in a copper shape used for the FADC die attachment. Attention must be paid to chip mounting on thermal vias, because the silver glue needs to fill up all micro-via holes to remove air between chip and copper shape. When the MCM is heated up with air underneath, the attachment will pop up, which is often referred to as 'popcorn' effect.

- **EMI shielding:** Use at least one filled ground plane in the layer cross-section as an electromagnetic shielding layer. On the top layer, use a cross-hatched ground shape surrounding bonding and SMD pads to improve the shielding further. This reduces the electromagnetic influence of signals to each other and it stabilises the ground potential. A cross-hatched shape is needed because drying moisture coming out of the cross-section can destroy the MCM (see Figure 5.18).

- **Chip attachment:** Copper shapes beneath each die are required to connect the die substrate with its voltage potential.

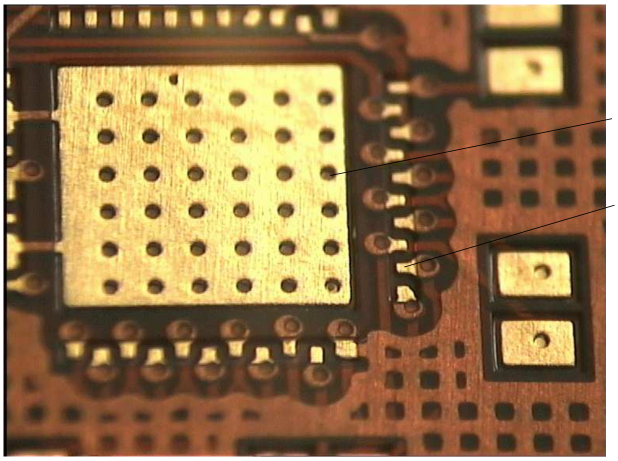

Figure 5.17: Thermal via layout.

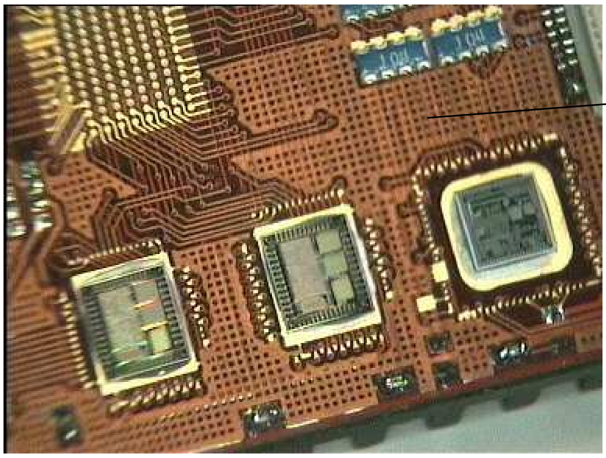

Figure 5.18: EMI shielding on surface, FeAsics, G-link, and Finco footprint.

- **Solder-mask:** A solder-mask can prevent short circuits during soldering of SMD components, and it protects bonding pads from the die attachment glue. It is also required for Flip-Chip pads to block solder at stubs connecting Flip-Chip pads.

The final PPrD-MCM is shown in Figure 5.19 after glob-top encapsulation of individual chips and prior to final hermetic encapsulation. The layout has a form factor of 4.3 cm $\times$ 3.7 cm enclosing an area of 15.9 cm^2. The total height is 1.21 cm including heatsink and SMD connectors.

The amount of silicon area is 80.15 mm^2 for 9 dies, resulting in a ratio of 5 % for the silicon to substrate area. This number is small compared to 30 %, which in general is considered to be achievable for MCMs. The reason for this is the inclusion of required SMD components on the MCM substrate. If SMD components were to be

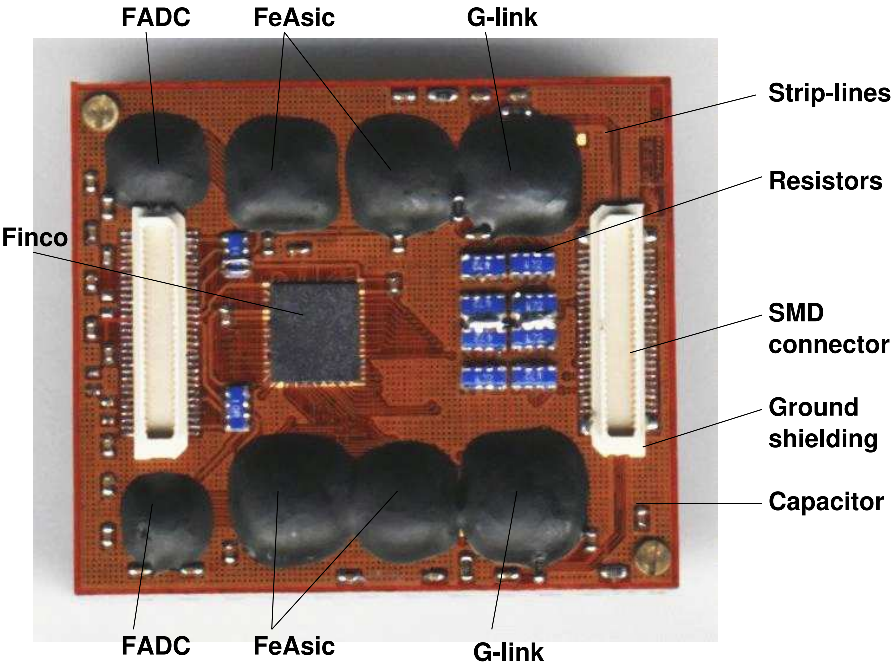

Figure 5.19: PPrD-MCM after glob-top encapsulation of individual chips and prior to final hermetic encapsulation.

placed outside the MCM they would of course increase the occupied board space further. The SMD connector pin-count is 120, whereas the internal pad count is 613. A total of 1380 vias were used for 271 signal nets with a total track length of about 5 m. See Table 5.4 for a summary of layout properties.

5.7 Thermal Management of the PPrD-MCM

One of the most challenging tasks during the Multi-Chip Module design process is the thermal management. It includes the selection, analysis, testing, and verification of a design for the purpose of producing a reliable end product. The increasing emphasis on thermal management in electronic designs stems from the large number of microelectronic devices in systems, their high heat densities, and the exponential nature of component failure rates with temperature. The classic means of cooling components through natural and forced convection air cooling is no longer a satisfactory solution to advanced packaging technologies. Cold plates, cutouts, thermal-vias and other enhanced techniques are needed in order to lower temperatures so that reliability goals can be achieved. At an early stage of the Multi-Chip Module design process, however, the application of basic heat transfer principles can indicate design problems and give assistance in selecting the most suitable Multi-Chip Module technology.

Layout Properties	
Substrate dimension (x, y)	4.3 cm, 3.7 cm
Package dimension (x, y, z)	4.3 cm, 3.7 cm, 1.21 cm
Substrate area	15.9 cm^2
Silicon area	80.15 mm^2
Ratio silicon/substrate area	5 %
Die count	9
Number of layers	4
Multilayer and substrate thickness	1070 μm
Total power consumption	max. 12.25 W, typ. ~9 W
Total thermal resistance R_{jc}	7.3 °C/W
Connector pin count	120
Number of nets	271
Number of vias	1380
Number of wire-bonds	470
Number of Flip-Chip bonds	143
Total bond count	613
Thermal vias	72
Total line length	~5 m
Impedance	93.7 Ω simulated
Resistance per square $R_\square$	0.069 Ω/□

Table 5.4: MCM Layout Properties.

The following section reviews the basic one-dimensional heat flow theory as applied to electronic packaging. Next, based on this theory, the total thermal resistance (chip-to-case) is calculated for the most power-consuming components of the demonstrator Multi-Chip Module. This is followed by an introduction to computer-based simulations, used to simulate three-dimensional heat transfer through a Multi-Chip Module multilayer. Finally, thermal measurements are presented which verify the computer-based simulation approach.

5.7.1 Basic heat flow theory

The process of heat transfer occurs between two points as a result of a temperature difference between them. Thermal energy may be transferred by three basic modes: conduction, convection and radiation.

Conduction

Heat transfer by conduction in solids occurs whenever a hotter region with more rapidly vibrating molecules transfers its energy to a cooler region with less rapidly vibrating molecules. The fraction of thermal energy ΔQ transported within a time interval Δt is called heat flow ($\frac{\mathrm{d}Q}{\mathrm{d}t}$). The fraction of heat flow which goes through a surface element ΔA is called heat flow density ($\frac{\mathrm{d}^2Q}{\mathrm{d}A\mathrm{d}t}$). In solids the heat flow density is proportional to the temperature gradient ($\frac{\mathrm{d}T}{\mathrm{d}z}$), which is orthogonal to the surface element ΔA. The proportional factor is a constant of the material called thermal conductivity (κ). This proportionality is expressed by the one-dimensional heat transfer equation 5.1.

$$\frac{\mathrm{d}^2Q}{\mathrm{d}A\mathrm{d}t} = -\kappa\frac{\mathrm{d}T}{\mathrm{d}z} \tag{5.1}$$

In case of a single solid plane, used as thermal contact between two temperature regions T_1 and T_2, with constant parallel surfaces A and thickness L, equation 5.1 can be written as:

$$\frac{\dot{Q}}{A} = -\kappa\frac{T_1 - T_2}{L}. \tag{5.2}$$

If we define a thermal resistance $R = \frac{L}{\kappa A}$, equation 5.2 can be written equivalent to ‘Ohms-Law’:

$$\dot{Q} = -\frac{1}{R}(T_1 - T_2) \tag{5.3}$$

or just

$$\Delta T = R \cdot \dot{Q}.$$

If we assume that a chip surface is heated uniformly, which leads to a constant temperature on the die surface, and if we assume that the heat flow path through a multilayer of homogeneous materials does not vary with time, then the temperature difference between the junction of a chip and its case can be calculated simply by adding up the thermal resistance of each layer of the cross-section and multiplying it with the heat dissipation P of the chip:

$$R_{jc} = \frac{L_1}{\kappa_1 A_1} + \frac{L_2}{\kappa_2 A_2} + \cdots + \frac{L_n}{\kappa_n A_n},$$

$$R_{jc} = \sum_{i=1}^{N} R_i,$$

$$\Delta T_{jc} = R_{jc} \cdot P. \tag{5.4}$$

Figure 5.20 shows an example of the total thermal resistance (junction-to-case), calculated from the materials of a package cross-section. If the case is held at a constant temperature, one can easily calculate the expected increase of the junction temperature from Equation 5.4. Some materials and their thermal conductivities are given in Table 5.5 for conductors and in Table 5.6 for insulators.

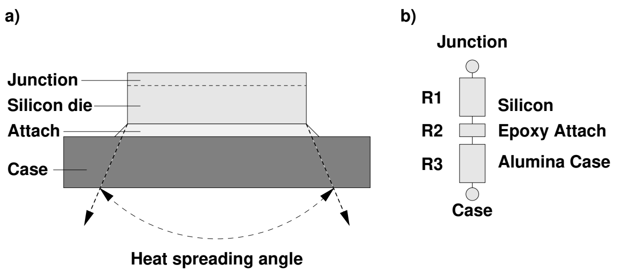

Figure 5.20: Thermal resistance chain for a single chip package, used to calculate the total resistance R_{jc} from junction to case.

Material	$\left[\frac{W}{m°C}\right]$
Copper	395
Aluminium 1100 H18	218
Silicon	118
Alumina (99% Al_2O_3)	25

Table 5.5: Thermal conductivity of electronic packaging materials at 44 °C.

Convection

In general, two types of heat transfer by convection can be distinguished: free and forced convection. Heat transfer by free convection occurs as a result of a change in the density of the fluid, which causes fluid motion. The process of heat transfer resulting from fluid flow across a heated or cooled surface is called forced convection. The general equation defining the convective heat transfer, either free or forced, from a surface at a temperature T_1 into gas or fluid at a surrounding temperature T_2 is given by:

$$\dot{Q} = -\alpha A \cdot (T_1 - T_2), \tag{5.5}$$

where the variable α is the heat transfer coefficient of the surface A.

Material	$\left[\frac{W}{m°C}\right]$
Epoxy (conductive)	0.35 - 0.87
Polyimide	0.33
Epoxy (dielectric)	0.23
Solder-mask	0.21
Air	0.026

Table 5.6: Thermal conductivity of electronic insulating packaging materials at 44 °C.

Radiation

Radiation refers to the transfer of energy by electromagnetic wave propagation. The wavelengths between 0.1 μm and 100 μm are referred to as thermal radiation wavelengths. The ability of a body to radiate thermal energy at any particular wavelength is a function of the body temperature and the material characteristics of the radiating surface. The total energy radiated by an ideal radiator (blackbody) at any particular temperature is given by the area under its specific temperature curve, which it is equal to σT^4 where σ is the Stephan-Boltzmann radiation constant. Materials or objects that act as perfect radiators are rare. Most materials radiate energy at a fraction of the maximum possible value. The ratio of energy radiated by a nonblackbody to that emitted by a blackbody at the same temperature is called the emissivity ε.

The heat flow density ($\frac{\mathrm{d}^2Q}{\mathrm{d}A\mathrm{d}t}$), which is radiated by a surface element ΔA is proportional to T^4, where T is the temperature of the radiating surface ΔA. The heat flow density can be calculated by the following law, which is called the Stefan-Boltzmann law:

$$\frac{\mathrm{d}^2Q}{\mathrm{d}A\mathrm{d}t} = -\sigma\varepsilon T^4.$$

A radiating body also absorbs energy which is emitted by its environment. If the surface of the environment is the same as that of the radiating body, one can write the Stefan-Boltzmann law as:

$$\dot{Q} = -\sigma\varepsilon A(T_1^4 - T_2^4), \tag{5.6}$$

where T_1 is the temperature of the radiator and T_2 the temperature of its environment.

5.7.2 Calculation of the heat-flow

In the early stages of design, a spreadsheet can be used to calculate thermal resistance for each chip on a Multi-Chip Module. From these early calculations, one can choose the optimal cooling approach for a given Multi-Chip Module technology. Besides technology dependent feature sizes such as multilayer thickness, materials, and the number of interconnection layers, other chip parameters such as chip dimensions and power dissipation are important to determine the expected thermal parameters. In case of the FeAsic, Table 5.7 shows a spreadsheet used to calculate the contribution of each cross-sectional layer to the total thermal resistance (junction-to-case). The calculation is based on the one-dimensional heat flow theory for conduction as described in Section 5.7.1. A constant heat spreading angle of 26.6 ° was used to take an increased heated area under each die into account [Dav77].

Layer	Conductivity $[\frac{W}{m \cdot °C}]$	Area $[mm^2]$	Thickness $[\mu m]$	Resistance $[\frac{°C}{W}]$
Air	0.027	-	-	-
Silicon Die	118	16.43	350	0.18
Epo-Tek H20S	1.57	16,43	50	1.94
Copper Layer 1	395	16,62	25	0.004
Polyimide	0.33	16.71	25	4.53
Epoxy	0.03	16.82	10	0.54
Copper Layer 2	395	16.86	25	0.004
Polyimide core	0.33	16.95	50	8.94
Copper Layer 3	395	17.15	25	0.004
Epoxy	0.03	17.25	10	0.53
Polyimide	0.33	17.29	25	4.38
Copper layer 4	395	17.38	25	0.04
Epoxy	0.03	17.48	50	2.60
Copper Substrate	395	17.68	800	0.12
Air	0.027	-	-	-
Junction-to-case resistance R_{jc} [°C/W]				23.77
Total temperature rise [°C]				7.13

Table 5.7: Heat flow through the layer cross-section beneath the FeAsic die. A heat spreading angle of 26.6° was used to take an increased heated area for each layer into account.

In case of the FeAsic die, the calculated thermal resistance R_{jc} is 23.77 °C/W. From this number, one gets the temperature rise from junction to case by multiplying it with the chip power dissipation. The FeAsic power dissipation is 300 mW and hence, the temperature rise is 7.13 °C. Assuming, that the heatsink, which is attached to the copper substrate, is cooled uniformally at a temperature of 45 °C, the expected temperature for the FeAsic junction is 52.13 °C. This temperature is in an acceptable operating range to ensure a low component failure rate, as can be seen in Section 5.7.8, Figure 5.31.

The maximum allowed numbers for thermal resistance and junction temperature can be found in technical data sheets. The dual 8-bit FADC from Analog Devices (AD9058) and the gigabit transmitter from Hewlett Packard (HDMP-1012), used for the demonstrator Multi-Chip Module, have a maximum junction-to-case resistance of 12 °C/W. The maximum allowed junction temperature for the FADC and for the G-link is 130 °C.

For each chip on the demonstrator Multi-Chip Module, the combined thermal conductivity is illustrated in Figure 5.21. From that figure, one can see that the FADC cooling is improved by about 87 % by using thermal vias. Furthermore,

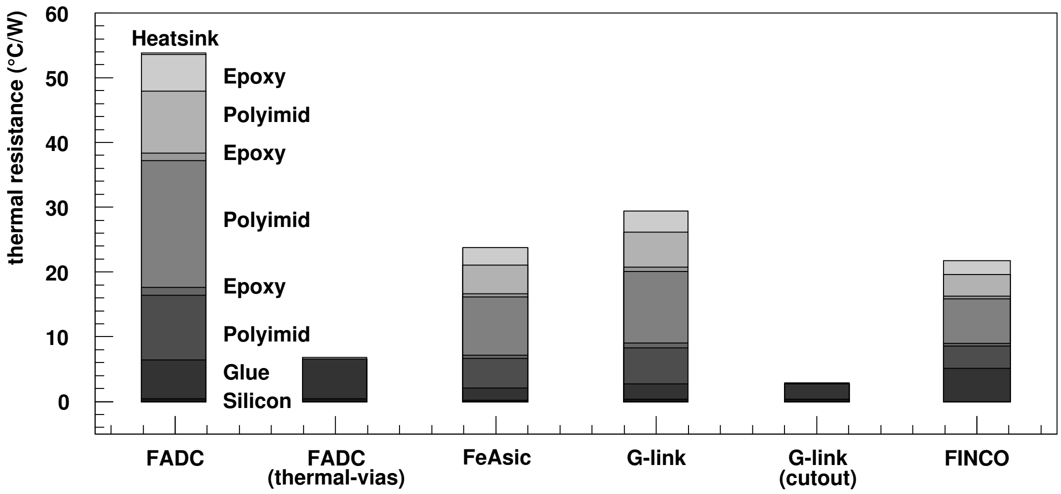

Figure 5.21: Combined thermal resistance for each chip on the demonstrator Multi-Chip Module, calculated using one-dimensional heat flow theory for conduction as described in Section 5.7.1. The cooling effect of thermal-vias and cutouts is illustrated by reduced bar chart height in case of the FADC and G-link.

the cutout technology used for the high power G-link die has improved the cooling mechanism by about 90 %. The heat flow through thermal vias was calculated by introducing a copper layer with a reduced area of 26 % of the original Polyimid insulating material. This corresponds to the area fraction occupied by micro-vias of 100 μm diameter. The attachment inside a cutout was calculated by using only the attachment material and the copper substrate as cross-section layers. Flip-Chip mounting of the Finco ASIC was calculated using 3.32 % of the Finco chip area occupied by solder bumps of 100 μm width and 80 μm thickness. Figure 5.22 illustrates the combined temperature rise for each chip, which is the product of the thermal conductivity and the power dissipation. The actual numbers for each junction temperature are given in Table 5.8.

There are not only advantages which arise from thermal via and cutout cooling approaches. There are also disadvantages, for example a loss of all routing area underneath each chip which makes use of such a cooling mechanism. This can lead to a further increase of the number of layers needed to route all connections within a given Multi-Chip Module package size.

5.7.3 Computer-based thermal analysis

The one-dimensional approach for solving heat transfer problems, as used to calculate the heat flow for each chip on the demonstrator Multi-Chip Module individually, is satisfactory in an early stage of the design process. However, in high temperature Multi-Chip Module designs a detailed temperature distribution is required to determine the overall Multi-Chip Module reliability and to prove the

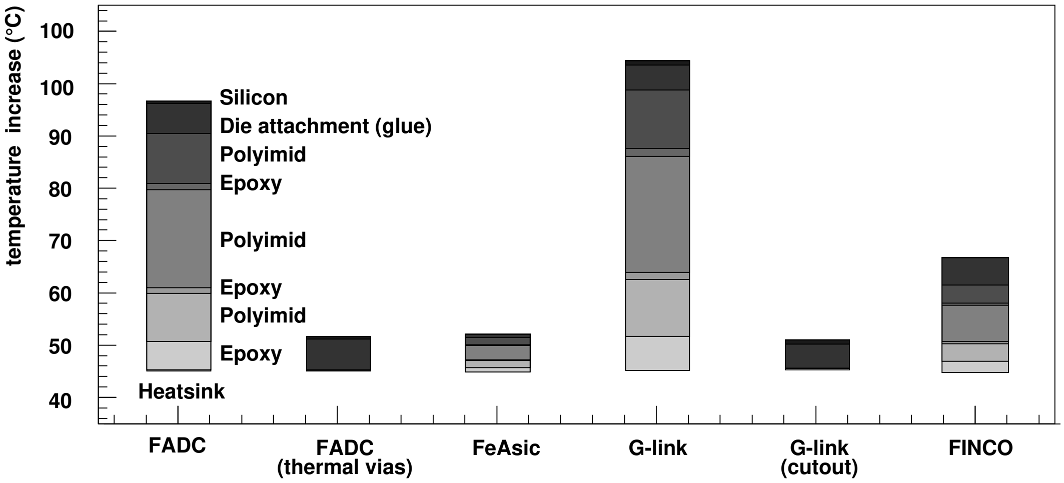

Figure 5.22: Combined temperature, calculated by using one-dimensional heat flow theory for conduction as described in Section 5.7.1. The cooling effect of thermal-vias and cutouts is illustrated by reduced bar chart height in case of the FADC and G-link.

Thermal parameter	FADC	FADC 26 %	FeAsic	G-link	G-link cutout	Finco 3.3 %
R_{jc} [°C/W]	53.81	6.84	23.77	29.47	2.89	10.88
Chip power [W]	0.96	0.96	0.30	2.02	2.02	2.0
HS temp. [°C]	45	45	45	45	45	45
Temp. rise [°C]	51.66	6.57	7.13	59.53	5.84	23,20
Junction temp. [°C]	96.66	51.57	52.13	104.53	50.84	68,20
Chip area [mm^2]	7.37	7.37	16.43	13.25	13.25	43.10

Table 5.8: Expected junction-to-case resistance and junction temperature, calculated for the FADC, FeAsic and G-link die. Thermal vias occupy 26 % of the FADC area and Flip-Chip solder bumps occupy 3.3 % of the Finco area. Expected numbers for the FADC and G-link without any cooling mechanism are given for comparison.

adequacy of the thermal design.

For analysing the thermal characteristics of the demonstrator Multi-Chip Module, the analysis tool DF/Thermax^{TM} [The97] was used. It is part of the Cadence product distribution [Cad] and it is integrated in the APD (Advanced Packaging Designer, [Apd97]). After a DF/Thermax^{TM} simulation is complete, one can view two- and three-dimensional maps that can be overlaid on the APD design. The algorithms implemented allow thermal analysis by finite differences. They are based on an iterative numerical technique, which involves the subdivision of the design into a grid network consisting of nodal elements having a specific volume and which are interconnected to adjacent nodes through a network of conductances. The thermal properties and temperatures are assumed to be uniform throughout each segment. The heat transfer through each surface of a node is expressed by either conductance or convection (see Equation 5.1 and 5.5), and radiation (see Equation 5.6).

The following section describes thermal setup conditions and results obtained by DF/Thermax^{TM} simulations of the demonstrator Multi-Chip Module.

5.7.4 Thermal simulations during the design process

Performing temperature analysis during the design placement and routing have optimised the thermal performance of the demonstrator Multi-Chip Module. The calculation of die temperatures has helped to find optimal die locations on the Multi-Chip Module substrate and the thermal behaviour as a function of the module environment has identified the operating temperature range. Heatsink sizes and materials were changed in order to find an improved cooling approach for a given environment. A transient analysis was performed to predict the thermal behaviour when the environment changes at a predefined point in time, for example when the system is powered up or a fan fails.

The DF/Thermax^{TM} simulator uses simplified information of the module build-up as input. It needs, derived from the detailed description given in Section 5.4, a simplified cross-section of materials arranged in adjacent layers, a description of the individual die attachment, and information about the geometry and the materials of the package encapsulation. Once all this information is provided, the boundary conditions must be defined. The layer cross-section and the package encapsulation with its boundary conditions are depicted in Figure 5.23.

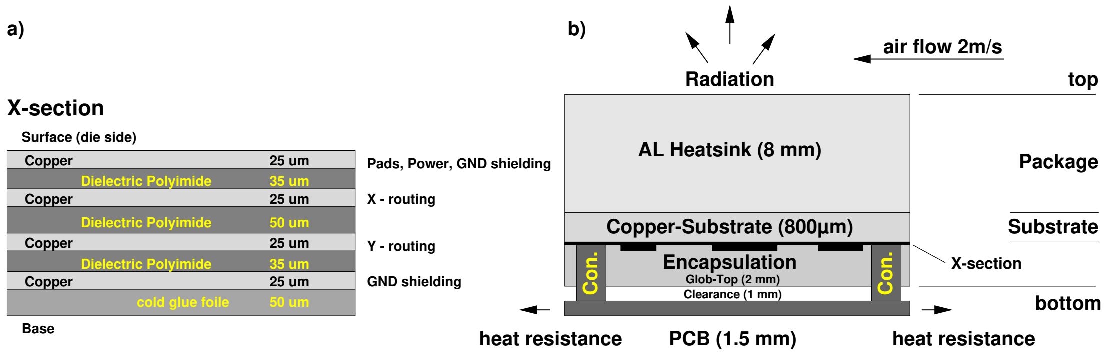

Figure 5.23: Simulation model used as input to the DF/Thermax^{TM} simulation. The cross-section is shown in a), and the package definition and the boundary conditions are shown in b).

For each die, material, size, and thermal resistance from the junction to the die attachment were defined and provided in a component library for simulation. Because of equal treatment of each component, the simulation cannot take special cooling approaches such as thermal vias or cutout regions of the cross-section into account.

Hence, the simulation results for the demonstrator Multi-Chip Module must be seen as a worst case system simulation to estimate maximum ratings, and to observe the thermal behaviour dependent on variation of boundary conditions.

Boundary conditions	
PCB edge guide resistance (North, South, East West)	5 °C/W
Constant temperature	25 °C
Initial temperature	25 °C
Force convection top	2 m/s
Force convection bottom	1 m/s
Radiation emissivity top	0.9
Radiation emissivity bottom	0.1
Distance to next MCM	10 cm

Table 5.9: Boundary conditions as applied to the MCM package.

Table 5.9 summarises the boundary conditions as they were used for the simulations. For the Pre-Processor PCB, on which the MCM will be mounted, an edge guide resistance of 5 °C/W was adjusted in each direction. The initial temperature, the temperature of the board surrounding, and the fan air temperature were set to 25 °C. Cooling with forced convection was simulated using fans blowing air over the package top with a velocity of 2 m/s, and through the clearance between the package bottom and the PCB at 1 m/s. This velocity is within the normal operation range of fans sitting underneath modules in VME crates.

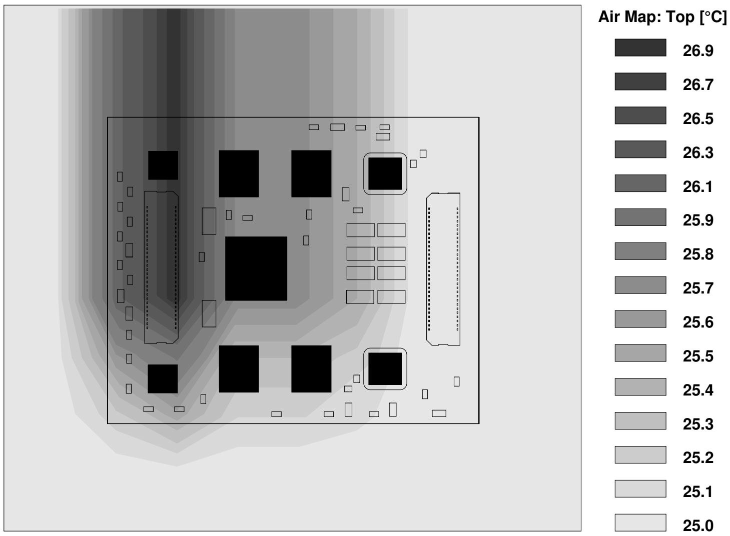

Figure 5.24: Temperature simulation (air map).

Figure 5.24 shows a two-dimensional air map above the package heatsink. A fan is blowing vertically from the bottom to the top and the air map is overlaid on the MCM layout to identify the component positions. From the initial air temperature of 25 °C the air is heated up by 1.9 °C to 26.9 °C. Its maximum value is reached above the FADCs and the Finco die, whereas it is less heated above the G-links.

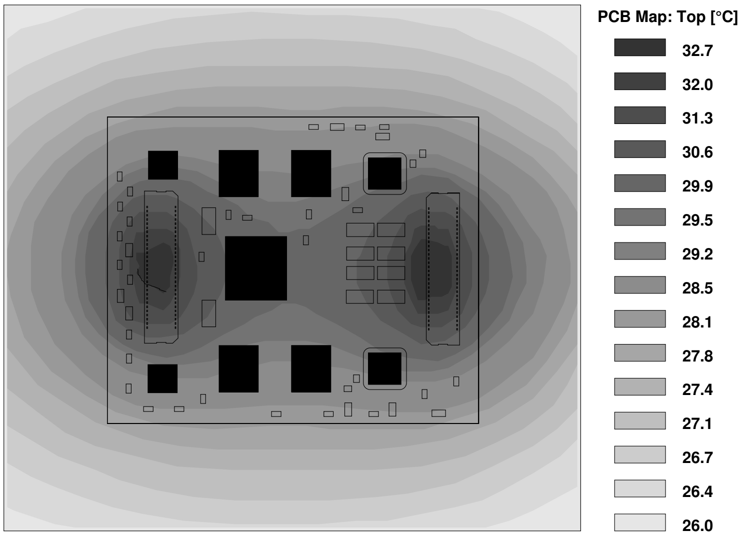

Figure 5.25: Temperature simulation (board map).

Figure 5.25 shows the heated area of the Pre-Processor PCB underneath the MCM package. The PCB is heated up from the initial themperature of 25 °C by 7.7 °C to 32.7 °C. The maximum PCB temperature is reached around the SMD connectors, where the SMD pins deliver the main contribution to the heat transfer.

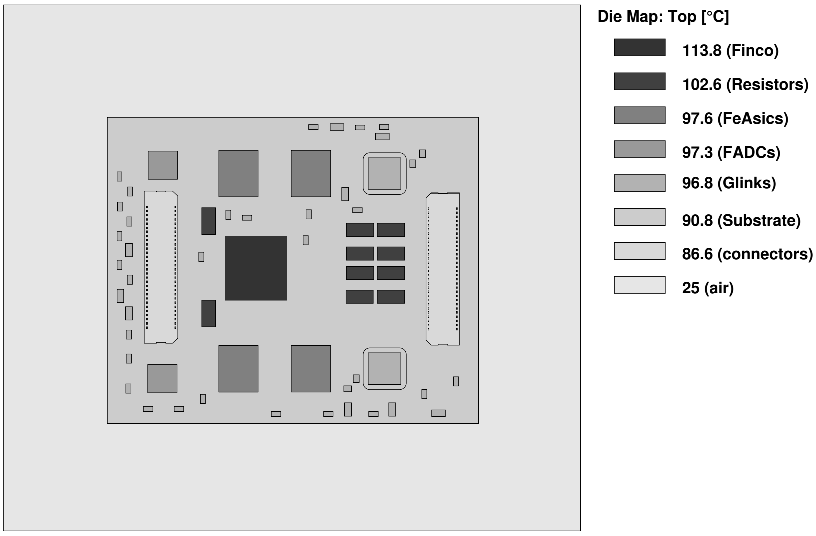

Figure 5.26: Maximum temperature simulation (die map).

A thermal map, showing the temperature of each component inside the MCM package, is given in Figure 5.26. The copper substrate and capacitors are heated up to a maximum temperature of 90.8 °C, whereas resistors, due to their power dissipation, are heated up to 102.6 °C. The maximum junction temperature is 97.6 °C for the FeAsics, 97.3 °C for the FADCs, 96.8 °C for the G-links, and the highest temperature is reached by the Finco die, which is 113.8 °C. All die temperatures are below their maximum operating temperature, as specified by technical data sheets.

Transient temperature simulation

A transient temperature simulation, starting at the time when the MCM is powered up for a duration of 40 minutes, is given in Figure 5.27. During that simulation constant fan-cooling was assumed, as given by the boundary conditions in Table 5.9. Afterwords a fit function of the form

$$T(t) = T_{max}\left(1 - \frac{2A}{1 + t/\tau}\right) \tag{5.7}$$

was applied using an error estimation of 5 % for the simulation result. The fit parameter T_{max} is the maximum junction temperature, the parameters A and τ can be used for a comparison of the transient temperature behaviour, where A gives the distance from T_{max} (in %) after the time $t = \tau$.

	T_{max} [°C]	A [%]	τ [min]
Finco	113.85	27.9	5.2
FADC	97.13	32.3	5.1
FeAsic	98.33	32.4	5.4
G-link	96.42	32.3	5.1

Table 5.10: Fit parameters of the transient junction temperature simulations.

In addition to a transient and a two-dimensional heat flow simulation, the DF/Thermax^{TM} analysis can be used to calculate a thermal model of the MCM package. This model can then be used in board-level simulations to simulate the thermal behaviour of the final Pre-Processor PCB. The junction-to-case resistance and the case-to-board resistance of the MCM package are listed in Table 5.11. This table also includes some maximum temperature values for which two-dimensional maps were generated.

5.7.5 Comparison between calculation and simulation

Table 5.12 gives a summary of the junction temperatures T_j for each die on the MCM. The temperatures were taken from the two approaches: the one-dimensional

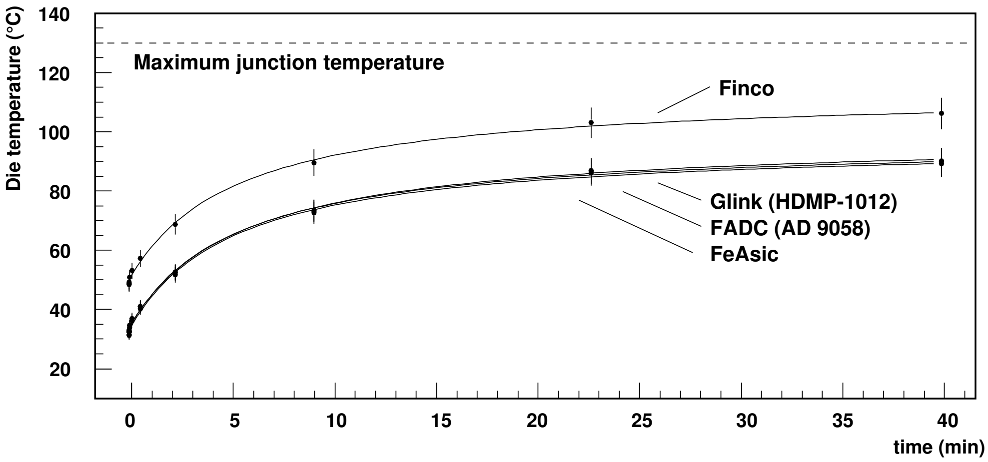

Figure 5.27: Simulated and measured transient die temperature.

Maximum thermal MCM package parameter	
Junction-to-case resistance R_{jc}	7.3 °C/W
Case to board resistance R_{cb}	28.4 °C/W
Total MCM power	12.25 W
Board temperature	31.13 °C
Clearance temperature	60.79 °C
Connector temperature	86.60 °C
Glob-top temperature	90.84 °C
Substrate temperature	90.80 °C
Heatsink temperature	90.80 °C

Table 5.11: Simulation results and model parameter for the MCM package.

spreadsheet calculation described in Section 5.7.2, and the two-dimensional simulation described in Section 5.7.4. The calculation assumed that the heatsink balance temperature is 45 °C, and the simulation assumed fan cooling with an air temperature of 25 °C and a velocity of 2 m/s. Results from the two-dimensional simulation were shown as die temperature maps or transient temperature curves.

The junction temperatures from the two-dimensional simulation are higher than the calculated temperatures, see Table 5.12. This is partly due to the inclusion of the heat dissipation of passive termination resistors and capacitors. Their heat dissipation was simulated as specified by their data sheets. Because of passive components, the total MCM power dissipation simulated was rather high (12.25 W). Measurements of the actual clocked MCM have shown that the total power dissipation is approximately only 9 W. Furthermore, the two-dimensional simulation does not allow the definition of a heatsink geometry. Only its dimension and its area can be defined. Therefore, the adjustment of the MCM boundary conditions are a critical issue for a two-dimensional simulation set-up.

T_j	One-dimensional calculation [°C]	Two-dimensional simulation die map [°C]	transient T_{max} [°C]
Finco	68.20	113.8	113.85
FADC	51.57	97.3	97.13
FeAsic	52.13	97.6	98.33
G-link	50.84	96.8	96.42

Table 5.12: Comparison of the junction temperature obtained from a spreadsheet calculation, a two-dimensional die map, and a transient temperature simulation.

The following section describes MCM temperature measurements, which can be used in future temperature simulations to optimise the heatsink definition and the boundary conditions. Then the MCM cooling will be simulate more realistic.

5.7.6 Transient temperature measurements

The Finco ASIC on the demonstrator MCM allows measurement of its own junction temperature. This is a convenient method to prove the calculation and simulation results. Figure 5.28 shows the transient behaviour of the Finco junction temperature. The temperature was measured inside a standard 9U VME crate, with a fan speed adjusted to 3000 rpm[5]. The fan blows parallel to the MCM substrate. This number gives roughly the same air velocity of 2 m/s as used for the simulations. The exact air velocity was not measured for this test. For a description of the VME test system see Chapter 7.

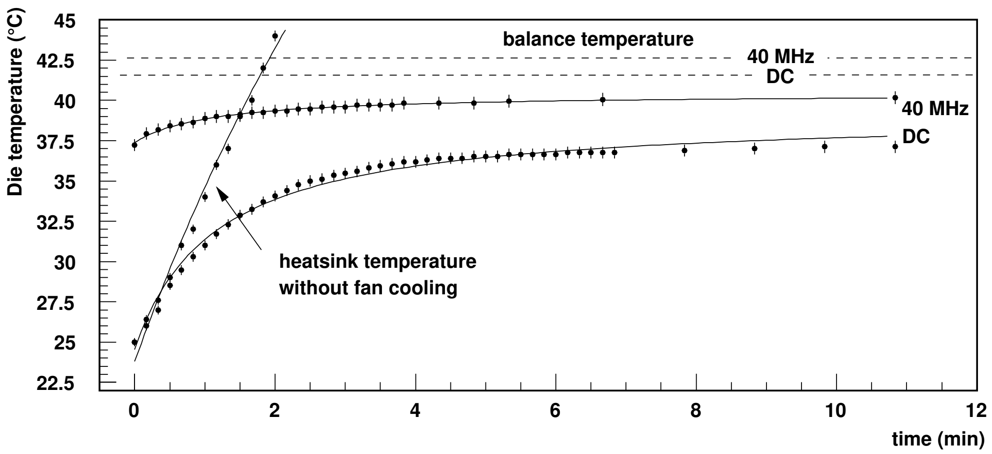

Figure 5.28: Measured transient junction temperature of the Finco die. The transient MCM heatsink temperature without fan cooling is also shown.

[5]rpm: revolutions per minute

Figure 5.28 shows two independent measurements. Each measurement was started at the power-up time, at a room temperature of 25 °C. The first measurement consists of two curves: one curve shows the Finco junction temperature for a duration of 11 minutes, with the 40 MHz bunch-crossing clock switched off (DC measurement), after that, the clock was switched on. This second curve was shifted back to the plot origin in order to enlarge the plot resolution. Two dashed lines indicate the long-term equilibrium temperatures of the Finco ASIC after half a day of running in the test system (40 MHz on/off). The next measurement included was done without fan cooling. During this measurement an external temperature sensor monitored the temperature of the heatsink. For each temperature curve a fit function, as given in Equation 5.7, was applied.

	T_{max} [°C]	A [%]	τ [min]
Finco un-clocked	39.2	19.1	1.1
Finco clocked	40.4	4.0	1.0
heatsink un-cooled	119.3	40.4	7.8

Table 5.13: Fit function parameters.

The un-clocked long-term equilibrium temperature of the Finco junction is about 41.5 °C. This is increased by about 1 °C for a clocked MCM. The MCM heatsink long-term equilibrium temperature is about 32 °C, which leaves a difference of 10.5 °C for the Finco junction. This is only about 45 % of what was calculated for the Finco temperature rise (23.2 °C) in Table 5.8, see Section 5.7.2. There the area of the Flip-Chip solder bumps were estimated to be about 3.3 % of the chip area. A deviation is therefore not surprising, because the thermal behaviour of Flip-Chip bonding is difficult to describe in a thermal calculation. Of course, a better thermal contact is welcome for the reliability of the final PPr-ASIC bonding. The PPrAsic may also be Flip-Chip bonded. Table 5.13 shows a comparison of parameters obtained from the fit function given in Equation 5.7.

Figure 5.29 shows the extrapolation of the fit function applied to the MCM heatsink measurement. This figure also includes the simulated transient-temperature curves shown previously in Figure 5.27. The extrapolation of the heatsink temperature is required because the MCM will be damaged if higher temperatures were to be measured. The maximum heatsink temperature from the extrapolation is about 119 °C. After adding the temperature rise for each junction the die temperature will be higher, as allowed by their data sheets. These measurements show, that fan cooling is essential for the MCM. The MCM will be damaged after a few minutes without a fan. The comparison with the simulated temperature curves show that the cooling effect of the fan was not simulated exactly enough. The simulated and measured temperature shapes are quite similar, and therefore this suggests that the cooling effect was not simulated strong enough.

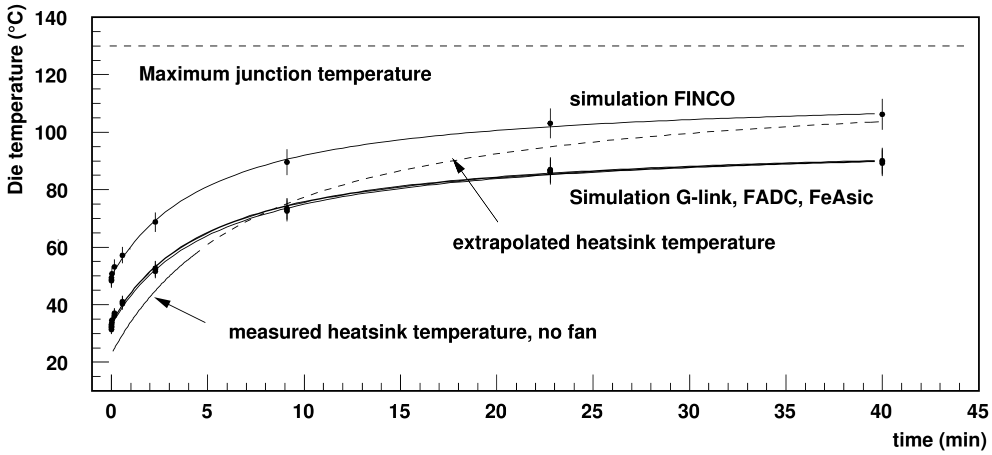

Figure 5.29: Extrapolated heatsink temperature without fan cooling. The simulated junction temperatures with fan cooling are included for a comparison.

5.7.7 Two-dimensional temperature measurements

Two-dimensional temperature measurements were done using an infrared-sensitive camera (Amber Radiance) [IWR]. Its wavelength sensitivity range from 3–5 μm with a resolution of 256×256 pixels and a dynamic range of 12-bits. This camera displays temperature as grayscale and calculates a mean value which can be used for measurement of the MCM mean temperature. Two measurements of the un-clocked MCM were performed: one measurement with fan cooling, and a second measurement which was first un-cooled and after 2.1 minutes the fan was turned on. In both cases the MCM was un-clocked. Figure 5.30 shows that transient temperature behaviour started at power-on time for a duration of 4 minutes.

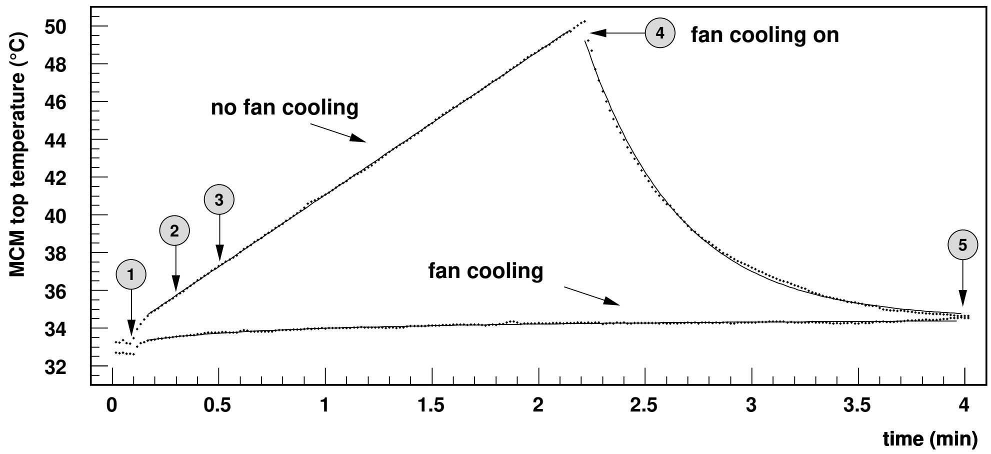

Figure 5.30: MCM two-dimensional temperature measurements.

Three fit functions were applied to the measurement results. One fit function was already shown in Equation 5.7 and the other two functions are as follows:

$$
\begin{aligned}
T(t) &= m \cdot t + T_0 \qquad \text{and} \\
T(t) &= T_0 \cdot e^{-\alpha \cdot t + \delta} + T_\infty .
\end{aligned}
$$

The fit functions are plotted on top of Figure 5.30, using the following parameters:

Fan always on:	$T_{max} = 34.6$ °C, $A = 0.76$ %, $\tau = 0.65$ min
Fan first off:	$m = 7.67\,$°C/min, $T_0 = 33.46$ °C
Fan on at $t = 2.1$ min:	$T_0 = 75.34$ °C, $\alpha = 2.26\,\text{min}^{-1}$, $\delta = 3.39$, $T_\infty = 34.50$ °C

In the following, some pictures are shown which were taken after 6 s, 13 s, 24 s, 2.1 min, and 4 min, after the MCM was powered on. These points are numbered in Figure 5.30 respectively.

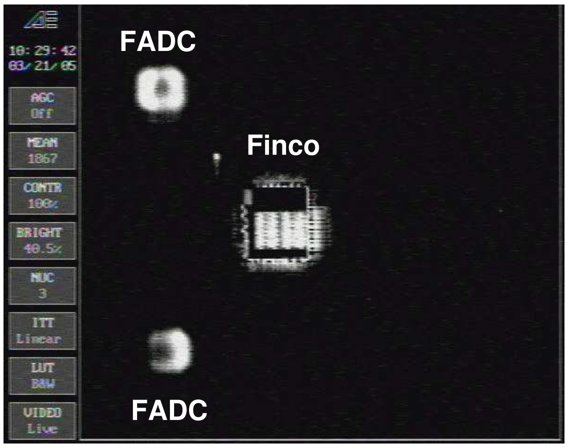

Point 1 ($T = 6$ s): This picture was taken just after the MCM was powered on. The FADCs are shining through their glob-top material. In case of the Finco ASIC the TTL to PECL conversion block is heated first, whereas the rest of the chip is still cold. The MCM mean temperature at this time is 33.5 °C ±1.5°C. This temperature is high, because the MCM was still warm before it was turned on.

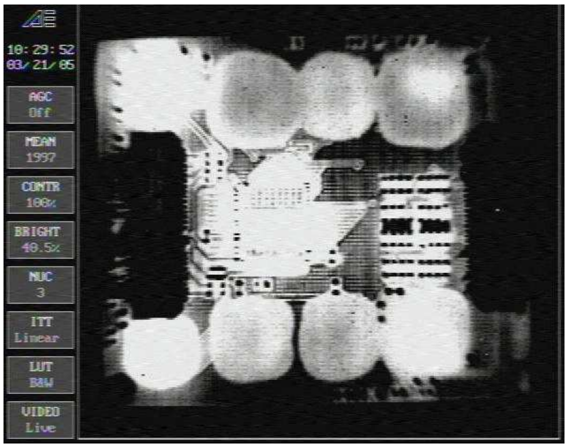

Point 2 ($T = 13$ s): After 13 seconds the MCM mean temperature is 35.9 °C ±1.5°C. From that figure one can see, that the glob-top material is heated first. The MCM substrate is cooler because of the good thermal contact to the heatsink on the backside. Copper tracks are colder because of their good thermal conductance and the termination resistors between the G-link chips contribute to the heat.

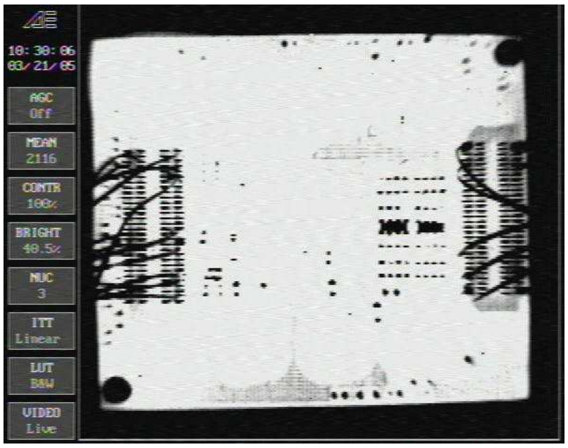

Point 3 (T = 24 s): After 24 s the MCM mean temperature is increased by 3.8 °C to 37 °C ±1.5°C. At this time the MCM substrate and the SMD connectors are heated as well. The thin wires on top of the SMD connectors were used to connect the MCM to its power supply.

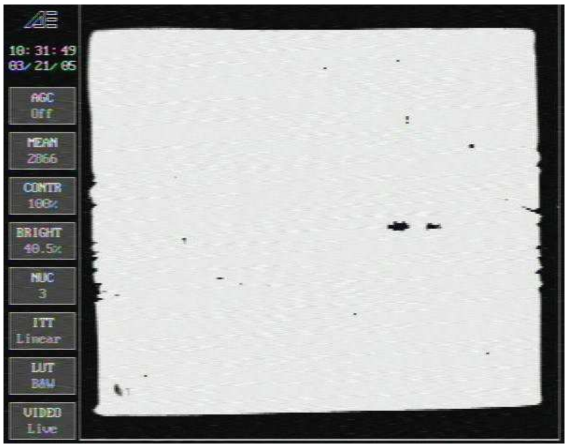

Point 4 (T = 2.1 min): After 2.1 minutes the fan was turned on. The mean temperature has now reached 50.2 °C ±2.0°C. Due to an analogue dynamic range of 8-bit for the display, the MCM looks heated uniformely, but the grayscale numbers are still within the 12-bit range of the camera.

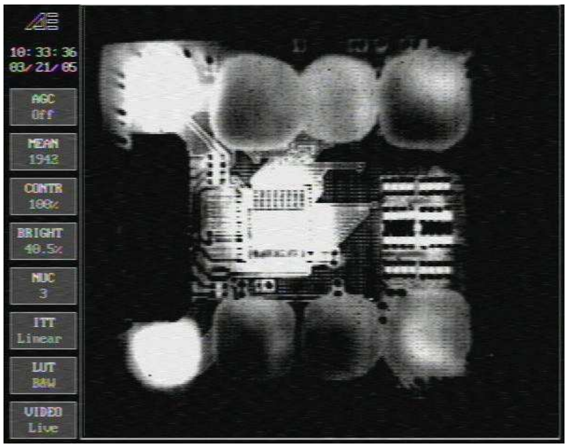

Point 5 (T = 4 min): After 4 minutes the MCM has reached its equilibrium temperature which is 34.6 °C ±1.5°C. The fan blows from bottom to the top of the picture. Hence the cooling effect is better for the chips close to the fan. All the chips on the bottom side of the picture are slightly colder than their adjacent neighbours at the top side. This is particular true for the FADC sitting in front of the SMD connector.

The absolute numbers of the two-dimensional temperature measurements are slightly different from the simulations, but it is nice to see, that qualitatively the air map simulation, shown in Figure 5.24, is quite similar to what was measured in the last figure.

5.7.8 Reliability and temperature — system aspects

When designing large electronic systems, such as the Pre-Processor system of the ATLAS Level-1 Calorimeter Trigger, great emphasis has to be placed on system reliability, especially when the heat density and packaging density is high. Component failure rates are strongly dependent on temperature and hence on the cooling efficiency. Figure 5.31 illustrates the failure rates of a mix of electronic components, digital, analogue, and radio frequency devices, as a function of the junction-to-case temperature. These curves are based on military hardware specifications [MIL-HDBK-217], for more details see [Har97]. The chips shown are normalised to size, complexity, and quality of their environment.

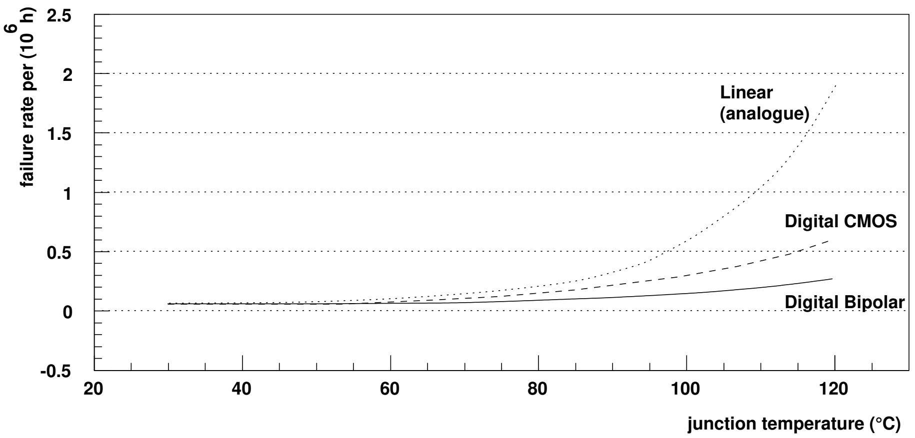

Figure 5.31: Component failure rates versus temperature for digital and analogue components, taken from military hardware specifications [MIL-HDBK-217] (see [Har97]).

Analogue circuits are seen to be much more sensitive to temperature than either digital CMOS or bipolar ones. This suggests a division of cooling resources for optimum reliability. In addition, system failures are often caused by a combination of temperature and vibration. Therefore, low temperature operation is one of the major objective when designing electronic systems as a means of improving system reliability. The reliability of an electronics system, consisting of a group of components, is the probability of operating continuously over a specific period of time with no failure. The reliability is expressed in percent.

Considering an initial number of components $N_i(0)$ of type i at time 0, one can calculate the number of components $N_i(t, T)$, which have operated failure free over a time period t at a temperature T, based on an exponential distribution of the form:

$$N_i(t,T) = N_i(0) \cdot e^{-\alpha_i(T)\cdot t},$$

where $\alpha_i(T)$ is the temperature-dependent failure rate of component i. The life time of that component τ_i is often referred to as Mean Time Between Failures (MTBF):

$$\alpha_i(T) = \frac{1}{\tau_i(T)} = \frac{1}{\mathrm{MTBF}_i}$$

The probability of failure-free operation at time t and temperature T is:

$$p_i(t,T) = \frac{N_i(t,T)}{N_i(0)} = e^{-\alpha_i(T)\cdot t}.$$

For a group of components constituting a system, the probability of failure-free system operation is the product of the individual component probabilities:

$$p^{system}(t,T) = \prod_i e^{-\alpha_i(T)\cdot t} = e^{-\sum_i \alpha_i(T)\cdot t}.$$

For the system $\mathrm{MTBF}_{\mathrm{MCM}}$ follows:

$$\mathrm{MTBF}_{\mathrm{MCM}} = \frac{1}{\sum_i \alpha_i(T)},$$

and the system reliability R can be defined as:

$$R := e^{-\frac{t}{\mathrm{MTBF}_{\mathrm{MCM}}}}. \tag{5.8}$$

As shown in Example 1 below, the module reliability of zero failures for a Pre-Processor Module over a time period of one year is 24.6 %, assuming that each chip in a Pre-Processor Multi-Chip Module has a failure rate of 1 part per one million hours. This example makes it obvious that connectors for the Pre-Processor Multi-Chip Module are essential. An exchange of a Pre-Processor Module as a whole is of coarse not affordable in terms of spares and costs.

Example 1

For a MCM consisting of 10 chips, each having an assumed failure rate $\alpha = 1.0$ part per million hours, the $\mathrm{MTBF}_{\mathrm{MCM}}$ is 100.000 hours. The probability of zero failures over one year is 91.6 % (see equations 5.8). Assuming there are 16 MCMs on a Pre-Processor board, the overall board reliability is only 24.6 %. The inclusion of 15 additional Multi-Chip modules lowers the reliability of a board by 67 %.

5.8 Signal integrity on the PPrD-MCM

In contrast to low-speed printed circuit board designs (PCBs), where the main design consideration is to provide loss-less interconnections between components, high-speed Multi-Chip Module designs behave in a more complex manner. The reduced size makes electrical loss and internal coupling mechanisms more important. Signals are carried in both the conductor and its dielectric surrounding. For this reason, the signal transport is influenced by the electrical properties of the dielectric, especially at high frequencies. This requires that the PPrD-MCM tracks must be designed as micro-strip lines carrying gigabit signals of a wavelength in the near-microwave frequency range (300 MHz – 3 GHz).

This section describes the signal analysis applied to the PPrD-MCM. It starts with a short introduction to the functionality of the DF/SigNoise$^{\mathrm{TM}}$ simulation tool that Cadence provides, followed by simulation results for G-link signals. From such simulations, one can derive electrical properties to characterise the signal quality and integrity with respect to an electronics system during the design process. Hence, the signal simulations performed are referred to as *signal integrity* simulations. For the basic theory of signal transport simulation see Appendix A.

5.8.1 Signal integrity analysis using DF/SigNoise

Analysis requirements

Before running a DF/SigNoise$^{\mathrm{TM}}$ simulation [Sig97], the design needs to be properly prepared with regard to properties, constraints, layer definitions and device models. Device models are a listing of all the pins on a certain device, and an assignment of which pin uses which I/O cell model. They contain behavioural information about an I/O buffer, such as its voltage thresholds and V/I curves. It also include the bonding wire parasitics (R,L,C) associated with each pin. The device models can be:

- a default I/O cell model;
- an IBIS device model (see [IBIS] for a specification);
- a SPICE model (see [SPICE] for a user guide).

Multilayer definition

The multilayer definition provides the geometries, which are required by the DF/SigNoise™ simulator. The conductor/dielectric thicknesses and dielectric constants give the z-axis information, the x- and y-axis information is obtained from the conductor routings in the design. This information is the database for characterising the interconnection during signal analysis.

Equivalent-circuit model generation

To generate an equivalent-circuit model for a selected net, a coupled network of multiple conductors is generated by:

1. performing a detailed geometry extraction, including the primary net, its neighbour nets, and the mutual coupling between them. The neighbour nets are found by the geometry of a user-defined search window for a local area without long-range effects;
2. a subdivision of each net into unique individual cross-section elements;
3. invoking a 'field solver' for each cross-section element to generate an individual trace model for it, if it does not already exist.

The 'field solver' uses a two-dimensional boundary element technique to break up the cross-section into a fine mesh. Charge distribution over the mesh is then solved to derive a capacitance matrix for the conductors in the cross-section element. The capacitance matrix is then used to derive the inductance matrix. Resistance and conductance matrices are also generated to take into account conductor losses such as skin effect, and dielectric losses.

The output of the 'field solver' are frequent-dependent RLGC matrices (resistance, inductance, conductance, and capacitance) per unit length, which get stored along with the geometry as a trace model. This minimises the number of field solutions to characterise a particular layout. From these RLGC matrices an interconnect network is derived, representing a simulation circuit to simulate the impedance and propagation delay of a signal, including mutual couplings to other tracks.

Simulation types

DF/SigNoise™ uses a simulator which is based on a subset of public domain SPICE[6] with extensions developed by Cadence to optimise it for Multi-Chip Modules. For example, the simulator uses macro models for micro-strip lines, in general

[6] SPICE: Simulation Program with Integrated Circuit Emphasis

referred to as transmission lines, and I/O cells. The different analysis types are distinguished by what is included in the simulation circuit and how a stimulus is applied to it. For a more detailed description of the simulator see [Sig97]. The simulation types which can be applied are as follows:

- **Reflection simulation:** For a reflection simulation, the selected primary net is characterised by its trace geometries extracted from the layout. Device models for the driver and receiver pins are used to build a 'single-line' simulation circuit, disregarding neighbour nets. A stimulus signal is applied to the driver pin (rise, fall, pulse, inverted pulse) and an output waveform is produced for all driver and receiver pins in the circuit. Thermal drift of the receiver pin is derived by $V_{td} = (T_j - T_{ref}) \cdot \frac{\mathrm{d}V}{\mathrm{d}T}$, where V_{td} is the thermal drift voltage, T_j is the junction temperature, T_{ref} is the reference temperature, and $\frac{\mathrm{dV}}{\mathrm{dT}}$ is the voltage temperature drift.

- **Crosstalk simulation:** The simulation circuit is built in such a way that it includes the primary net selected for simulation, neighbours and mutual coupling. The primary net is held at the high or low state as appropriate, and the neighbours are simulated inverse to the primary net state. Crosstalk is simulated in the time domain, producing waveforms and report data.

- **Comprehensive simulation:** In a comprehensive simulation, a 'multi-line' circuit is built similar to that used for crosstalk analysis. Power and ground parasitics are taken into account. The primary net is stimulated, and the neighbour nets are stimulated inverse to the primary net. The drivers on the primary net and the neighbour nets are switched simultaneously. This analysis represents the worst-case transient scenario.

- **System-level signal analysis:** Nets that span multiple designs can be analysed using a design link, e.g. a Multi-Chip Module that is mounted on its motherboard. The design link contains a model to represent the connector on either side and, if required, a cable that physically connects the two layouts. One simulation circuit is generated for an entire net, which is referred to as a full system-level simulation.

Electrical signal properties such as: propagation delay, final settle delay, over- and undershoot, can be derived from the simulation types described above. The *propagation delay* is the transmission line delay, the time required for wave propagation from the driver to the receiver. The *final settle delay* is determined by first measuring the time from the start of the simulation (time zero) to when the receiver voltage crosses its high threshold level for the final time. Then the buffer delay of the driver circuit is subtracted. This delay characterises the time required for the signal transmission including the time to set the receiver circuit to its high logic state. An *overshoot* is the voltage swing beyond the steady state voltage level and an *undershoot* is the voltage swing back into the midrange after the nominal steady state level has been crossed. The overshoot is important because a voltage which is greater or lower than the maximum logic range can damage a device. These definitions are illustrated in Figure 5.32.

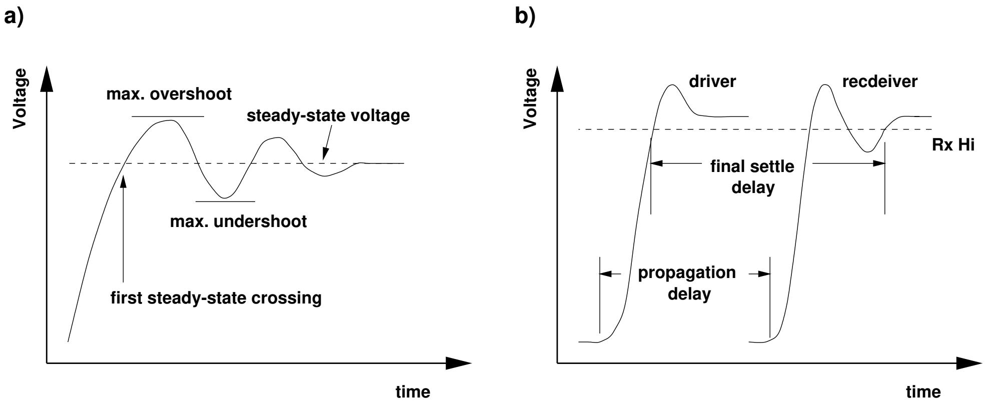

Figure 5.32: Illustration of over- and undershoot is shown in (a), propagation delay and final settle delay in (b).

5.8.2 Signal integrity simulation of gigabit signals

One of the critical signals on the demonstrator Multi-Chip Module is the high-speed serial output of a G-link transmitter (HDMP-1012). This signal is supposed to be sent via an SMD connector to a line driver sitting on a motherboard. The line driver acts as a buffer to drive a single-ended coax cable over a distance of a few meters. In order to simulate the circuit in between the G-link bonding pad and the line receiver input pin, a system-level simulation was performed (see Section 5.8.1). Figure 5.33 shows the extracted schematic topology used to analyse the influence of the MCM layout on that circuit. This schematic includes an IBIS device model for the G-link output driver including capacitance and inductance of the wire-bond connection, a transmission-line model extracted from the MCM layout, an RLGC matrix model for the connector, a transmission line model extracted from the motherboard PCB layout, and an IBIS device model of the line receiver, as well as passive components needed for signal termination. The signal is transmitted on a 100 μm wide track on the MCM and on the PCB top routing layer. An input stimulus was applied to the G-link driver model depending on the type of simulation performed.

Reflection on G-link signals

Figure 5.34 shows a reflection simulation result of the schematic shown in Figure 5.33. The input to that simulation was a rising pulse edge which changes the

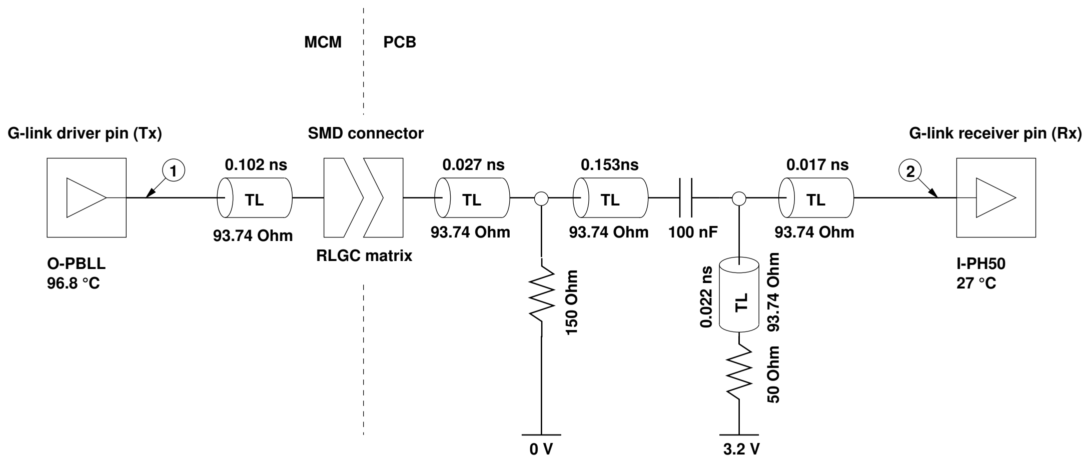

Figure 5.33: Extracted simulation topology (schematic) used for the system-level signal simulation. Device models for the G-link driver and receiver pin, connector RLGC matrix, transmission line models (TL), and passive components used for termination, are shown.

PECL[7] logic state of the G-link driver output from low to high. Dashed lines are drawn for the PECL levels assigned to the logic state of the driver and the receiver. The voltage of the G-link driver pin on the MCM and the voltage of the destination receiver pin on the motherboard are shown as transient waveforms. Their recording positions are labeled in the schematic with '1' and '2' respectively. From the plot one can get undershoot and overshoot (~25 mV each), first switching (~0.5 ns) and final settle delay (~0.8 ns). The actual numbers are summarised in Table 5.14 for a variable track length ranging from 2 to 12 cm.

The reflectivity r can be estimated from that plot, which in general is a complex value, see Appendix A. For tracks which have smaller ohmic losses, the track impedance Z_o can be considered as pure ohmic, see Equation A.20, which results in a real number for the reflectivity. The initial voltage pulse transmitted at $t = 0$ can then be simplified and calculated from:

$$U(0) = \frac{Z_0 \cdot U_0}{Z_i + Z_0}$$

where Z_0 is the track impedance, U_0 the open circuit voltage of the driver, and Z_i is the output impedance of the driver. From the G-link specification one gets $Z_i = 50\,\Omega$ and $U_0 = 4.2\,V$, and from the layout extraction one gets $Z_0 = 93.74\,\Omega$. The voltage pulse $U(\tau)$, which arrives at first at the receiver circuit after a propagation time τ, can be calculated by (A.19):

[7]PECL stands for positive-ECL, where the power supply ground pins are connected to +5 V and the negative power supply pins (VSS) are connected to ground (0 V). The logic levels are shifted respectively.

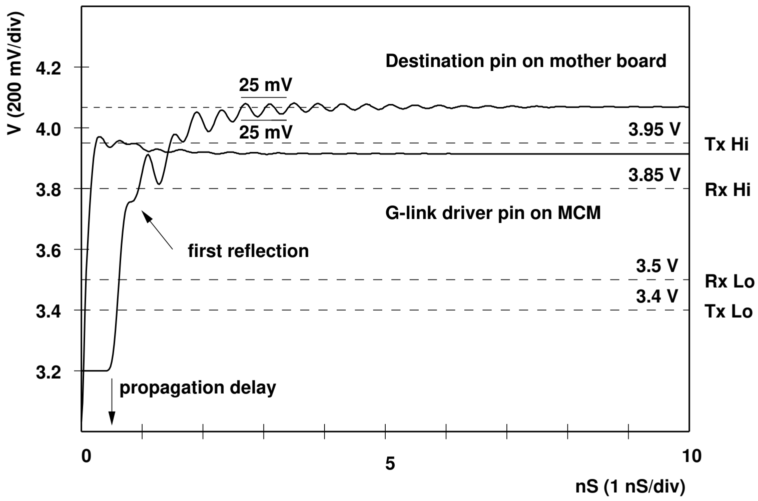

Figure 5.34: Rise reflection as transient simulation for a tracklength of 2.95 cm. Includes: layer definitions, stripline length, thermal driver shift.

$$U(\tau) = U(0) + r \cdot U(0)$$

This voltage can also be taken from Figure 5.34, which is 3.75 V, and finally it can be used to calculate an approximation for the reflectivity r, which is:

$$r = 0.37.$$

The reflectivity is larger than zero, corresponding to an impedance load greater than the output impedance of the G-link driver circuit.

Figure 5.35 shows a pulse reflection simulation, where a rectangular waveform with a frequency of 480 MHz was used as simulation input. This frequency[8] is equivalent to a maximum serial Baud[9] rate of 960 MBd for the G-link output. The serial Baud rate includes 4 additional protocol bits, whereas the maximum utilisable data rate is 800 MBit/s for 20-bits. From the figure, one can see that a reflection affects the rising and the falling edge of the transmitted rectangular waveform. The signal is distorted in such a way that the high state of the signal falls briefly below the level which the receiver requires to identify the level correctly as a logic '1' state. For the low state this level is just reached. This reflection could therefore be a source of transmission errors depending on the point in time when the data is latched into

[8]The number of bit/s is equivalent to twice the frequency, if the negative half wave corresponds to a logic '0' and the positive half wave corresponds to a logic '1'.

[9]The unit Bd (Baud) is the maximum number of the shortest codes (bits) per second in a transmission system, including protocol information.

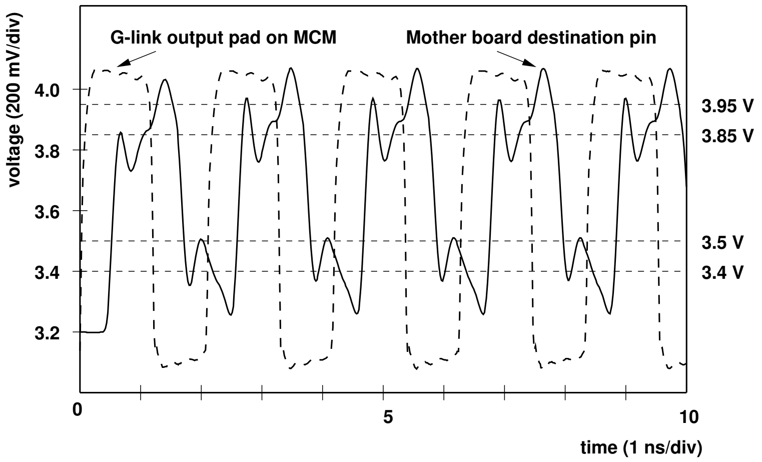

Figure 5.35: Pulse reflection of the G-link transmitter running at 960 MBd. A pulse stimulus of 480 MHz was used for a track length of 2.95 cm.

the cable driver. The simulation was performed for a track length of 2.95 cm. The position of the first reflection depends on the propagation delay, and therefore on the track length. As a remedial action, the track length can be optimised and the signal termination should be further improved to reach a reflectivity which is closer to $r = 0$.

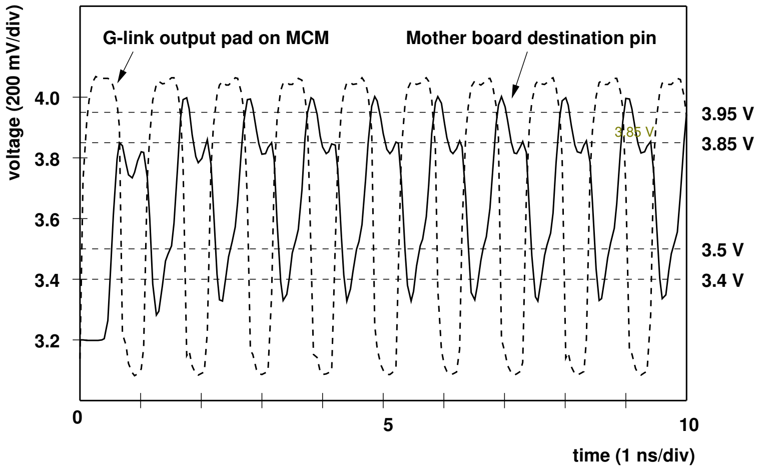

Figure 5.36: Pulse reflection of the G-link transmitter running at 1920 MBd. A pulse stimulus of 960 MHz was used and the track length was 2.95 cm.

Figure 5.36 shows a pulse reflection simulation where the serial data rate was doubled (1920 MBd). The simulation input was a 960 MHz rectangular waveform. At this serial data rate, one further reduces the peak-to-peak voltage at the receiver. The duration at the corresponding logic levels is reduced, which makes the transmission more sensitive to bit errors. The normal operating range of the G-link is from 150 MBd up to 1500 MBd with a bit error rate of less than 10^{-14} [HP]. This range can be extended up to 1800 MBd for a maximum allowed operating temperature of +60 C°. In this range the bit error rate is increased by a factor of a thousand to 10^{-11}. These numbers confirm the distortion effect seen, and that one has to expect an increased bit error rate when running at 1600 MBd.

Propagation delay sweep

Figure 5.37 shows the final settle delay versus the propagation delay for a G-link signal at 960 MBd. The simulation result for the propagation delay, which depends on the track length, is 5.55 ns/m. For each of the simulation points shown, a transient reflection simulation was performed, with a different track length as parameter. The length was varied in a range from 2 cm up to 18 cm. The correlation one would expect is a linear slope of the final settle delay. This is only the case from 5 cm upwards, whereas below 5 cm, reflections which occur on the rising signal edge cause an increase of the final settle delay.

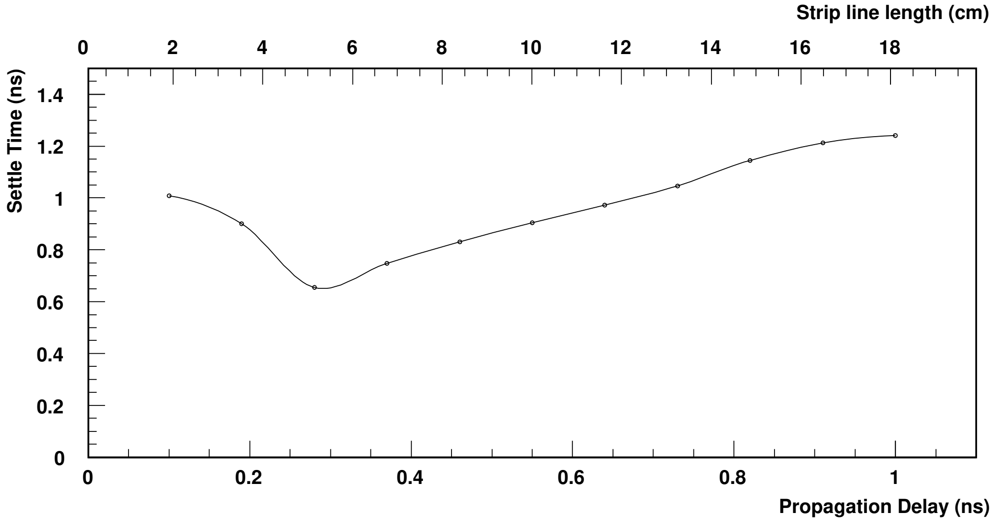

Figure 5.37: Final settle time versus propagation time delay, introduced by a variation of the MCM track length.

Signal overshoot and undershoot can be derived from the same simulation as described before. The results are given in Figure 5.38. Again the simulation results are dependent on the track length. This suggests an optimisation of the track length in a range between 5 cm and 10 cm, to minimise the over- and undershoot.

Crosstalk on G-link signals

Figure 5.39 shows the crosstalk on a G-link signal from all neighbours. This simulation uses the same topology as shown in Figure 5.33. The PCB layout consists

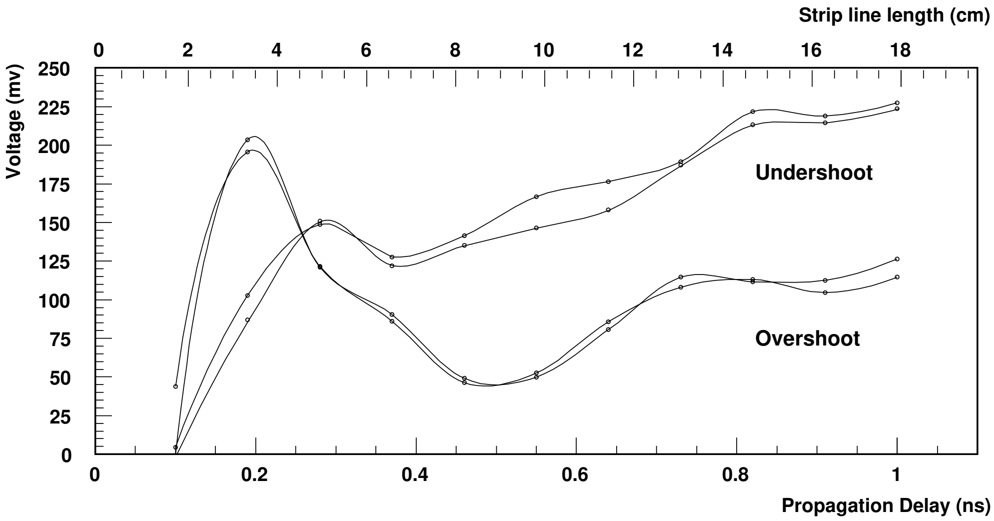

Figure 5.38: Propagation delay sweep (Over-, Undershoot)

only of those parts which are required for signal termination and of the coax cable driver. Hence, neighbouring signals are only present on the MCM layout and are the only source for 'direct' crosstalk, whereas mutual coupling is present for both layouts. The simulation includes the G-link net, all the neighbours found in a distance of 500 μm (search window), and mutual coupling. The primary net was held high during this simulation, whereas the neighbour nets go in the inverse direction. The crosstalk on the G-link driver pin is about $\sim$10 mV in the high state, and on the PCB it is about $\sim$50 mV for the low state.

Figure 5.40 shows a comprehensive simulation. It includes the same setup as for the crosstalk simulation as shown before. In addition, power and ground parasitics, rise reflection on the primary net, and the neighbour net are simulated in the inverse direction. This comprehensive simulation is a combination of the rise reflection shown in Figure 5.34 and a crosstalk simulation shown in Figure 5.39. It represents a worst-case transient scenario for a rising pulse signal. The G-link output on the MCM is more affected in this simulation. The overshoot is about $\sim$260 mV, measured from the steady state voltage level. On the destination pin it is only about $\sim$20 mV. The undershoot is about $\sim$60 mV on the MCM pin and $\sim$90 mV on the PCB pin.

System-level simulation summary

A summary of electrical parameters derived from the system-level simulations shown above is given in table 5.14. The properties were recorded at the point labelled '2' in Figure 5.33. The table contains the propagation delay for a strip-line of 100 μm. Maximum values for under- and overshoot and final settle delay are

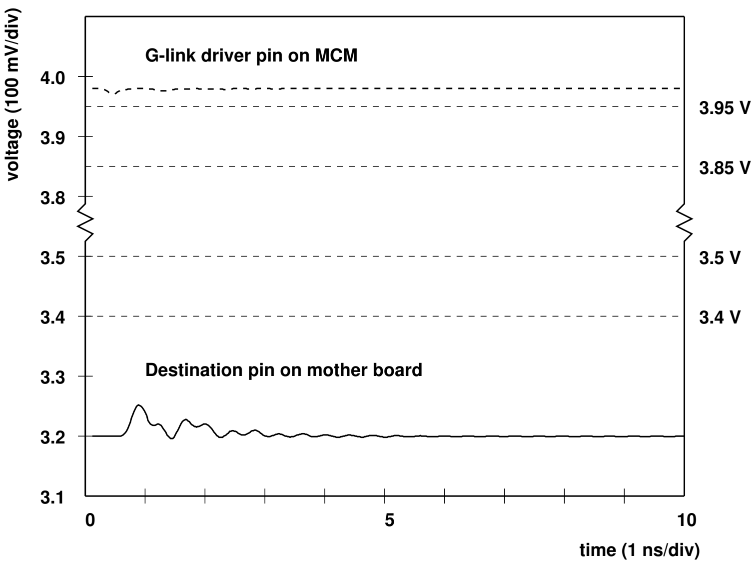

Figure 5.39: Crosstalk from all neighbours. Includes primary net, all neighbours, and mutual coupling.

given, in a range where the track length is varied (2 cm up to 12 cm). The crosstalk on the cable driver pin was ~50 mV for a 2.95 cm long track.

Simulation results for a 100 μm strip-line	
Propagation delay	5.55 ns/m
Undershoot	$\leq$ 225 mV
Overshoot	$\leq$ 200 mV
Final settle delay	0.65 ns - 1.25 ns
Crosstalk	~50 mV

Table 5.14: DF/SigNoise^{TM} simulation results, at the cable driver input pin of the motherboard.

5.8.3 Signal integrity measurements of gigabit signals

This section presents measurements which were made to gain confidence that the MCM strip-lines behave in a similar way to what observed in the signal integrity simulations. A network analyser, HP 4396A manufactured by Hewlett-Packard, was used to measure the electrical characteristics of a strip-line, such as transmittivity, reflectivity, track impedance, crosstalk, and phase-shift. The same high-speed serial strip-line which was simulated before was connected to a transmission/reflection test-set. This test-set consists of a directional coupler to separate reflected or transmitted parts from an electrical signal. Measurement results pre-

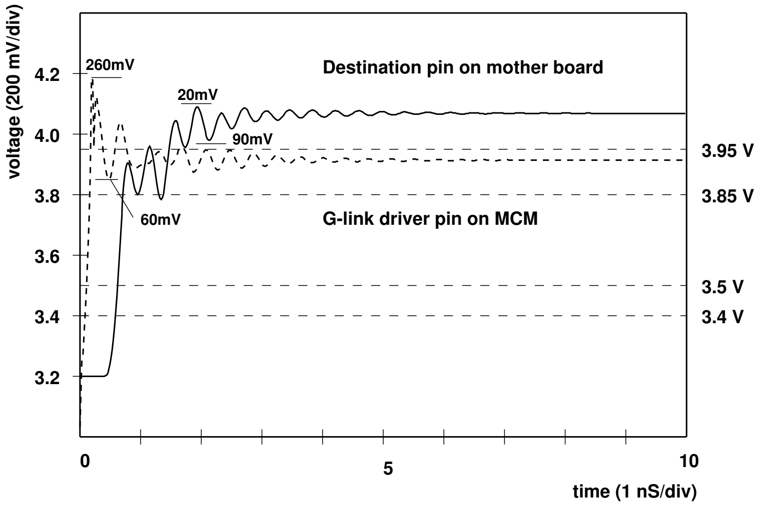

Figure 5.40: Comprehensive simulation (worst-case transient scenario).

sented here look different than results from simulations because the network analyser performs a frequency sweep in a range up to a maximum of 1.8 GHz. Therefore, the measurement results are shown as spectra and not as transient waveforms for a specific frequency.

The MCM is connected to the network analyser via cable connectors and termination resistors. The influence of these components must be calibrated before measurements in a microwave frequency range can be performed. The MCM under test must be shielded from its environment inside a metal box and the strip-line must be connected to a 50 Ω coax cable. Standard needle probes cannot be used for measurements in a micro-wave frequency range. High-speed needle probes, which would work of coarse, are expensive and where not available for this test. Hence, soldering of the strip-line as close as possible to the input cables of the transmission/reflection test-set was used. This has given the best results.

In the following, the frequency-sweep of the network analyser was adjusted from zero to 1 GHz. Beyond 1 GHz, resonances from the hollow body of the shielding box occurs.

Measured transmission/reflection

Figure 5.41 shows the transmittivity of a 100 μm strip-line. The strip-line has a length of 1.8 cm. It starts at the G-link bonding pad and it ends after the SMD connector. Compared to the simulations, which were described before in Figure 5.33, this measurement does not include the G-link driver circuit nor the termination circuit on the motherboard.

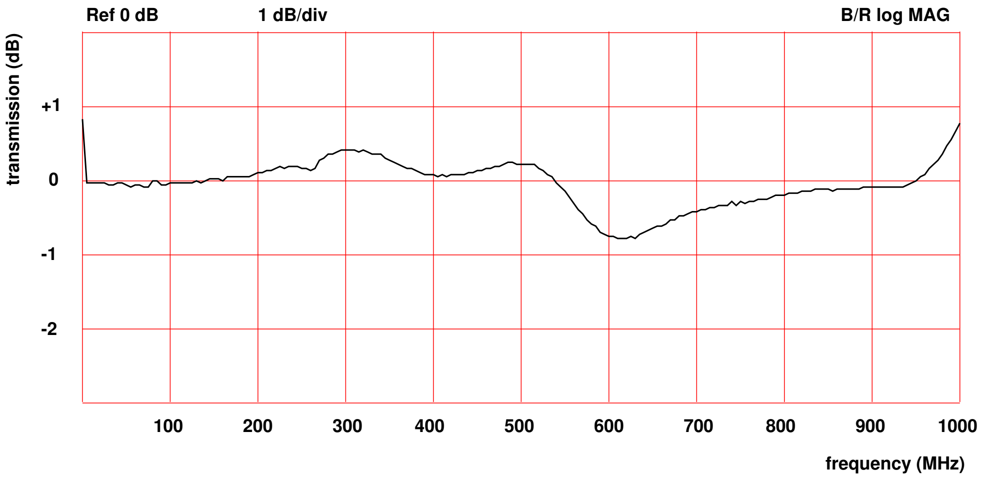

Figure 5.41: Transmittivity of a 100 μm strip-line.

The strip-line transmittivity is measured in decibels (dB). This is calculated from the voltage ratio of the transmitted signal U_t to the source signal U_s as follows:

$$t\,[\mathrm{dB}] = 20 \cdot \log \frac{U_t}{U_s}.$$

The measured transmittivity is better than –1 dB. This corresponds to a transmitted signal amplitude of 89.1 % of the original source. Two frequency points are of interest for the two G-link transmission modes: 480 MHz (960 MBd) and 960 MHz (1920 MBd).

Figure 5.42 shows a polar diagram of the transmission measurement. The transmittivity is plotted in the radial direction, whereas the angle gives the phase-shift between the signal source and the transmitted signal. From the figure one can see that the phase-shift is between about zero and –40°. The phase is zero for DC voltages and goes negative to about –40° at 850 MHz. Then the phase-shift is reduced again until 1 GHz is reached. This is due to the inductive behaviour of the strip-line.

Measured crosstalk on G-link signals

The G-link device has differential PECL output signals. These signals are transported via tracks 100 μm apart. For the crosstalk measurement, one track was connected to the source and the neighbouring track was connected as the transmitted signal. The crosstalk to the neighbouring track in a distance of 100 μm can then be obtained by a transmittivity measurement.

Figure 5.43 shows how the crosstalk depends on the frequency. It increases to about –20 dB at a frequency of about 250 MHz. Then it stays constant up to a

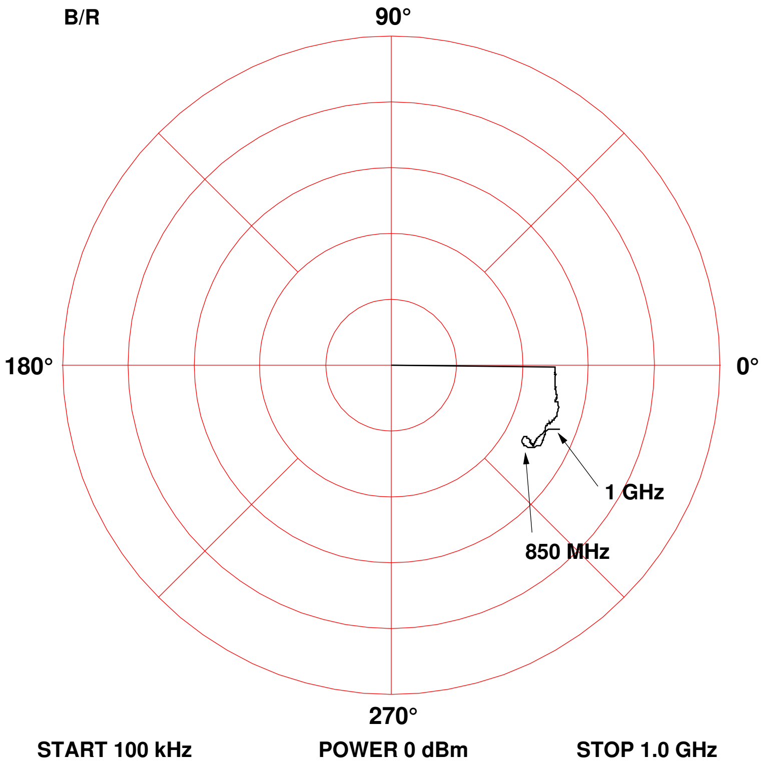

Figure 5.42: Polar diagram for the transmission measurement

frequency of 1 GHz. A crosstalk of –20 dB corresponds to 1/10 of the input signal amplitude. For a PECL voltage swing of ± 400 mV, this corresponds to 40 mV crosstalk. This is similar to what was simulated (~ 50mV).

Strip-line impedance

The impedance Z_{track} of an MCM strip-line can be calculated from the reflectivity r or from the transmittivity t directly by:

$$\begin{aligned} Z_{track} &= \frac{r+1}{r-1} \cdot Z_0 \qquad \text{or} \\ &= \frac{t}{2-t} \cdot Z_0. \end{aligned}$$

This assumes, that $t = 1 + r$ and r is negative. See Appendix A for a derivation of these equations (Equation A.19 and A.21). The real and imaginary parts of the track impedance are shown in a Smith chart in Figure 5.44. The circles represent values with constant real part and the curves represent a constant imaginary part. The impedance curves start at the singularity at ∞, which means that for DC

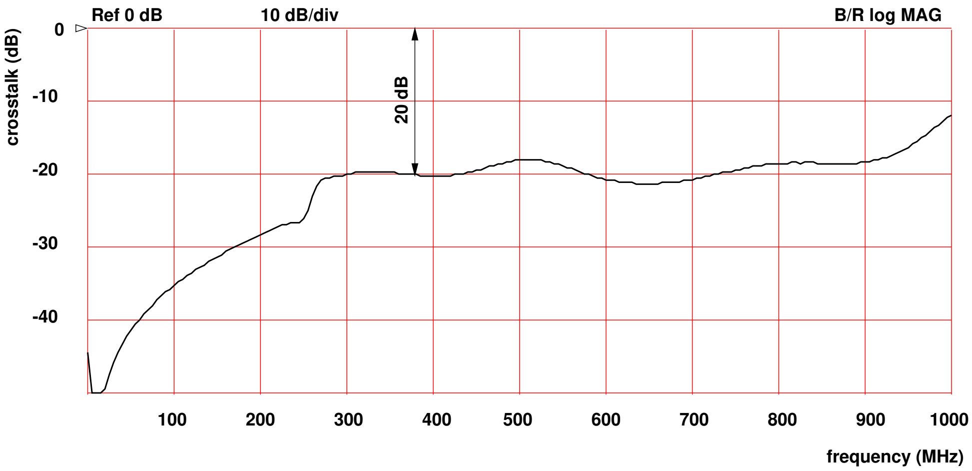

Figure 5.43: Crosstalk of differential strip-lines 100 μm apart.

voltages the track impedance between the track and the ground plane is ∞. For higher frequencies the impedance curve moves along the constant 25 Ω real circle. This means the real part of the track impedance is fairly constant. The imaginary part of the impedance increases to negative values, which means that the track impedance appears as a capacitor. At about 850 MHz the curve makes a loop and goes back in the direction of more positive imaginary values. At this frequency the track inductance seems to have an effect. The real and imaginary parts of the impedance curve at the frequencies of 480 MHz and 960 MHz can be used to calculate the track impedance by taking the root-mean-square of these values. The impedance for a 100 μm strip-line on the top MCM layer is 51.5 Ω (46.55 Ω + i21.35 Ω) and 50.5 Ω (46.25 Ω + i20.30 Ω) at a frequency of 480 MHz and 960 MHz. This is close to the optimum value of 50 Ω, which is best matched to the G-link output driver pin. However, the calculated track impedance using the MCM layer definitions was about 90 Ω for the smallest track width of 100 μm.

A summary of the signal integrity measurements for a 100 μm track on the top MCM layer is given in Table 5.15.

Measurement results for a 100 μm strip-line	
Phase-shift	$< 40°$
Maximum reflectivity -r	< 0.11
Minimum transmittivity t	> 0.89 (± 1 dB)
Track impedance Z at 480 MHz	51.5 Ω (46.55 + i21.35)
Track impedance Z at 960 MHz	50.5 Ω (46.25 + i20.30)
Crosstalk	$\sim$40 mV (20 dB)

Table 5.15: Measurement summary of a 100 μm strip-line in a frequency range from 0–1 GHz.

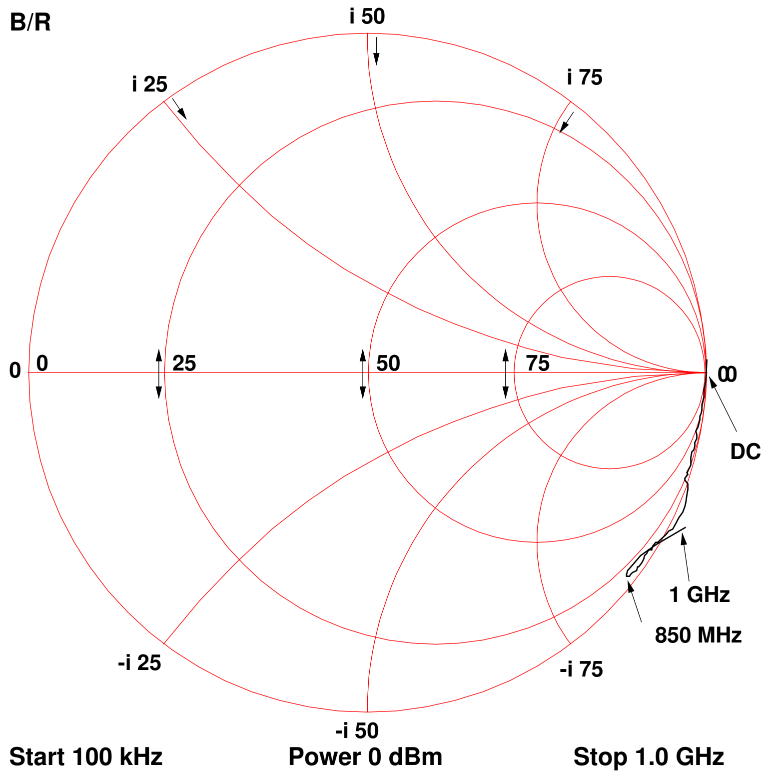

Figure 5.44: Smith chart for the track impedance measurement.

5.9 Considerations for the final PPr-MCM

The final PPr-MCM layout will probably be less demanding in terms of power dissipation than the demonstrator PPrD-MCM, see Section 3.3.1. The exact power dissipation is not yet fixed, but it could be reduced by about 3 W. This improvement would be due to the use of LVDS transmitters instead of G-links. Furthermore, the PPrD-MCM temperature measurements have shown that the heat can be efficiently removed by air cooling. The cooling effect is better then expected from the simulation results. This leaves the remaining critical issue: the MCM area must be small, and its aspect ratio optimised for the final board density of the Pre-Processor Module.

5.9.1 Mass production

The demonstrator MCM was designed at the University of Heidelberg and manufactured by Würth Elektronik. Except for the Flip-Chip bonding of the Finco ASIC, all assembly was done at Heidelberg. This includes the following procedures: die gluing, wire-bonding, SMD soldering, encapsulation, and testing. If one would do these steps one after another, an MCM assembly would take about three days. A final mass production requires at least 1824 PPr-MCMs, without spares. This will require the involvement of external companies to a further extent. Keeping the ATLAS Calorimeter Trigger time schedule in mind, more production steps should be performed by a single company. This would optimise the logistics and the supply times. The following sequence has turned out to be best suited for the assembly. Tasks which must be done by external companies are marked with 'external', others are marked with 'internal'.

1. Flip-Chip mounting (external);
2. Die mounting/gluing (may be internal);
3. Soldering of nearby SMD components. Only those SMD components which will be affected by the glob-top encapsulation of dies (may be internal);
4. Clean from solder remains (internal);
5. wire-bonding (external);
6. Short electrical DC chip tests (internal);
7. Glob-top encapsulation (external);
8. SMD soldering of the remaining SMD parts, e.g. the SMD connectors are to bulky if they are soldered before wire-bonding (may be internal);
9. Final clocked MCM test (internal);
10. Global encapsulation (external);

Chapter 6

A Flip-Chip interconnection ASIC

- Tasks for the PPrD-MCM
- Finco layout
- Test results

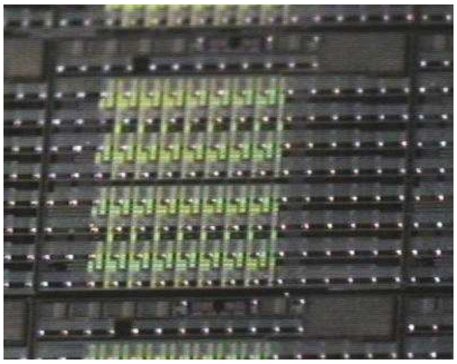

6.1 Tasks for the demonstrator Multi-Chip Module

The Finco[1] IC is an application specific design for the demonstrator Multi-Chip Module. It was fabricated in a 0.8 μm BiCMOS-process[2] offered by Austria Micro Systems [AMS]. The BiCMOS-process combines the advantages of both bipolar and CMOS[3] technology on the same chip, and hence the mixing of ECL and CMOS standard cell libraries. The Finco ASIC is a mixed-signal design, which consists of analogue and digital standard cells [AMS96]. It performs some of the Pre-Processor tasks for the demonstrator project: analogue baseline adjustment for four trigger tower signals, data multiplexing, and temperature monitoring.

The Finco layout was optimised for Flip-Chip[4] mounting, where small solder bumps are used to form the electrical contacts between pads on the chip and bonding pads on the MCM substrate. The chip I/O pads are arranged in a two-dimensional array, with a pad spacing of 350 μm in x-direction and 450 μm in y-direction. Figure 6.1 a) shows a photo taken from the Finco ASIC after the solder bumps where put-on, and 6.1 b) after mounting onto the demonstrator Multi-Chip Module.

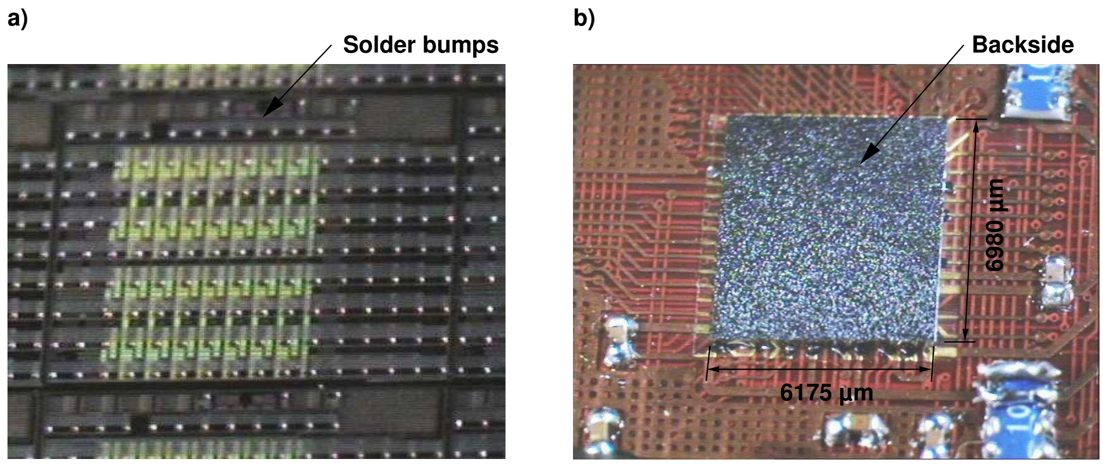

Figure 6.1: Finco ASIC pictures before and after mounting.

For testing of the demonstrator Multi-Chip Module, the Finco ASIC provides boundary-scan and multiplexing of FeAsic SpyBus signals. In addition, one analogue line-receiver circuit was designed. This circuit is used to investigate the reception of analogue trigger tower signals on an ASIC, which is sitting on a Multi-Chip Module. For more technical reasons, level conversion from TTL to PECL is

[1]Finco: Flip-chip INterCOnnection ASIC
[2]BiCMOS: Bipolar CMOS
[3]CMOS: Complementary Metal Oxide Silicon
[4]Flip-Chip: This name stems from the mounting technique, where the silicon die is turned upside-down.

required. The ASIC is controlled via its serial interface for loading of configuration data. Figure 6.2 shows a block diagram of the ASIC. Details about its tasks are given in the following sections and they are summarised as follows:

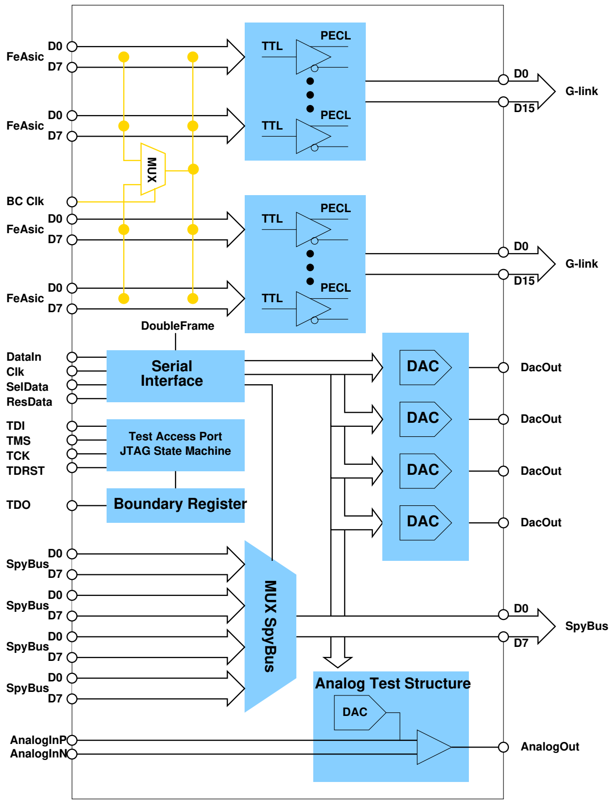

Figure 6.2: Finco block diagram.

- **Data multiplexing:** Data from four FeAsics (4×8-bit) are multiplexed from 40 MHz to 80 MHz (2×8-bit). This is required for data transmission of four trigger towers via one G-link transmitter at 1600 MBd.

- **Digital-to-analogue conversion:** Four 10-bit macro-cell digital-to-analogue converters (DACs) are included to adjust the baselines of analogue PPrD-MCM input signals. These DACs are part of an analogue library that AMS provides for their 0.8 μm CMOS process.

- **Signal level-conversion:** The Finco ASIC includes 32 level converters from TTL to PECL (4×8-bit data). PECL levels are required as inputs to the G-link transmitters.

- **Temperature monitoring:** Temperature is one of the MCM design parameters which needs to be investigated. The Finco ASIC measures its own junction temperature and therefore the temperature on the MCM. This feature was used for measurements of the PPrD-MCM temperature, and it will also be implemented in the PPrAsic for temperature monitoring.

- **Line receiver:** The Finco ASIC includes one line receiver test circuit for one analogue trigger tower signal. The circuit consists of a differential amplifier and one 10-bit DAC with a low impedance output buffer for baseline adjustment. The output from this circuit can be used directly as input to an FADC. This circuit can answer the question: how critical is it to include analogue line receivers on an MCM, e.g. in terms of crosstalk? This would reduce the Pre-Processor component density further, if required.

- **Flip-Chip mounting:** The Finco layout is optimised for Flip-Chip mounting. Its pads are arranged as an array. Flip-Chip mounting is being considered for the final PPrAsic to improve bonding reliability and to reduce the MCM size.

- **SpyBus multiplexing:** Each FeAsic provides an 8-bit test bus (SpyBus) to 'spy' on its internal register with a speed of 40 MHz. One of four SpyBuses can be multiplexed to the external MCM pins.

- **Boundary-scan:** This is a very useful test facility for a Multi-Chip Module and for digital Integrated Circuits (ICs). It allows testing of internal chip logic and the surrounding electronics. This was extensively used for a functional test of the Finco ASIC, where a large number of chip tester channels could be saved.

6.1.1 Data multiplexing for double-speed serial-link transmission

Data multiplexing is required to double the serial transmitter data rate from 800 MBd to 1600 MBd. This 'double-frame' mode can be configured via the serial interface. After a chip reset, the double-frame mode is turned off and the chip is in 'normal' operation. In normal operation no clock is required and each input pad is always assigned to the same output pad. The double-frame mode can be turned on and off using the LoadFrame token with data 0x001 or 0x000 (hex) respectively. For a definition of interface tokens see Table C.5.

Figure 6.3 shows a timing diagram of the double-frame mode. During the high state of the 40 MHz clock cycle, data from FeAsic 3 and FeAsic 4 are sent to the G-links, whereas during the low state data from FeAsic 1 and FeAsic 2 are sent. In this mode both G-links get the same input data and one G-link can be used as fan-out.

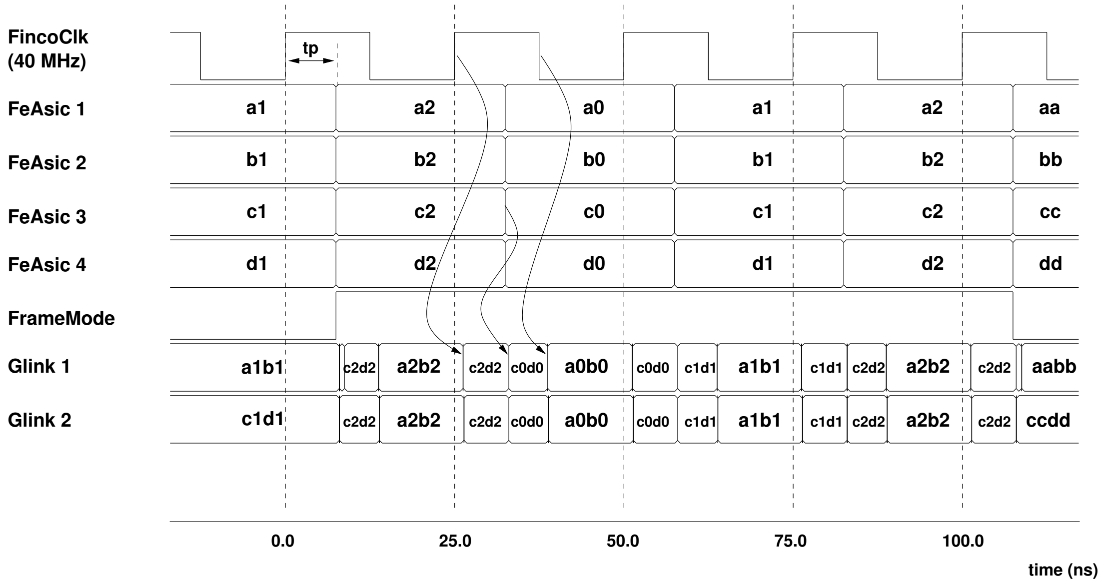

Figure 6.3: Timing diagram of the double-frame mode.

To avoide timing violations of setup and hold time, the Finco clock must have the same timing as the input data from the FeAsics. This is best matched if the delay t_p between FeAsic data and Finco clock is zero.

6.1.2 Integrated digital-to-analogue converters

Digital-to-analogue converters (DAC) are required for the first processing step in the Pre-Processor, see Section 3.2.1 for a description. The Finco ASIC includes four of these converters, made from analogue standard cell designs taken from AMS libraries. These DACs have a resolution and linearity of 10-bits, and their architecture is based on two resistor dividers optimised for small area (1/4 mm^2). Due to this area optimisation, the output impedance is high (21 kΩ) and requires a low-offset amplifier at the output. The digital data for each DAC can be loaded via the serial interface of the Finco. The upper V_{RP} and lower V_{RN} reference voltages are connected from outside. The output voltage V_{out} depends on the reference voltages and on the digital data N, loaded via the serial interface, as follows:

$$V_{out} = \frac{V_{RP} - V_{RN}}{1024} \cdot N + V_{RN}.$$

The measured 10-bit linearity of the DAC is shown in Figure 6.4. An error of ± 5 %, due to the use of a simple voltage multimeter, was assumed for a linear fit. This measurement was part of the PPrD-MCM system test as described in Section 7.

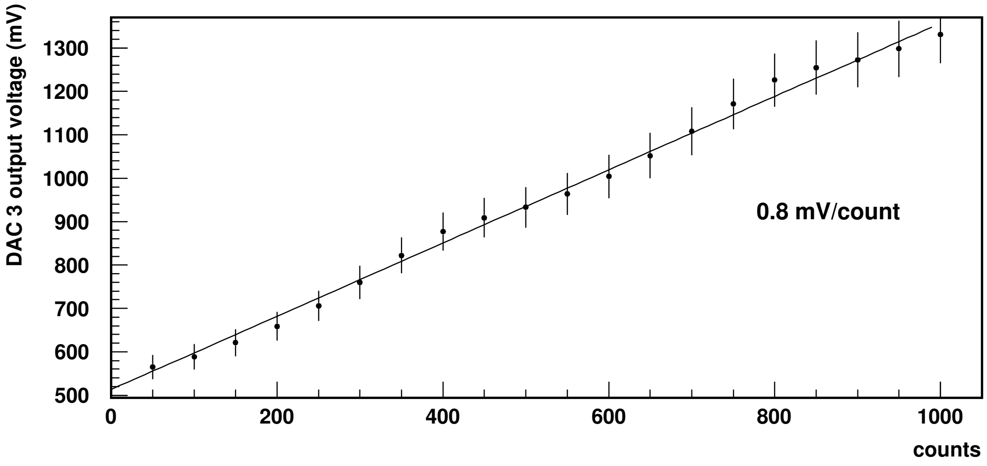

Figure 6.4: Measured DAC linearity.

6.1.3 TTL to PECL level-conversion

The FeAsic uses TTL logic levels for all of its I/O pads. In order to match the logic levels from the FeAsic to the G-link transmitter chips, the TTL levels must be converted to positive 100K ECL levels (PECL). This requires level-conversion for 4×8-bit data. Converter chips are commercially available, but only[5] with four channels per die. This would require eight converter chips on the demonstrator Multi-Chip Module. Each 4-channel converter is itself small in size, but the effective area including wire bonding represents a considerable fraction of the total MCM area. Furthermore, small quantities of commercial components are often expensive and not available as dies from distributors.

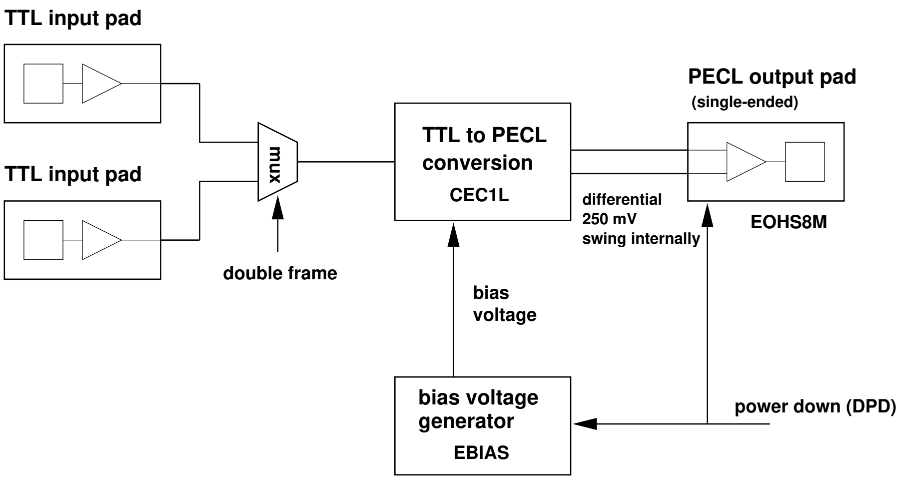

Figure 6.5: Illustration of the double-frame data multiplexing and level conversion.

[5]This was at least the case when the Finco ASIC design was started; it may have changed meanwhile.

Figure 6.5 illustrates the multiplexing and level-conversion inside the Finco. An AMS standard cell (CEC1L) is used for the conversion to differential PECL signals internally, and a PECL output buffer (EOHS8M) converts to single-ended signals with enough current for external signal termination. An active-high power-down signal (DPD) is used to power-down the driver section of standard cells and a bias voltage generator (EBIAS) is used to bias temperature and process dependent parameter drifts. One bias reference is used for 8 level-conversion cells.

6.1.4 Temperature monitoring

The layout of a bias voltage generator (EBIAS) was modified for temperature monitoring. The internal temperature compensating circuit was extended by a current mirror. This is shown in Figure 6.6. The temperature dependent current can then be measured externally. The simulated temperature drift is 0.284 μA/°C. Measurements are shown in Figure 6.7 for two Finco chips (chip 1 and chip 2). The measured current temperature drift is (0.26 $\pm$0.003) μA/°C.

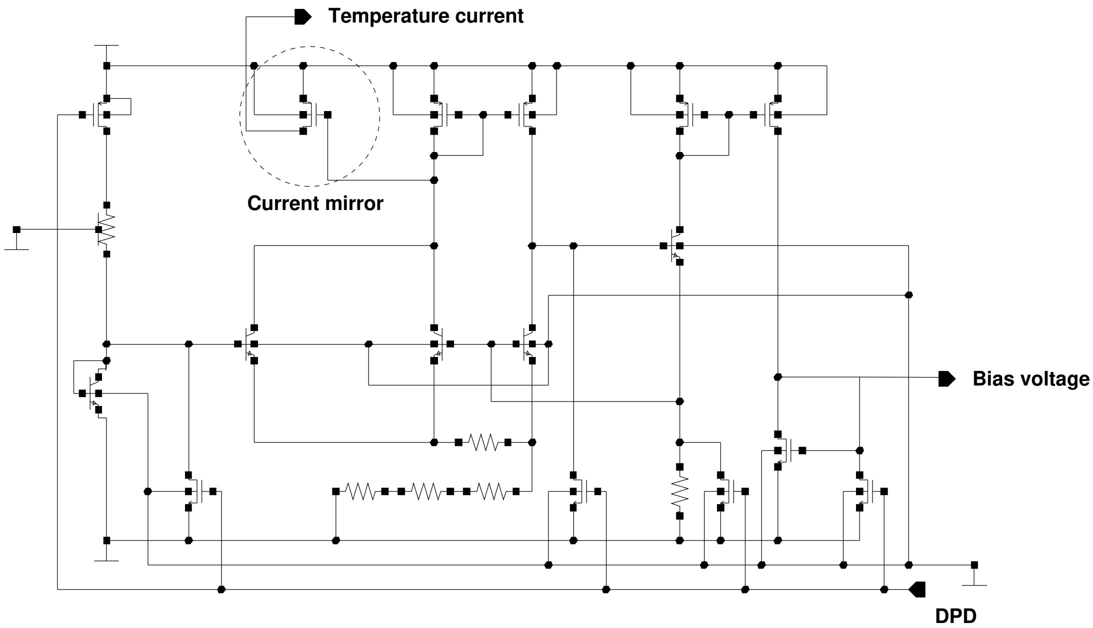

Figure 6.6: Schematic of the temperature sensor.

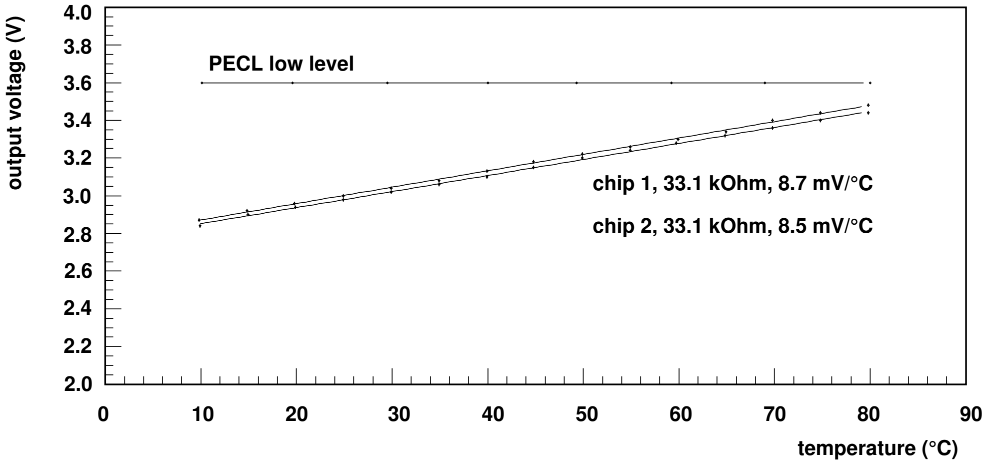

Figure 6.7: Temperature measurements of the integrated temperature sensor. The temperature dependent current was measured across an external 33.1 kΩ resistor.

6.1.5 Integrated differential line receivers

The differential line receiver circuit is shown in Figure 6.8. The circuit is based on a fast operational amplifier (OPVIDEO) with high unity-gain bandwidth of typical 47 MHz. The baseline offset is adjusted with a 10-bit standard cell analogue-to-digital converter (DAC10). The high impedance output is buffered with a second OPVIDEO amplifier.

A simulation of that circuit is shown in Figure 6.9, where a 10 MHz sine-wave was used as input stimulus. Depending on the loaded DAC count N, the output voltage can be calculated as:

$$U_{out} = (U_{pos} - U_{neg}) + \frac{5\,\mathrm{V}}{1024} \cdot N.$$

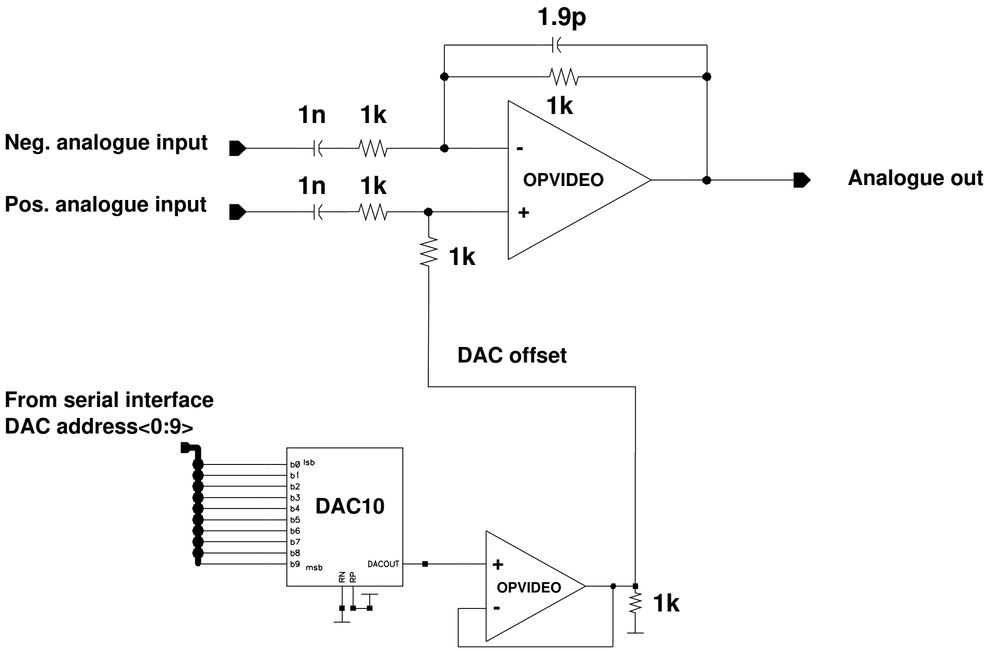

Figure 6.8: Differential line receiver circuit.

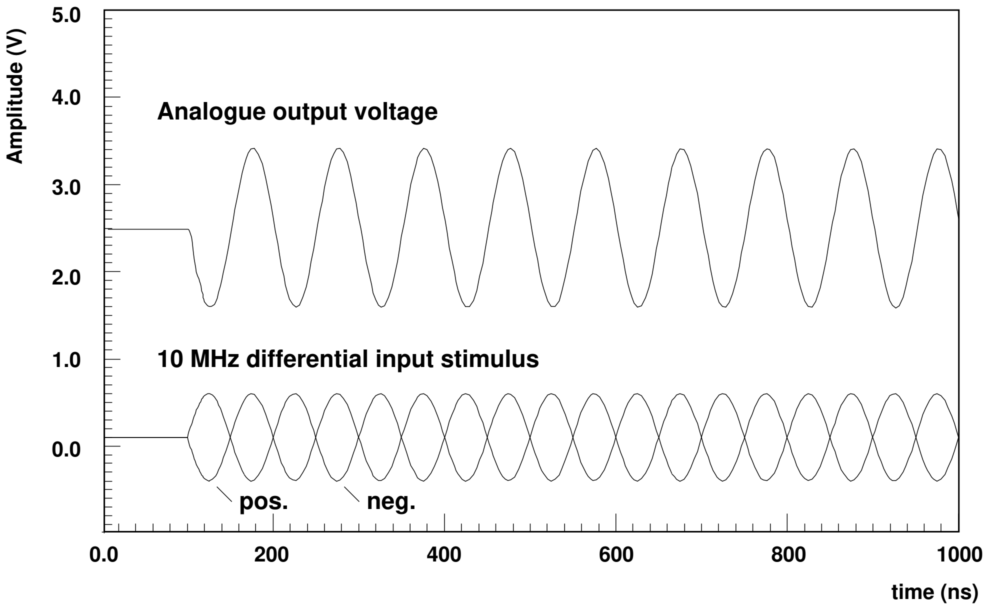

Figure 6.9: Differential line receiver circuit simulation.

6.1.6 Flip-Chip mounting

The standard bonding technique for chip mounting on a Multi-Chip Module is wire bonding. This technique uses about 25 μm thick aluminium or gold wires to connect chip pads to bonding pads on an MCM substrate. The wire bonding pad design requires design rules which must be considered for: pad size, pad pitch, distance to the chip edge, and the bonding angles. The most important constraint for the chip footprint size is the distance to the chip edge. About 1 mm space is required for a bonding tool to place a proper wire bond. These constraints must be taken into account when designing the footprint. Hence the actual area for each chip is increased and may be a limiting factor when high density MCMs are required.

An attractive bonding alternative is Flip-Chip bonding, which allows a high packaging density of ICs on an MCM substrate. Compared to wire bonding the size of a Flip-Chip footprint is equal to its chip area. Beside this packaging argument, Flip-Chip bonding has the following advantages:

- A Flip-Chip is capable of handling a higher number of I/Os because solder bumps can be arranged in an area array rather than being restricted to the chip's periphery. The MCM manufacturing process is therefore the limiting factor in I/O pad placement.
- Due to solder surface tension, this technique has a self-aligning capability during bonding.
- Shortest interconnection distances reduce electrical parasitics and provide excellent electrical performance for Flip-Chip solder joints.
- Flip-Chip bonding is more reliable, which improves bonding yield and therefore the overall MCM yield.

Disadvantages of Flip-Chip bonding are related to infrastructure and not technical issues. More process steps are required. The most efficient Flip-Chip bonding process in terms of cost and process expenditure is 'wafer bumping', where the chips need to be obtained as wafers. For application specific ICs (ASICs) this is often not of concern, because ASICs can be ordered as wafers anyway. If commercial chips are required in wafer form, it may not be possible to get them from distributors. In that case Flip-Chip bumping of single chips is required, where other processes are more suitable: thermocompression[6] or adhesive[7] bonding. Then chip alignment is needed to deposit solder bumps for each chip individually.

[6] Thermocompression is a bonding technique where joining requires a high bonding force and heat. This process offers no self alignment.

[7] Adhesive bonding requires apply of glue. Chip alignment and glue curing are required.

Wafer bumping process parameters	
Chip bumping yield	>95 % (typ. 98 %)
Chip mounting yield	>95 % (typ. 99 %)
Minimum bump pitch	200 μm
Bump diameter	~100 μm
Reflow temperature	~230 °C

Table 6.1: Flip-Chip process parameters for a 200 μm pitch [IZM]. The numbers can be improved if further process adjustment is performed.

The Finco wafer bumping, and the mounting on the PPrD-MCM substrate, was performed by the Fraunhofer Institute in Berlin [IZM]. Table 6.1 is a collection of process parameters they maintain for a 200 μm pitch, and the following is an enumeration of process steps which where applied to the Finco wafer:

1. **Deposition of solder under bump metallization (UBM):** Either electroplating or electroless (autocatalytic) deposition of a solderwettable material is necessary, e.g. Cu or Ni. Electroplating requires a potentialdifference between a sputtercoated wafer and a reference anode within a metal electrolyte. The UBM also acts as a diffusion barrier to the solder. In the present case, electroless deposition of Ni was used to form the under bump metallization.

2. **Deposition of solder paste:** Solder deposition was done using stencil printing technology. Through the apertures of a stainless steel stencil, solder paste is printed and thereby selectively deposited on the under bump metallization of the pads.

3. **Reflow soldering:** The bumped wafers are heated to promote solder reflow, with the result of formation of controlled-height solder bumps over the I/O bonding pads.

4. **Wafer cleaning:** This is required to remove remaining solder materials, e.g. solder flux.

5. **Wafer cutting:** At this stage the wafer is cut into individual chips and singulated into trays.

Process steps related to the MCM assembly are as follows:

1. **Chip mounting:** Once the die is positioned on the pre-fluxed substrate, the entire assembly is placed in an oven. The applied thermal profile will form the actual solder joints.

2. **Shear testing/X-ray inspection:** The adhesion of the solder bumps can be determined quantitatively by using a shear tester (destroying the joints), or qualitatively by using X-ray microscopy (nondestructive).

3. **Under-filling:** Finally, the chips need protection from the atmosphere, mechanical shock, moisture, and chemicals used in the manufacturing process. Most of this would already be taken into account by global MCM encapsulation, but this process cannot fill the small gap between chip and substrate. Viscose under-fill materials are required instead, to fill a gap of about 100 μm. Further more, underfill materials with thermally conductive fillers can improve the thermal conductance for the chip, e.g. AlN, BeO, diamond powder.

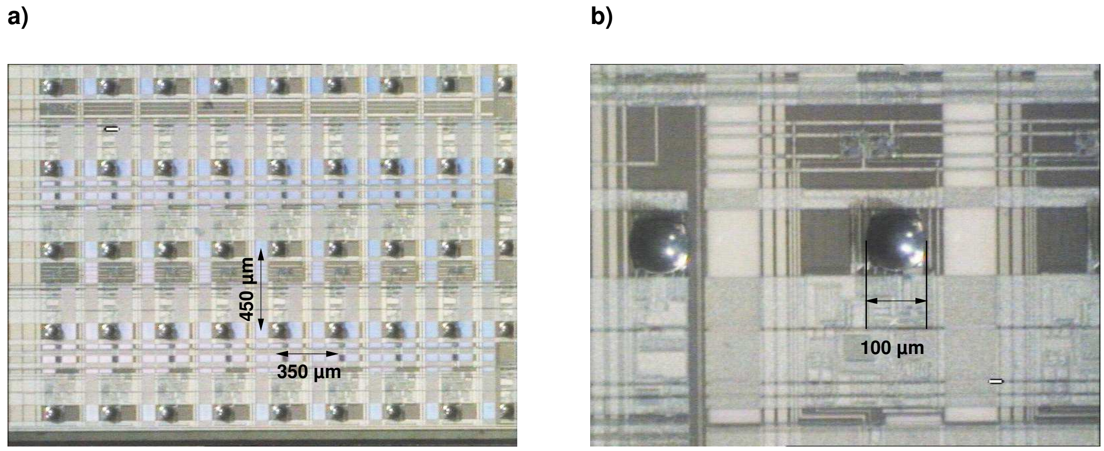

Figure 6.10: Flip-Chip solder bump on a Finco pad. Pad spacing is shown in a) and the solder bump size in b).

For Flip-Chip bonding it is not necessarily required to have I/Os arranged in an area array. Often a chip design is ‘core limited’, where the pad pitch is wide enough and most of the chip size is taken up by the inner chip logic. In contrast to that, the Finco ASIC layout is ‘pad limited’. It has 143 I/Os and hence its MCM footprint size was optimised by an array placement of I/Os. The minimum pad pitch is then given by the design constraints of the MCM-L process, which is 350 μm and 450 μm due to the rest-ring of micro-vias. Solder bumps and their arrangement are shown in Figure 6.10 a) and b).

The Finco layout was designed with Composer for schematic entry, Synergy for Verilog HDL synthesis, Preview for floorplaning, and Analog Artist for mixed-signal simulations. These tools are part of the Cadence design workframe [Cad]. The floorplaning software tool Preview is convenient to use, if digital standard-cell designs are made. But for auto placement and routing, it assumes that the user wants to have the I/O pads at the chip periphery. It does not support the placement of I/Os in the chip centre, where standard cells have to be placed around I/Os. This fact made it necessary to do most of the Finco chip layout 'by hand', which is more time consuming and requires extensive design rule checks.

Considerations for the PPrAsic

The PPrAsic will be a purely digital design, and it uses other floorplaning tools which may support placement and routing of standard cells between I/Os. Even then, it is much easier for floorplaning to have a single digital core block in the chip centre. Wafer bumping requires at least a 200 μm pad pitch. The chip perimeter defines the maximum number of I/Os. In the case of wire bonding, this number is limited by the pad cell width, which is about 110 μm in an AMS 0.6 μm CMOS process.

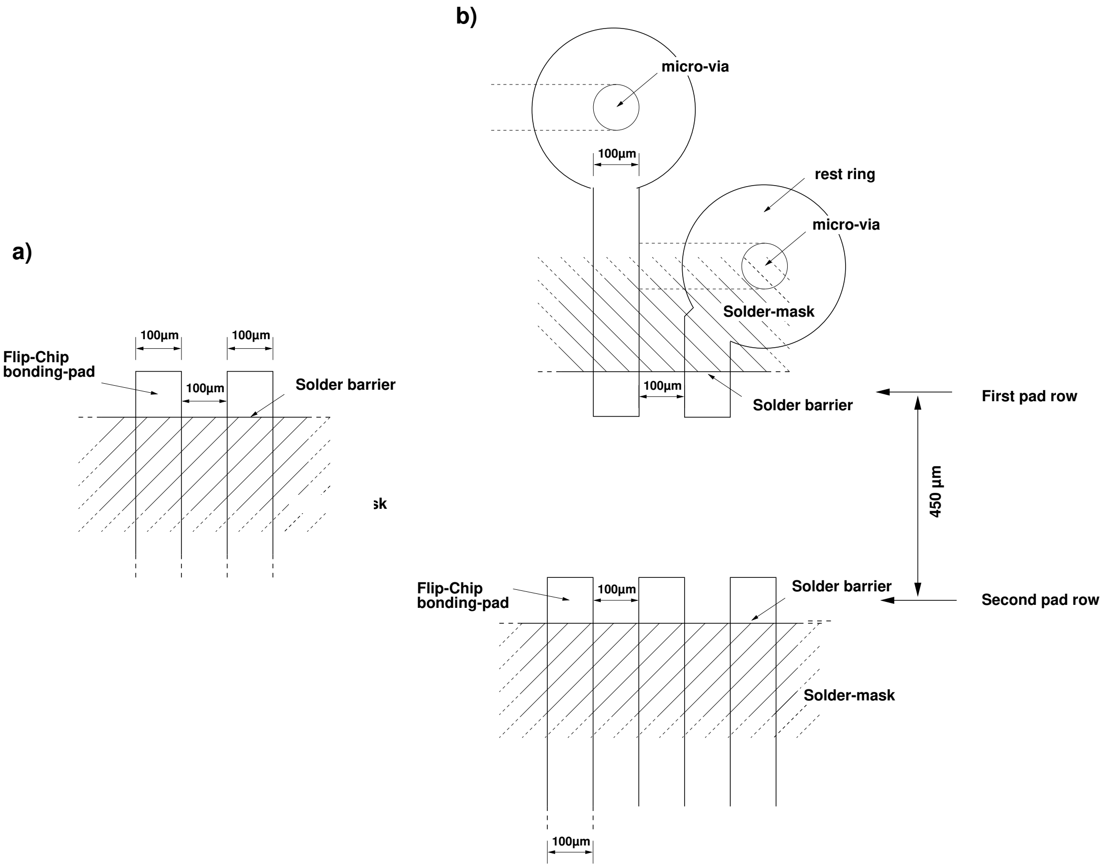

Figure 6.11: Smallest Flip-Chip pad layout suggested for the final Pre-Processor MCM. A single pad ring layout is shown in (a) and a staggered pad layout in (b), where a second pad ring is used at a distance of 450 μm. Both layouts have a minimum pitch of 200 μm.

If the pad pitch is smaller than 200 μm, the pads could be staggered in order to enlarge the pad-to-pad distance. This could be done by sliding all even or odd numbered pads in such a way that a second pad ring is formed. Then the pitch is enlarged but the chip area is slightly increased. A second pad row on one side of a chip would require an additional 450 μm space, which is about the length of a pad in 0.6 μm CMOS technology. In the MCM-L technique vias can then be placed in the middle of the chip to make contact with deeper layers. Figure 6.11 a) shows an example for the MCM bonding pad layout if only a single pad ring is required. Figure 6.11 b) shows the staggered bonding pad layout on the MCM substrate as an alternative. Because of the good solder wetability of bonding pads, a solder-mask is required to act as solder barrier.

Figure 6.12 shows the reduction of footprint area for an MCM substrate if Flip-Chip bonding instead of wire bonding is used. Four steps can be seen. Each step corresponds to the opening of a new pad row for each of the four chip sides. For a core area of 50 mm^2, the reduction of the footprint area is 36 % for more than 110 pads. In this region the design is core limited, and Flip-Chip bonding with a pitch of 200 μm is possible without additional pad rows.

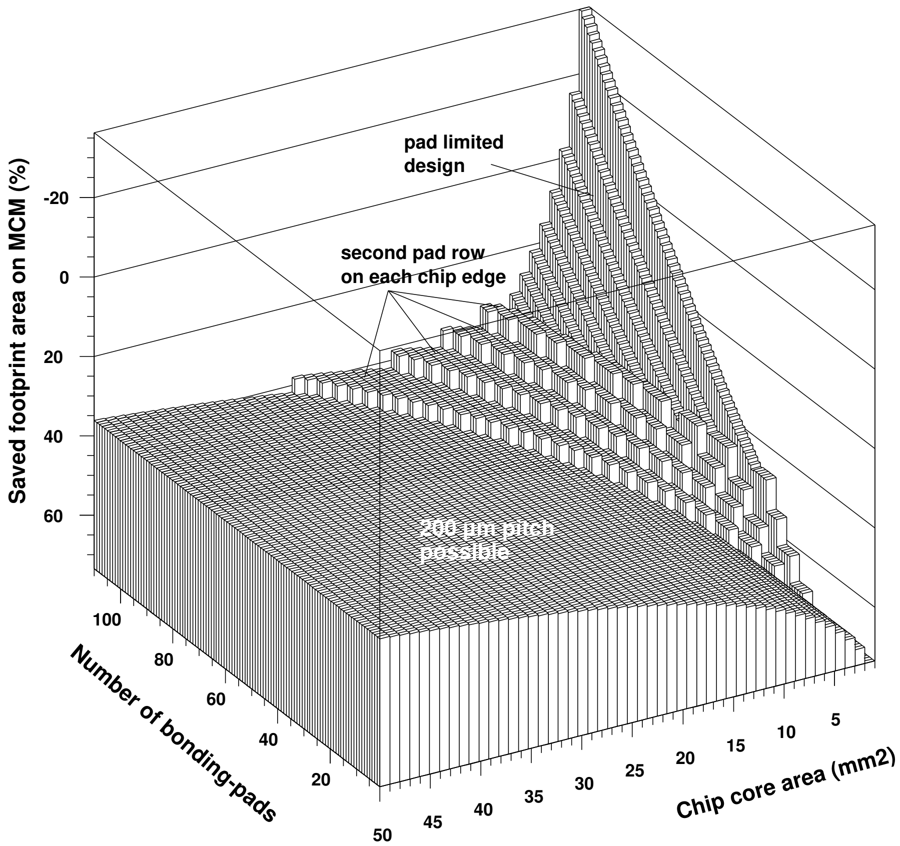

Figure 6.12: Saved space for chip footprint on a MCM layout.

Via-in-pad technology for Flip-Chip bonding

The Finco layout was based on the via-in-pad technology applied to all signals underneath the chip. Because of the 100 μm via hole, the Flip-Chip bonding pad needs to be filled with solder before Flip-Chip mounting. This is one extra process step, which should be avoided. Vias cannot be filled uniformly, and therefore the effect of self-alignment of solder bumps is reduced. Figure 6.13 shows a picture taken after via filling. Paint was put on track stubs connecting the bonding pads, in order to build a solder barrier. This must be avoided for Flip-Chip bonding of the PPrAsic, and a solder mask is required close to the bonding pad to act as a solder barrier.

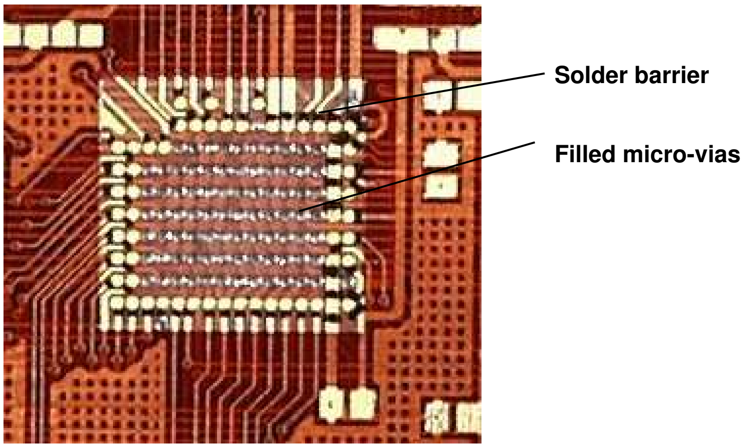

Figure 6.13: Filled micro-vias and solder barrier paint [IZM].

6.1.7 Boundary-scan (JTAG)

Boundary-scan is formally known as the IEEE/ANSI 1149.1-1990 standard. It is often referred to as JTAG[8] and defines a collection of design rules applied to a digital integrated circuit (IC), that allows software to perform automated printed circuit board testing [Ken94].

The Finco boundary-scan is controlled with a finite state machine, which is included in the digital 'core' logic block of the ASIC. This state machine is called Test Access Port controller (TAP). The TAP controller has three input pins: a Test Clock pin (TCK), a Test Reset pin (TRST), and a Test Mode Select pin (TMS). The states of the TMS controller are changed by defined sequences of the TMS pin state. A sequence allows switching between two internal registers: the *instruction register* and the *data register*. The instruction register is used for loading commands to the TAP controller, whereas the data register can be selected as either an *ID*

[8] JTAG: Joint Test Action Group, a group made up of companies in Europe and North America.

register or a *boundary-scan* register. From the ID register one can read a unique chip ID, whereas from the boundary-scan register one can read from, or write to. The boundary-scan register is made up from special pads containing flip-flops and multiplexers. The flip-flops are connected together to form the actual ring-like boundary-scan register. Multiplexers inside the pads are used to change between two alternative test modes:

- *Internal test* performs an internal chip test only. In this case the pad states are only visible to the internal chip logic and external asserted pad states are ignored; or
- *External test* of the chip environment. This mode scans the input to the chip and allows pre-loading of chip output pad states.

Once the boundary-scan register is selected as data register its content can be shifted out serially. During that, it is connected between the Test Data Input pin (TDI) and the Test Data Output pin (TDO). The Finco boundary-scan register gives access to 70 pins: all SpyBus input pins from the FeAsics, all serial interface pins, and all TTL input signals from the FeAsics are included. The PECL output pins are left out because the PECL output pads do not include boundary-scan logic.

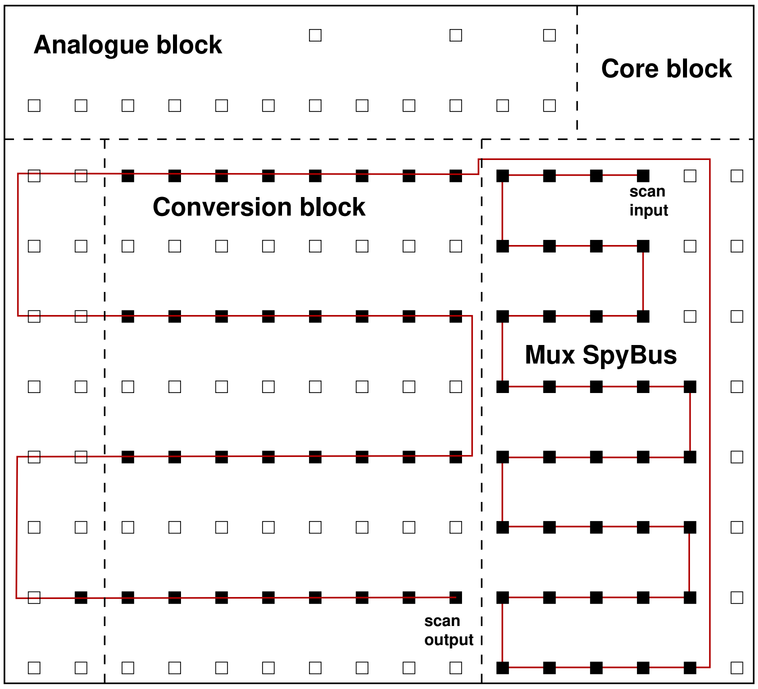

Figure 6.14: Illustration of chip layout and boundary-scan path. All scanned I/O pads (70) are shown in black.

Figure 6.14 illustrates the location of the boundary-scan path. Each pad in the chain contains a flip-flop which has a defined location in the boundary-scan register. The data are shifted in at the 'scan input', labelled TDI, and are shifted out at the 'scan output' pin labelled TDO. For a definition of the 70-bit-wide boundary-scan register content see Appendix C, Figure C.1. The finite state machine of the TAP controller is illustrated in Figure C.2 and control tokens for the instruction register are defined in Table C.6.

6.2 Finco layout

Figure 6.15 shows the layout of the Finco ASIC. It consists of four building blocks: one analogue block, one digital standard-cell block, one level-conversion block, and the SpyBus multiplexing block. Power pads are placed in a row on the left edge of the chip. Power and ground signal routing was of concern during the layout, because of the high power dissipation of the conversion logic and the PECL output buffers. The conversion block has the TTL inputs adjacent to the PECL output pads.

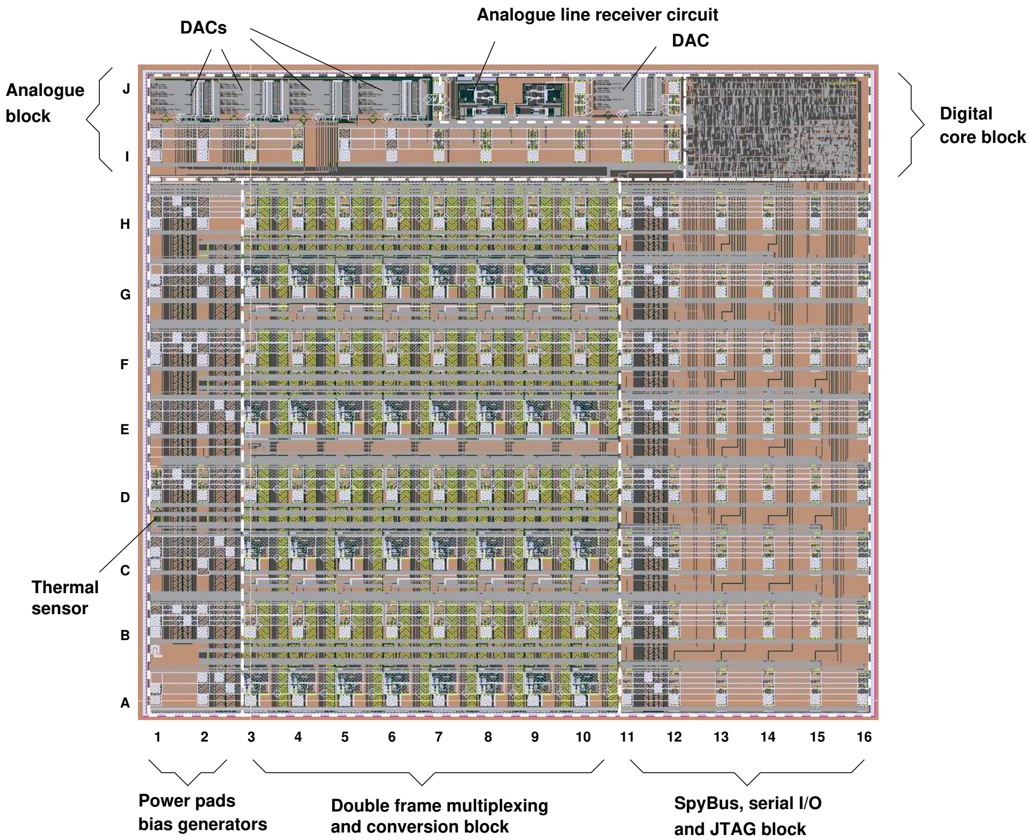

Figure 6.15: Finco layout.

6.3 Test results

For the Finco chip test a matrix needle-probe card was used to contact each of the 143 I/O pads to a chip tester. Test 'vectors' were created, and the chip tester allows them to be applied to the chip. Then the output is compared to response 'vectors', which were obtained from simulations of the chip circuit. A large number of pins, where each pad requires its own chip tester channel, is a problem for chip tests. Only 64 chip-tester channels were available, and hence the JTAG interface was used for loading of test vectors without connecting the pads to a tester channel. The use of the JTAG interface slowed down the chip test, but full-speed tests at 40 MHz were not performed anyway, because of a DC probe card and long cable connections from the probe card to the chip tester. Figure 6.16 a) shows a picture of the matrix probe card manufactured by GPS Prüftechnik mbH and Figure 6.16 b) shows the test results from a wafer test. The chip yield was 61 % for this wafer. Sixteen short circuits where found in the power supply, 36 defective boundary-scans were observed through the TAP controller, some stuck bits have appeared in the conversion block, and errors were found in the analogue circuits.

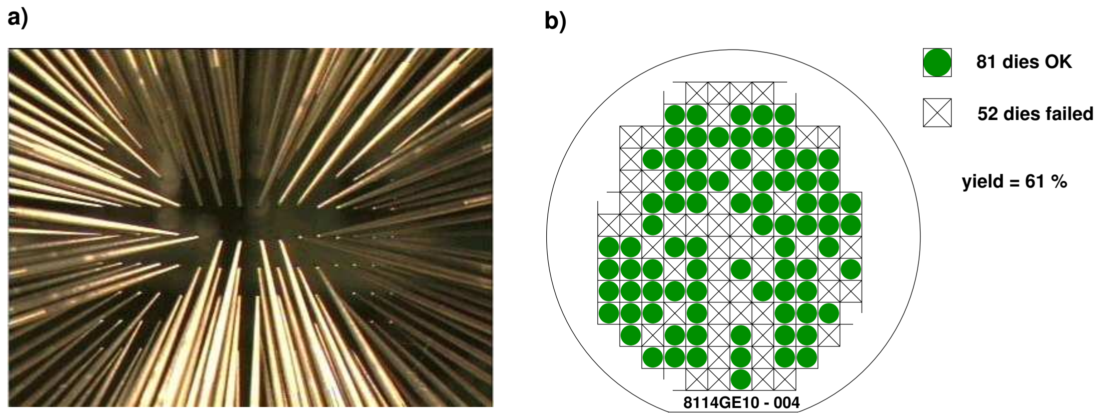

Figure 6.16: 143-needle matrix probe card used for the Finco chip test (a). A wafer test result is shown in (b).

The final chip test at a speed of 40 MHz was performed as part of the modular Pre-Processor test system, which will be described in the following Chapter 7. FeAsics data have passed this chip and regarding the Flip-Chip mounting, no connection problem were observed. A number of temperature cicles have been performed for MCM temperature measurements with no effect on the solder joints. The mixing of analogue with digital standard-cells had no effect on the quality of the analogue voltages used for baseline adjustment. Finally, the MCM test has shown the functioning of the Finco ASIC, because, without this chip the PPrD-MCM would not work at all.

Chapter 7

A modular Pre-Processor test system — measurement results

- Modular test system overview
- VME-board configuration
- Control and diagnostic software
- MCM system measurements

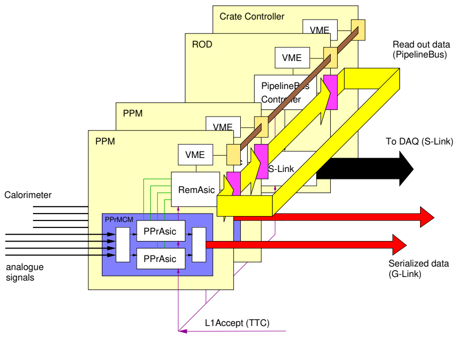

7.1 The modular test system

This chapter starts with a short overview of earlier electronic developments for the Pre-Processor. Then the modular Pre-Processor test system is described, which is one step closer to a final prototype system. This system is able to test the key components described in Section 3.3, and hence the two main tasks of the Pre-Processor: the readout of eight Pre-Processor Modules and the compact pre-processing of trigger tower signals on a Multi-Chip Module.

Section 7.1.3 describes the modularity of the hardware and Section 7.1.4 the control and diagnostic software, which is used for system tests. At the end of this chapter, system measurements will be presented which have proven the functioning of the PPrD-MCM.

7.1.1 Earlier electronic developments

In the year 1992, initial thoughts about the design of the LHC experiments have led to a research and development programme for the trigger system, RD 27. At this early stage, digitisation and preprocessing for the trigger system was supposed to be located on the detector. The first electronic development was a prototype ASIC for carring out an electron/photon cluster algorithm in a pipelined way. Heidelberg's development work on the Pre-Processor began in 1995 with a four channel flash analogue-to-digital converter (FADC) [Han95]. This FADC module has contributed to various laboratory and test-beam experiments. In parallel, alternative Calorimeter Trigger designs were investigated. A design based on a very fast and densly-packed processing ASIC was considered as an alternative Calorimeter Trigger implementation. An early test version of such an ASIC (TAsic) was tested [Wag96], with a input data rate of 800 Mbit/s.

Since then, the RD 27 program has evolved into the current ATLAS trigger program. The location of the digitisation electronics has been moved away from the detector to an external trigger cavern. Hence, a analogue signal transmission of trigger tower signals was needed. An analogue optical transmission link was developed, with a linear resolution of 8-bits [Pfe96]. An ASIC impementation of BCID for non-saturated trigger tower signals was build based, on previous FPGA implementations [LEB95]. A first Pre-Processor prototype ASIC (FeAsic) was manufactured in the year 1996 [Fea96], [LEB96]. A test module used as test platform for this ASIC and its readout concept has been designed and built. This module is still in use and is called the 'famous' Front-End Module (FEM) [Sch97]. It was integrated into the ATLAS trigger prototype data acquisition software and its performance has been proven in various laboratory tests. It was used in combination with the RD 27 FADC and it was part of a 'full-slice' test of Level-1 Trigger prototype modules [TDR98] at CERN. The performance of that module has been published at [LEB97].

Further developments were needed for the readout of the Pre-Processor system. The Readout Merger ASIC was designed and manufactured in 1998 [Rem98]. Since then, the system concept has evolved towards the final system design, as it was described in Chapter 3.

7.1.2 Modular test system overview

The philosophy of the modular test system is to build a small-scale version of the Pre-Processor based on the prototype components which exist so far. This includes the RemAsic, a PipelineBus ring implementation, and the preprocessing inside the demonstrator PPrD-MCM. The test system consists of a general purpose motherboard and several daughtercards, which provide the special functionality. Flexibility and modularity is achieved by the use of FPGAs and CMC[1] type daughtercards. Because of its modularity, this test system can act as a platform for new technology developments, with a short turn-around time, e.g. for the final PPrAsic, the final PPr-MCM, or LVDS link tests.

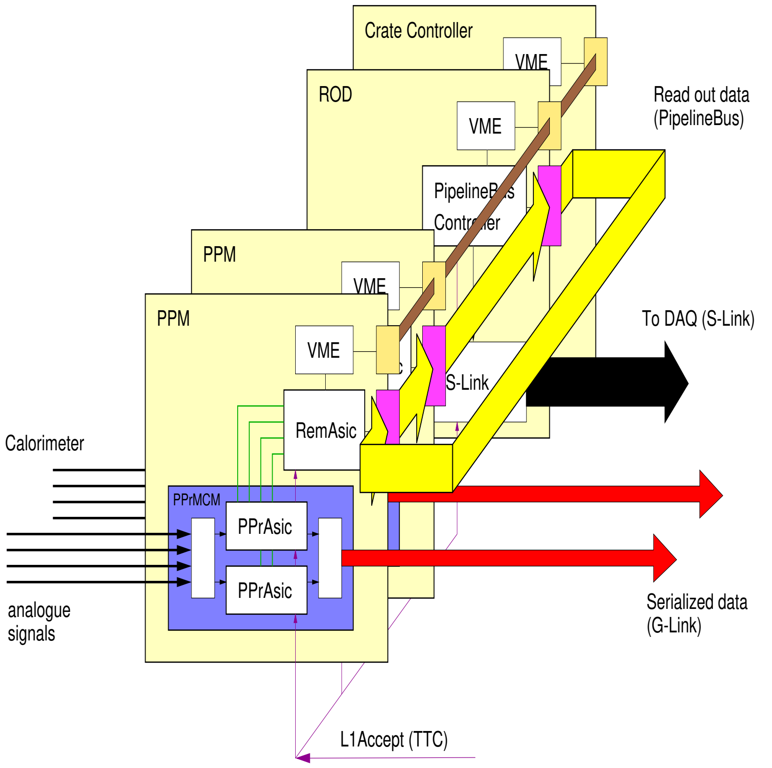

Figure 7.1: Illustration of the modular Pre-Processor test system [Sch98].

Figure 7.1 illustrates the scope of the test system. It consist of up to eight PipelineBus readout nodes. Three nodes can be equipped as Pre-Processor units,

[1]CMC: Common Mezzanine Card

whith a PPrD-MCM and a RemAsic CMC card on the general purpose motherboard. The other nodes are spectator nodes, which increase the PipelineBus length in an economical way. Another motherboard is configured as a Readout Driver module (ROD). This module has a standard S-link CMC card [Sli] and a PipelineBus master CMC card on it. The crate controller is a PowerPC computer (RIO2) manufactured by CES, running under the Lynx Os operating system.

Emphasis will be placed on the timing and synchronisation of the modules. A level-1 accept signal can be generated to trigger the readout of all Pre-Processor units at the same time. A G-link receiver CMC card can generate these level-1 accept signals dependent on its input data. The level-1 accept signal, the 40 MHz LHC bunch-crossing clock, and the level-1 bunch-crossing number will be fanned-out by an I/O control CMC card.

The following section describes the general purpose motherboard and the set of CMC daughter cards which exist. Then the configured VME-boards are described.

7.1.3 Modular VME-board configuration

The basis of the modular Pre-Processor test system is a general-purpose motherboard. Several daughtercards, which provide the special functionality, can be plugged into two CMC card slots. The top slot 1 can carry CMC cards with either a standard size of 149 mm×74 mm or with a wider front panel extension of 100 mm. The bottom slot 2 can only carry CMC cards which have a standard form factor. The slot 1 extension is needed for PipelineBus connectors. Two Compact PCI connectors, with 110 pins each (5 rows×22 pins), are used for 70 PipelineBus I/Os (35 inputs, 35 outputs), control signals, and grounding pins. The motherboard has the following features:

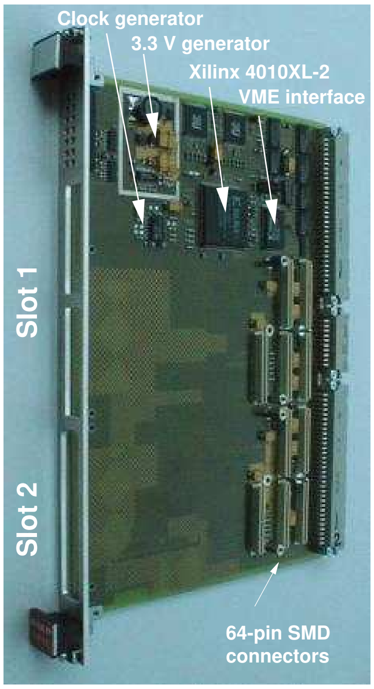

VME motherboard: This VME module is used as a common CMC card carrier. Its dimension is 6U (233 mm) in height and 160 mm in depth. It provides two CMC card slots. One slot with extended front panel space and a second slot of normal height (149 mm). Its functionality is programmable through a Xilinx FPGA (4010XL-2). For each CMC card configuration an associated FPGA design is required. The address space of an on-board 32 kByte Dual-Ported memory is mapped to the VME bus through the use of an A24D16 VME bus interface. An on-board clock generator provides two independent clocks up to 120 MHz. The module requires a –5 V supply from the backplane and it generates its own +3.3 V supply on-board.

CMC cards	CMC components	Mother-FPGA
RemAsic	1 RemAsic, bus interface	required
Master/Spectator	bus interface	required
I/O control	1 CPLD, bus interface	—
Pre-Processor	1 PPrD-MCM	required
S-link	1 S-link	—
Readout Driver	1 FPGA, 1 DP-memory, bus interface	—
G-link tx/rx	2 G-link chips (2tx, tx+rx, 2rx)	required
LVDS tx/rx	2 LVDS chips (2tx, tx+rx, 2rx)	required

Table 7.1: CMC daughtercard overview. CMC components and the motherboard FPGA resources are listed.

Table 7.1 is a listing of all CMC daughter cards, which will be used to add the special functionality for the test system. At this point one has to stress that, without the enthusiastic activities of the Heidelberg electronics group, this test system with many different electronics modules would have not been possible [KS99]. Most of these CMC cards already exist, and are described in the following:

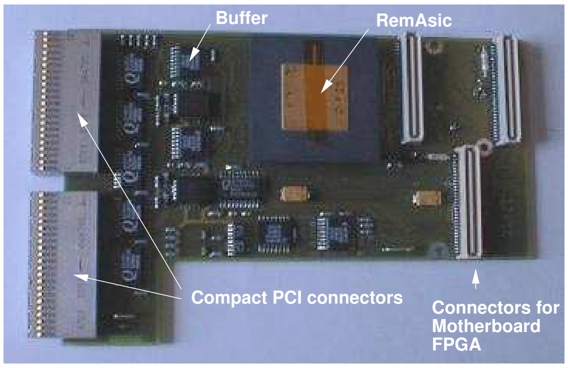

RemAsic CMC: The RemAsic CMC card consists of one Readout Merger ASIC. Two compact PCI connectors are interfaced to the PipelineBus. Buffers are used to provide enough current for a fast bus-signal risetime. The communication (readout and configuration) of the PPrD-MCM is done through a direct connection via the motherboard.

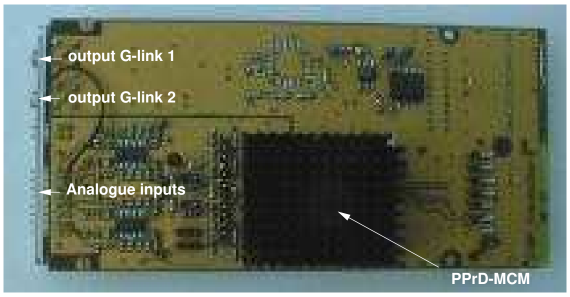

Pre-Processor CMC: This CMC card carries one demonstrator MCM. It does all the preprocessing required for four trigger tower signals. It receives differential analogue signals, it performs digitisation, digital preprocessing, and the high-speed serialisation at 800 MBd or optionally at 1600 MBd.

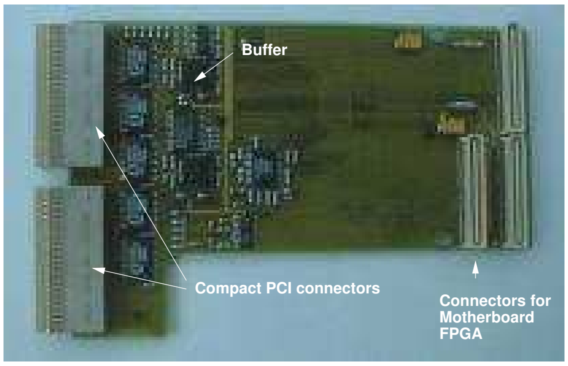

Master CMC: The Master CMC card has two compact PCI connectors for the PipelineBus interface at the front panel. The Motherboard FPGA adds the functionality of a master node to that card. All signals are buffered to provide enough current for a fast signal risetime of the 5 V TTL bus signals.

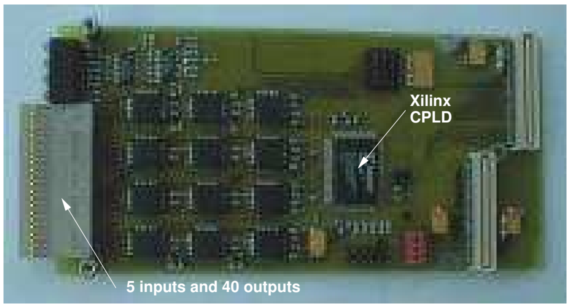

I/O control CMC: Trigger timing and control signals are generated from this CMC card: the level-1 accept signal, the level-1 number, the 40 MHz bunch-crossing clock, and PipelineBus control signals. Five programmable ECL inputs can be fanned-out to 40 ECL outputs in any combination.

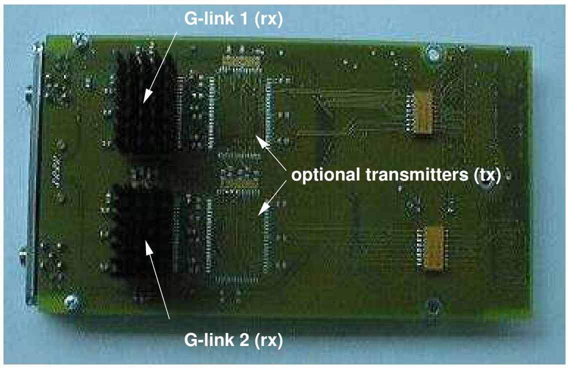

G-link CMC: This CMC card can either be used as a dual G-link transmitter or receiver card, or it can be used as transceiver, with one transmitter and one receiver chip on it. What is chosen depends on the card assembly. Both G-link data rates are supported: either the 800 MBd or the 1600 MBd.

Slot configurations for a motherboard are summarised in Table 7.2. Not all combinations are possible because of the limited number of I/O pins of a motherboard FPGA, which must be shared between two CMC slots.

VME-board unit	Slot 1 (bus slot)	Slot 2
Master	Master/Spectator CMC	I/O control CMC
Pre-Processor	RemAsic CMC	Pre-Processor CMC
Readout Driver	Readout Driver CMC	S-link CMC
Spectator	Master/Spectator CMC	—
G-link	G-link tx/rx	—
LVDS	LVDS tx/rx	—
Fan-out	I/O control CMC	I/O control CMC

Table 7.2: VME-board units and their slot configuration.

For each board configuration the motherboard gets different identity. In the following, these boards are referred to as VME units. The VME units fit into a standard VME crate. All input and output signals are fed in at the front panel. The PipelineBus connects all the modules along the front panel, which is convenient for tests. The functionality of the VME units is described as follows:

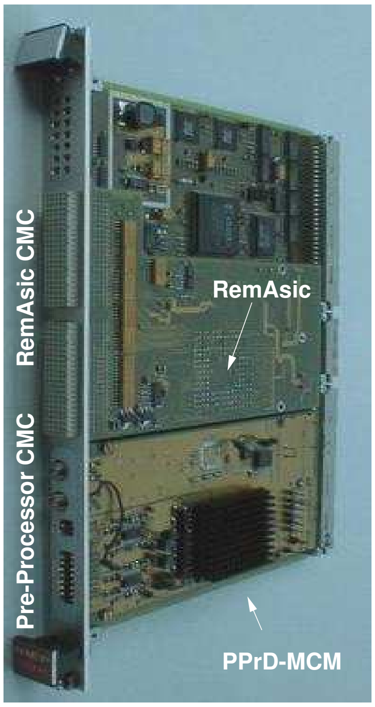

Pre-Processor unit: This module has a similar functionality to the final Pre-Processor Module. For a reduced number of four trigger tower channels it performs the preprocessing and the readout. It consists of a RemAsic CMC card in slot 1 and a Pre-Processor CMC card in slot 2. A motherboard FPGA design is required for the configuration of the PPrD-MCM and for an alternative configuration of the Readout Merger ASIC through a serial implementation of the PipelineBus interface. This configuration of the RemAsic is only used in stand-alone tests without a bus master present. In normal operation, configuration data are received via the PipelineBus. Readout data from the PPrD-MCM are read out, compressed, and moved on the 35-bit wide PipelineBus, upon a request from the bus master.

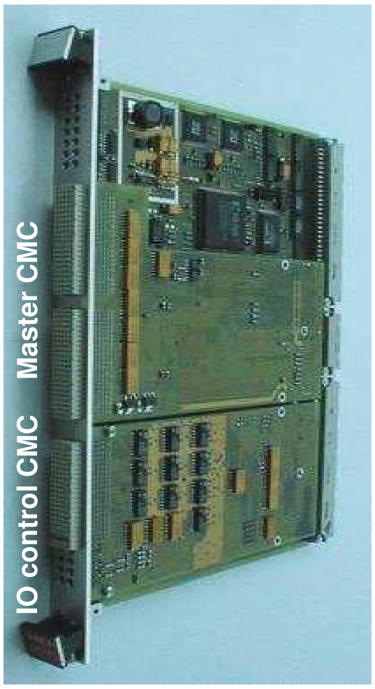

Master unit: This module consists of a Master CMC card in slot 1 and a general purpose I/O control CMC card in slot 2. In this configuration, the motherboard FPGA design is either used for the implementation of a PipelineBus master node or for the implementation of a spectator. A spectator node can be used to spy on the PipelineBus or it can act as a place holder to enlarge the ring for an extended PipelineBus test.

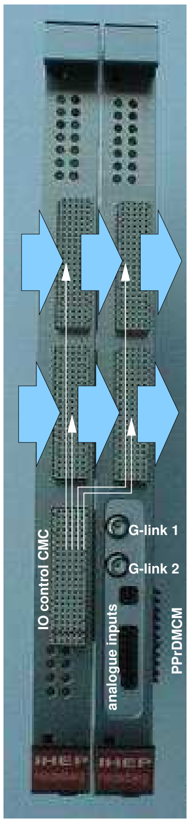

PipelineBus implementation: The PipelineBus ring will be realised as one-slot connections along the front panel. All 35 bus output pins from one RemAsic CMC card are fed into a neighbouring module by the use of one-to-one connections. The one-to-one connections are formed by small printed circuit boards, which are plugged into the Compact PCI connectors at the front panel. Control signals for each bus node are generated by the bus master and fanned-out via the I/O control CMC card. A flexible ribbon cable is used for the transmission of control signals, which are the 40 MHz bunch-crossing clock, the level-1 accept signal, and the level-1 number.

7.1.4 Monitor and control software

An object-oriented software package written in C++ is under development for the test system. This software (HDMC[2]) gives easy access to the hardware with a user-friendly graphical user interface [Sch99]. It provides data histogramming for monitoring and it controls modules present in a test configuration.

A 'part' collection class allows the assembly of modules for definition of a test set-up. Each electronics module requires the definition of dependencies, which accomplish a module in terms of memories, registers, and components. Dependencies can themselves depend on others, e.g. a register definition depends on a memory, which depends on a module, and on the definition of a VME bus. All module definitions are compiled at run time. The bit definitions of registers are read in at run time from text-based configuration scripts. This has the advantage that the software does not need recompilation if bit-field definitions change. For the actual hardware access, a network client/server connection can be used. For interactive hardware access, this network connection is not a bottleneck which slows down the

[2]HDMC: Heidelberg Monitor and Control software

software speed, because only a little data will be transmitted for monitoring and control purpose. Figure 7.2 shows a picture of the GUI interface. The features of the software development can be summarised as follows:

- Purely object-oriented software package written in C++;
- Interactive hardware access via TCP/IP client/server connections;
- Graphical user interface (GUI) is based on QT 2.0 class libraries [QT];
- Platform independent source code (LynxOS, Linux, Solaris, HP-UX);
- Integration of new hardware requires no software recompilation (configuration-script based approach).

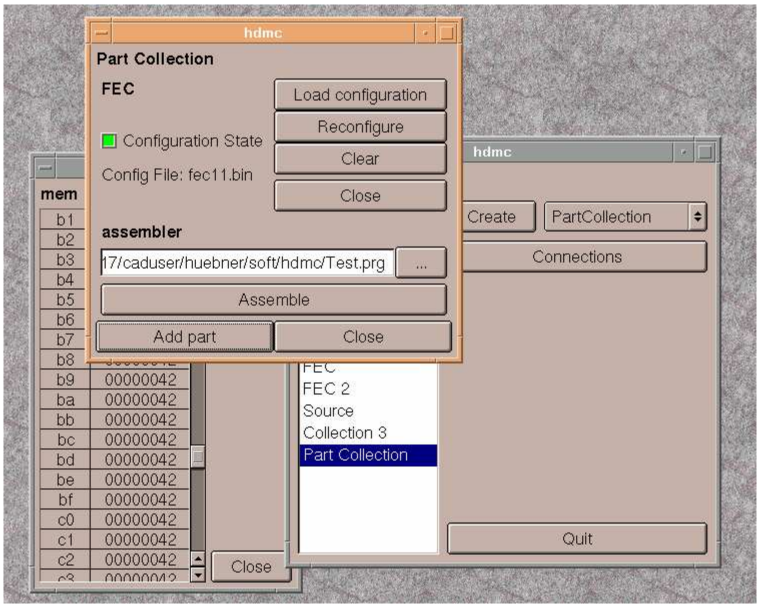

Figure 7.2: The user's view of the HDMC software [Sch99].

7.2 MCM system measurements

This section describes measurements of the PPrD-MCM as part of the modular Pre-Processor test system. The aim was to demonstrate the functioning of the PPrD-MCM, with all its real time preprocessing and its high-speed serial data transmission of trigger towers.

The MCM test set-up consists of two VME motherboards, each equipped with a CMC daughter card. One Motherboard carries a Pre-Processor CMC card and the other one carries a G-link receiver CMC card. Figure 7.3 shows a picture of the test set-up. As input to the Pre-Processor card, a liquid argon-shaped calorimeter signal was generated by an Arbitrary Function Generator (AFG). The G-link output signals from the Pre-Processor card were connected via a 1 m long coax cable to the G-link receiver card. A diagram of that test set-up, with all the processing steps between, is illustrated in Figure 7.4.

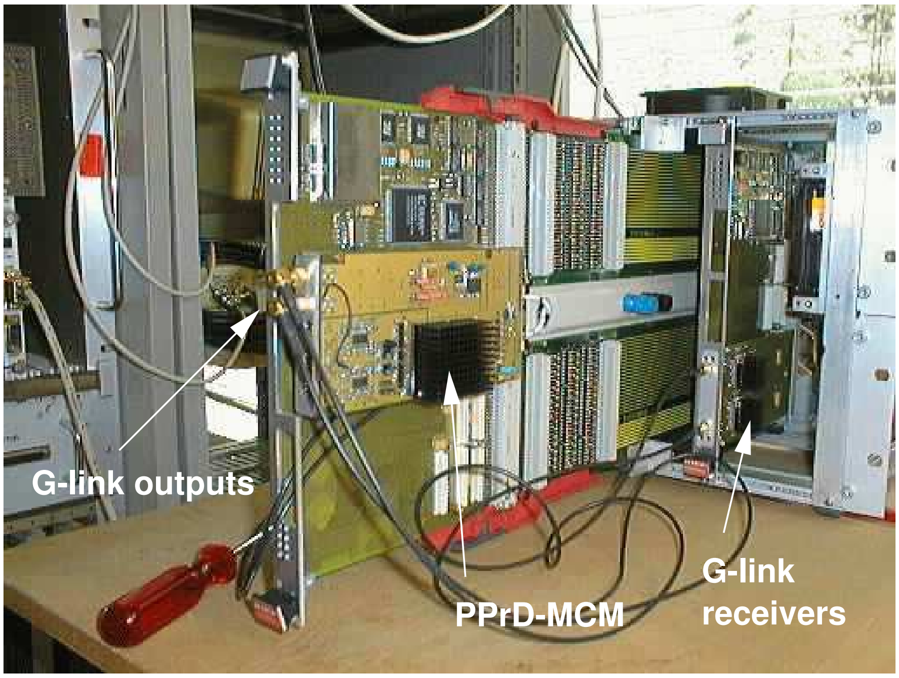

Figure 7.3: PPrD-MCM test set-up.

A differential output signal from the AFG is feed into the line receiver circuit. There, it is mixed with an analogue baseline offset, generated from the Finco ASIC on the PPrD-MCM. This offset-adjusted signal goes straight into the MCM. It is processed by MCM components in the following way: first it is digitised to 8-bit precision, next, the FeAsic performs BCID for non-saturated trigger-tower signals, and then the Finco ASIC converts logic levels from TTL to PECL before the data are serialised at 800 MBd by the G-link transmitter chip. The high-speed G-link output signal is regenerated on the CMC card by the use of a PECL buffer chip. Then it is transmitted along tracks on the printed circuit board up to a SMA front panel connector.

The G-link signal is connected via a 1 m long coax cable to the G-link receiver board. There it is again regenerated by a PECL buffer chip before it goes into the G-link receiver. The G-link receiver requires a 40 MHz reference clock to synchronise to its input data. Its internal phase-locked-loop (PLL) first performs a frequency synchronisation and then a phase synchronisation. If that process

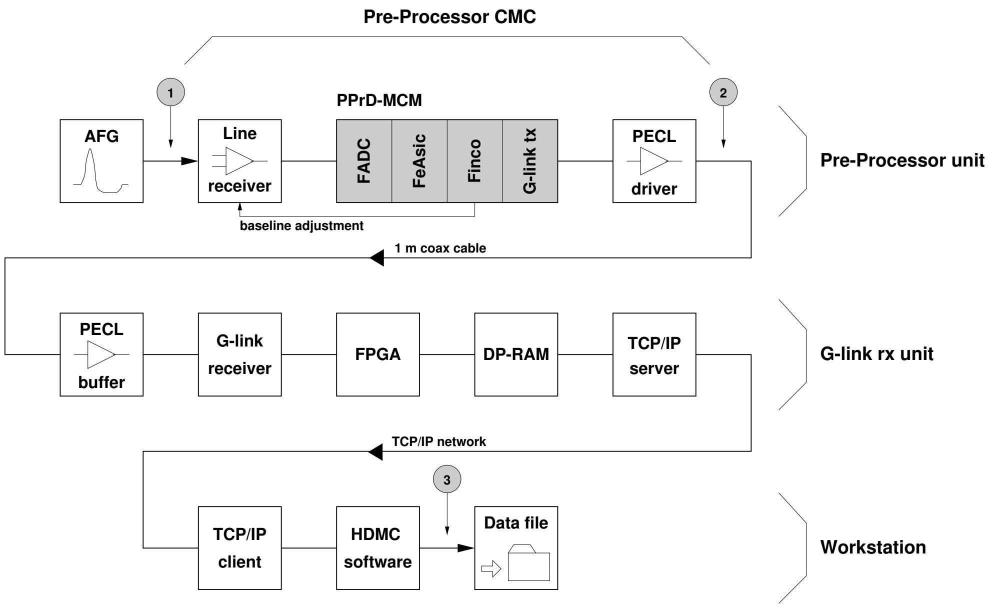

Figure 7.4: Processing chain of the PPrD-MCM test set-up.

succeeds, the G-link receiver will 'lock' to its input data. It decodes the serial bit-stream and provides the parallel output data to the motherboard FPGA. The FPGA latches the G-link data from both G-links (2×16 bits) and writes the data to the motherboard dual-port memory. This is done at a speed of 40 MHz.

The task of the monitor and control software is to read the data from the dual-port memory. A server program running on the crate controller decodes write commands and read commands, which were sent from the HDMC client software. The server program performs the actual hardware access. Data packages are sent back via the TCP/IP network to the HDMC software. There they are stored in a simple text file format.

The serial data stream of the G-link device consists of control bits (C-fields) and of data bits (D-fields). For each clock cycle of 25 ns these bit fields are sent out serially. The G-link device calculates the sign of its input data, for each data frame, in order to maintain 'DC balance' of the transmission line. Based on the data history and the current sign, the G-link determines whether or not the current data frame should be inverted, to ensure a 50 % duty cycle for DC balance. The G-link protocol bits of the C-field indicate this inversion to the receiver chip. In the case of constant input data, e.g. FFFF (hex), the serial output stream will look similar to a 50 ns clock, due to this inversion process. This frequency can easily be observed on an oscilloscope. Individual bits are difficult to see on a 'slow' oscilloscope, because their time duration is only 1.25 ns.

Figure 7.5 shows an oscilloscope picture of both G-link output signals from the Pre-Processor CMC card. Bit-frames, consisting of C-field and D-field, can be

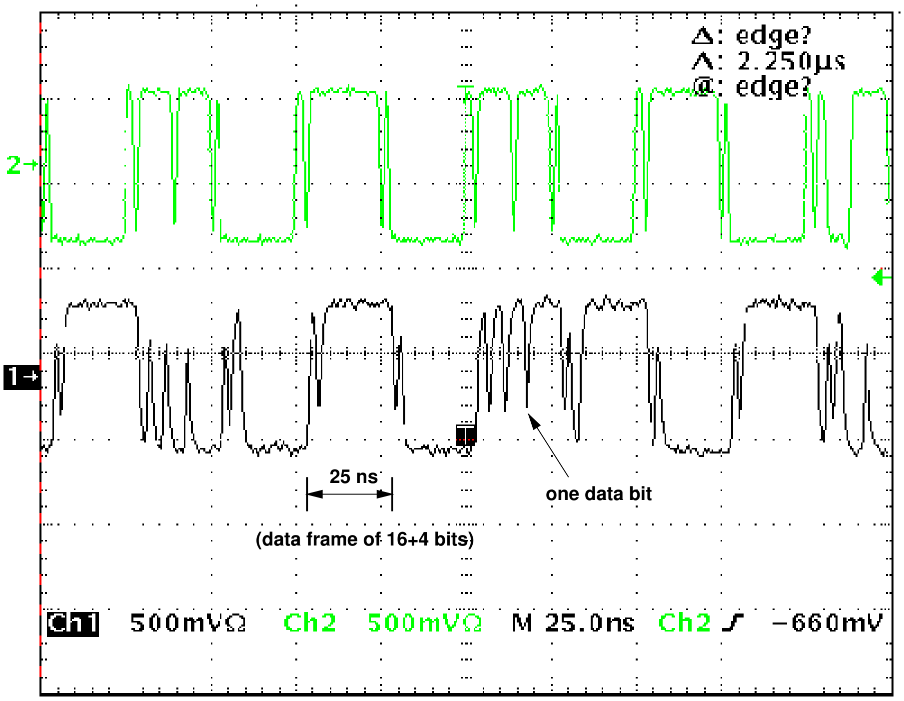

Figure 7.5: G-link output signal waveforms on a 'fast' oscilloscope.

seen. Their time duration is 25 ns. Individual data bits can be observed as well, because of the high analogue bandwidth and sampling frequency of the Tektronix oscilloscope used (1 GHz and 2 GS/s).

It is possible to observe correlations between an analogue input signal and this high-speed serial bit-stream, if the baseline of the input signal is shifted to negative values. Then, only the pulse peak will be digitised and the baseline level is always clipped at zero. The peak-finder will suppress the rest of the signal, except at the peak maximum, which will pass the non-saturated BCID logic. Note that zero FADC values will alternately be inverted by the G-link transmitter to ensure DC balance.

Figure 7.6 shows the correlation of the analogue input signal with the high-speed serial bit-stream. The bunch-crossing-identified data occurs after a latency of 9 bunch-crossings (225 ns) in one G-link bit-stream. This latency attributed as follows: one tick from the FADC, seven ticks from the FeAsic[3], and one tick from the G-link. The Finco ASIC does not contribute to the latency. The positions, where these waveforms are recorded were labelled in Figure 7.4 as '1' and '2' respectively.

The serial bit-stream needs to be received in order to prove that it represents the

[3]The FeAsic design was not optimised for a minimum latency. The FeAsic buffers its input and output data, which requires two additional clock ticks (2×25 ns) compared with the final PPrAsic latency.

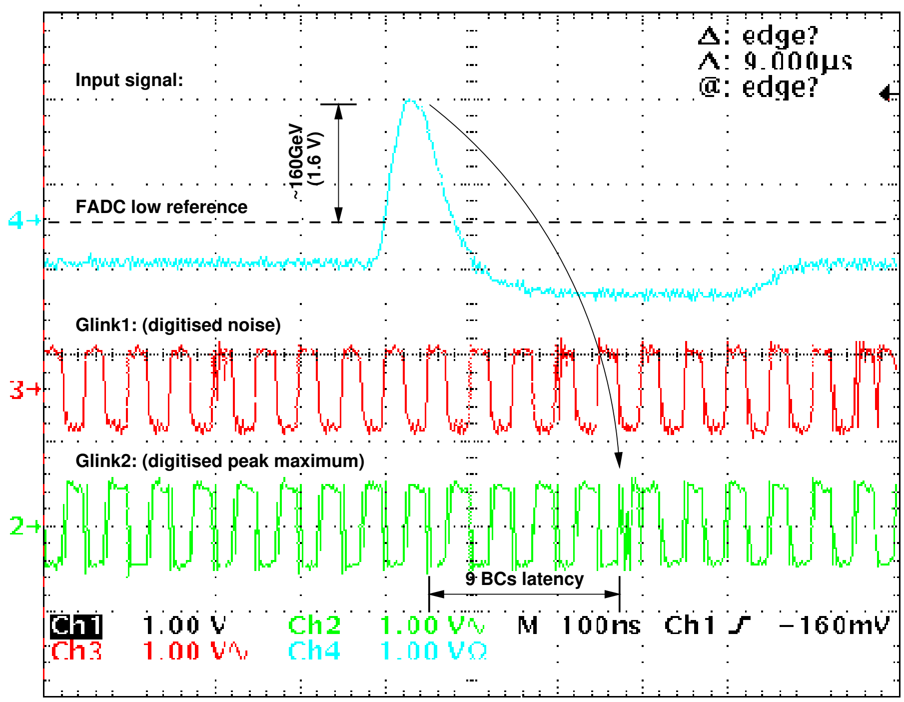

Figure 7.6: The bipolar shaped calorimeter signal is shown in the top waveform. Its bunch-crossing identified correlation in the high-speed serial bit-stream at 800 MBd can be seen in the bottom waveform.

bunch-crossing-identified peak maximum of the analogue input signal. Figure 7.7 shows a histogram of the data, read out by the monitor and control software. The location of this data was marked with '3' in the test set-up Figure 7.4. The histogram shows FADC counts versus time slices of 25 ns. The peak height corresponds to the analogue pulse maximum and the distance between the peaks corresponds to the repetition time, adjusted at the AFG generator.

The MCM system measurements described here have demonstrated the functioning of the PPrD-MCM. For one analogue trigger-tower signal, all the preprocessing has been shown. This includes all the preprocessing from the analogue line receiver circuit up to the reception of high-speed serial data at 800 MBd.

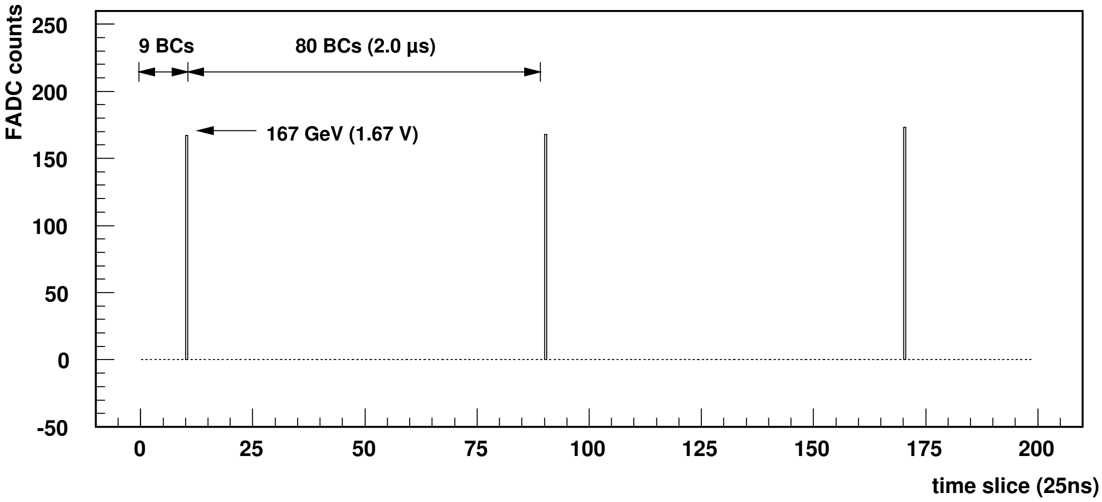

Figure 7.7: Data file content received by the G-link receiver chip, and readout by the monitor and control software.

Chapter 8

Conclusions and outlook

This thesis has described the research whose aim was to develop a compact Pre-Processor system for the ATLAS Level-1 Calorimeter Trigger. The compactness and complexity of the final Pre-Processor is crucial to the architecture of the ATLAS Level-1 Calorimeter Trigger. In addition, its reliability is of importance for the running of the ATLAS experiment, since all data of the Calorimeter Trigger system have to go through the Pre-Processor before any event can be accepted.

Contributions to the performance and the architecture of the Pre-Processor were made. A BCID algorithm for saturated trigger-tower signals was developed, which will be implemented in the Pre-Processor ASIC. This algorithm combines BCID efficency with enormous simplicity in design and implementation. It takes only two samples from the leading pulse edge and compares their values against programmable thresholds. Even if the FADC digitisation strobe is mis-aligned within ± 3 ns, the algorithm will be 100 % efficient up to a peaking times of 63 ns, without re-adjustment of thresholds.

A bunch-crossing multiplexing scheme (BC-mux) which doubles the effective bandwidth of the high-speed serial data transmission from the Pre-Processor to the Cluster Processor were developed. The attraction of the BC-mux scheme is, that it halves the number of transmitter and receiver chips, connectors, and cables for the high number of serial data links to the Cluster Processor.

A demonstrator Multi-Chip Module (PPrD-MCM) was successfully designed and build. It includes most of the final preprocessing and the readout of the Calorimeter Trigger, for four trigger tower signals. The preprocessing includes digitisation to 8-bit precision, identification of the corresponding bunch-crossing in time (BCID), calibration of the transverse energy, readout of raw trigger data, and high-speed serial data transmission to the Calorimeter Trigger processors. The MCM has a size of 4.3×3.7 cm^2 and it consists of 9 dies. The MCM was designed with a smallest feature size of 100 μm and it was fabricated in a laminated MCM-L process offered by Würth Elektronik. It was tested as part of a modular Pre-Processor test system, where transmission and readout tests have shown the feasibility of building a compact Pre-Processor system. Reliability and temperature aspects have been investigated. The clocked MCM temperature is about 42.5 °C, with variations from chip to chip. The un-clocked MCM mean temperature is 34.6 $\pm$1.5 °C.

A Flip-Chip interconnection ASIC (Finco) was developed for the PPrD-MCM and fabricated in a 0.8 μm BiCMOS-process offered by Austria Micro Systems (AMS). This ASIC was designed for analogue baseline adjustment, level-conversion, and to doubled the serial date rate of G-links. It has shown the feasibility of Flip-Chip mounting on the PPrD-MCM and the mixing of analogue with digital components for the final Pre-Processor Multi-Chip Module. It includes temperature monitoring and in-circuit testing (boundary-scan) which gives experience for the final Pre-Processor testability.

The established MCM design technique and experience will now be used for the final Pre-Processor Multi-Chip Module. All Details about the final MCM will be specified and then the MCM size can be optimised to fit 16 PPr-MCMs on a VME board aimed to process 64 trigger tower signals.

Appendix A

Theory of signal transport simulation

The purpose of this section is to provide the theoretical basis of signal integrity analysis used, to model the behaviour of a Multi-Chip Module layer cross-section and its metallisation by a network of resistors, capacitors and inductances. These are the basic elements of an equivalent circuit used to perform reflection and crosstalk simulation. The signal transport theory allows transformation of the continuum set of electromagnetic integral and differential equations into a set of purely algebraic equations which can be solved on a computer using standard numerical techniques.

Simulation method overview

The theory is based on electromagnetic field equations, developed by Maxwell and others. The techniques used for simulation can be subdivided into *differential* and *integral* methods [Sco94]. The *differential* methods deal directly with the field equations in their differential form to compute the electric and magnetic fields for the entire simulation area. These equations are discretised by a grid of nodal elements, to produce a set of linear equations with an appropriate set of boundary conditions. The *integral* method uses the ability to rewrite these differential equations in integral form by the use of Green's theorem [Nol93]. Both methods proceed by subdividing the conductors present in the layout, into a set of elements, and defining a set of basis functions for the charge and current densities in each element. A set of linear equations is then assembled, using the condition that the tangential electric field on perfect conductors is zero, together with a set of voltage and current boundary conditions. These equations are solved to yield the complete charge or current distribution in the layout.

Equivalent-circuit model

An equivalent-circuit model for a Multi-Chip Module and its metallisation is required to associate the real charge and current distribution with a set of inductors, capacitors and resistors. This model is the interface between the electromagnetic theory, used to simulate the electromagnetic behaviour of a Multi-Chip Module layout, and the circuit theory required to test the behaviour of the implemented circuits. If the metallisation can be regarded as thin relative to the dimension of the geometry, a two-dimensional current flow can be assumed. The equivalent-circuit model can be derived from the metallisation polygons by subdividing those into rectangular shapes associated with a node of the circuit. The individual values of the equivalent-circuit components, connected between nodes, are calculated from the charge and current distribution for each of the rectangular shapes. Because of the subdivision of the metallisation into a set of similar shapes, the number of required field solutions is reduced. A calculation is only required for those which are different.

Basic electromagnetic theory as applied to MCMs

The starting points for signal integrity analysis are Maxwell's equations:

$$\begin{aligned} \operatorname{div}\mathbf{D} &= \rho && \text{(A.1)}\\ \operatorname{rot}\mathbf{E} &= -\dot{\mathbf{B}} && \text{(A.2)}\\ \operatorname{rot}\mathbf{H} &= \mathbf{j}+\dot{\mathbf{D}} && \text{(A.3)}\\ \operatorname{div}\mathbf{B} &= 0 && \text{(A.4)} \end{aligned}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic field vectors, $\mathbf{D}$ and $\mathbf{B}$ are the electric and magnetic flux density vectors, ρ is the density of free charge, and $\mathbf{j}$ is the current density.

The field and flux vectors are linked by the following three equations:

$$\begin{aligned} \mathbf{B} &= \mu\mathbf{H} && \text{(A.5)}\\ \mathbf{D} &= \varepsilon\mathbf{E} && \text{(A.6)}\\ \mathbf{j} &= \sigma\mathbf{E} && \text{(A.7)} \end{aligned}$$

where μ is the permeability, ε is the permittivity and σ is the conductivity at a given point. The current-continuity equation, which demonstrates the conservation of charge and which has introduced the current density $\dot{\mathbf{D}}$ in Equation A.3, is:

$$\operatorname{div}\mathbf{j} = -\frac{\partial \rho}{\partial t}. \tag{A.8}$$

For circuit simulations it is more convenient to work in terms of potentials rather than the field vectors, because of the interaction between field and circuit theory.

Imposing the Lorentz Gauge [Nol93], which sets:

$$\operatorname{div}\mathbf{A} + \frac{1}{c^2}\dot{\phi} = 0,$$

one can express the vectors **E** and **B** in terms of an electric scalar potential ϕ and a magnetic vector potential **A**:

$$\mathbf{E} = -\nabla\phi - \dot{\mathbf{A}} \tag{A.9}$$
$$\mathbf{B} = \operatorname{rot}\mathbf{A}. \tag{A.10}$$

Using these equations, the Maxwell equations result in a Lorentz-invariant form for the electromagnetic potentials:

$$\Box\,\mathbf{A} = -\mu\mathbf{j} \tag{A.11}$$
$$\Box\,\phi = \frac{-\rho}{\varepsilon}, \tag{A.12}$$

where $\Box \equiv \Delta - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}$ is defined.

If one assumes that the dominant interaction effects take place at spatial separations of much less than a wavelength, one can ignore the time-dependent terms of Equations A.11 and A.12. This has proved to be a very good approximation for circuit boards and simular structures, which are typically electrically small but geometrically complex [Sco94]. The time dependence and coupling between electric and magnetic fields are still retained in the solution through Equations A.8 and A.9, which allow for propagation effects. This approach can also be adapted to eclectically large problems if there is no long-range coupling. For example, if a large metal area is present, such as a ground plane on one surface of a Multi-Chip Module, rapid attenuation of the potentials is achieved and thus the coupling mechanisms are sufficiently localised for this simulation approach. See Section 5.8.1, which mentions the geometry extraction of a local area without long-range effects using the DF/SigNoise™ [Sig97] tool. The time-independent versions of Equations A.11 and A.12 are:

$$\Delta\,\mathbf{A} = -\mu\mathbf{j} \tag{A.13}$$
$$\Delta\,\phi = \frac{-\rho}{\varepsilon}. \tag{A.14}$$

Each of these equations is then solved to find the potential for a given distribution of charge or current density. Rewriting these equations into their integral representations

$$\mathbf{A}(\mathbf{r}) = \int\int\int G_A(\mathbf{r}|\mathbf{r}')\mathbf{J}(\mathbf{r}')d^3\mathbf{r}' \tag{A.15}$$
$$\phi(\mathbf{r}) = \int\int\int G_\phi(\mathbf{r}|\mathbf{r}')\rho(\mathbf{r}')d^3\mathbf{r}' \tag{A.16}$$

was previously referred to as an *integral* method, where the integrals are formed over a volume which includes all the sources under consideration. The vector $\mathbf{r}$ is the field point vector and $\mathbf{r}'$ is the source point vector. The Green's functions G_ϕ and G_A can be considered to be the potential due to a unit point charge and a unit point current respectively. For the simplest possible case of a source in free space, the solution to this problem is well known as:

$$G_A(\mathbf{r}|\mathbf{r}') = \frac{\mu_0}{4\pi|\mathbf{r}-\mathbf{r}'|}$$
$$G_\phi(\mathbf{r}|\mathbf{r}') = \frac{1}{4\pi\varepsilon_0|\mathbf{r}-\mathbf{r}'|}.$$

In more general situations, with arbitrary arrangements of dielectric material and grounded metal bodies, the Green's functions are rather more complicated. Hence, for a Multi-Chip Module layer cross-section, one can calculate the potentials by either:

- obtaining the Green's functions for a simplified layer structure,
- or by the use of an approximate model for the charge and current distribution using a set of basis functions.

In either case the physical knowledge of the behaviour of charge near conductor edges and above ground planes can reduce the size of the problem. The charge-density distribution of a nth element, which has a set of N_n charge-density basis functions $\Psi_i(\mathbf{r})$ assigned to it, is defined as [Sco94]:

$$\rho(\mathbf{r}) = \sum_{i=s_n}^{s_n+N_n-1} \xi_i \Psi_i(\mathbf{r}).$$

The first basis function is labelled s_n, and the ξ_i is the amplitude of the i^{th} basis function. To determine this solution, the amplitudes need to be sampled over conductors by the use of weighting functions, as described in [Sco94]. The choice of basis functions has a considerable effect on the accuracy and efficiency of the simulation.

Once a set of basis functions is calculated for a given voltage distribution over the layout, a solution for the charge-density distribution is found. This leads to an approximation from which capacitances between elements can be calculated. Equivalent-circuit inductances can be calculated from the actual current-density distribution, with the restriction that in this case the basis functions must be divergence-free, which ensures that there is no charge associated with inductive elements. A confirmation of the validity of the approximations by comparing the results with experimental data can be found in [Sco94]. The final results from that theory are matrices for capacitances (C) and inductances (L), as well as for resistance (R) and conductance (G) required for a equivalent-circuit model. Resistance and conductance matrices are required to take conductor loss and dielectric loss into account. These matrices are referred to as an RLGC matrix combination, which is the basis of the signal integrity analysis performed using the DF/SigNoise^{TM} simulation tool.

Reflection and Crosstalk

The simulation method described above includes reflection and crosstalk effects to a primary net which is under simulation. The simulation results are transient voltage curves from which the influence of reflection and crosstalk can directly be seen, in order to give a qualitative impression on how a signal would look. Since the analysis performed is a simulation of voltages rather than a calculation of reflection and crosstalk parameters, electrical properties from the voltage curves can give additional information. Hence, a short introduction to reflection and crosstalk is given here, as required for the analysis of the following simulation results.

The Lorentz-invariant expression for the electromagnetic potentials are given in Equations A.11 and A.12. The general task for those equations is to solve an inhomogeneous differential equation of the form:

$$\Box \, \psi(\mathbf{r}, t) = -\sigma(\mathbf{r}, t),$$

where $\sigma(\mathbf{r}, t)$ represents a general source function and $\psi(\mathbf{r}, t)$ a general waveform. A solution for such an equation can be obtained by the use of a retarded Green's-functions. Here the electromagnetic potentials get a similar structure as for the electrostatic or magnetostatic potentials [Nol93].

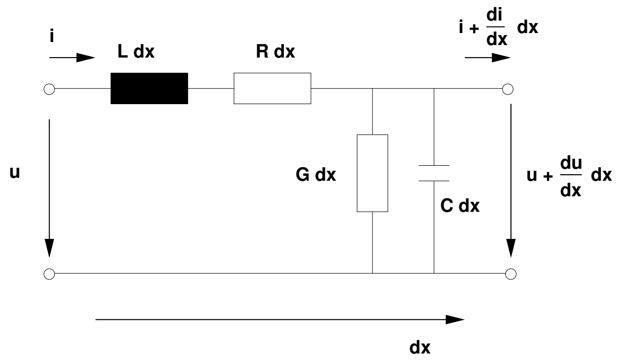

Figure A.1: Equivalent circuit model of an infinitesimal track.

A more practical solution is to start from an equivalent circuit for infinitesimal tracks of length dx. The circuit is shown in Figure A.1. By the use of Kirchhoff's-Law, one can calculate the voltage drop over R and L by:

$$-\frac{\mathrm{d}u}{\mathrm{d}x} = R \cdot i + L \cdot \frac{\mathrm{d}i}{\mathrm{d}t}$$

and for G and C by:

$$-\frac{\mathrm{di}}{\mathrm{dx}} = G \cdot u + C \cdot \frac{\mathrm{du}}{\mathrm{dt}}.$$

These differential equations can be transformed into:

$$-\frac{\mathrm{d}U}{\mathrm{dx}} = (R + i\omega L) \cdot I \tag{A.17}$$

$$-\frac{\mathrm{d}I}{\mathrm{dx}} = (G + i\omega C) \cdot U \tag{A.18}$$

by using a complex voltage expression $U(x)$ and a complex current expression $I(x)$. They are defined by: $u(x, t) = Re\{U(x)e^{i\omega t}\}$ and $i(x, t) = Re\{I(x)e^{i\omega t}\}$.

After differentiation of Equation A.17 and combining it with Equation A.18, the result is a wave propagation equation of the form:

$$\frac{\mathrm{d}^2 U}{\mathrm{dx}^2} = \gamma^2 U,$$

where $\gamma = \sqrt{(R + i\omega L)(G + i\omega C)} = \alpha + i\beta$ is defined. The exponent α can be considered as a damping coefficient and β as a phase. A general solution for that equation using complex amplitudes U_0^f and U_0^b is:

$$\begin{aligned} U(x) &= U_0^f e^{-\gamma x} + U_0^b e^{\gamma x} \\ &= U_f(x) + U_b(x), \end{aligned}$$

and for the current it is:

$$\begin{aligned} I(x) &= \frac{U_0^f}{Z_0} e^{-\gamma x} - \frac{U_0^b}{Z_0} e^{\gamma x} \\ &= I_f(x) + I_b(x), \end{aligned}$$

The index f refers to wave propagation in the forward direction, whereas b refers to wave propagation in backward direction. A track termination (Z_{load}) at a distance $x = l$ will cause reflections on the track if Z_{load} does not match the impedance of the track. The fraction of the backward and forward wave intensity is expressed by the reflectivity r. It can be calculated by taking constraints present at $x = l$, which are $U(l) = U_{load}$, $I(l) = I_{load}$, and $U_{load} = Z_o \cdot I_{load}$:

$$r = \frac{U_b(l)}{U_f(l)} = \frac{Z_{load} - Z_0}{Z_{load} + Z_0}. \tag{A.19}$$

where $Z_0 = \sqrt{\frac{R+i\omega L}{G+i\omega C}}$ is the track impedance.

Therefore the reflectivity depends only on the impedance of the track (Z_0) and on the impedance of the track termination (Z_{load}). In general, the reflectivity is a complex number changing the amplitude and phase of the reflection wave. For tracks which have fewer ohmic losses ($G \ll i\omega$, $R \ll i\omega L$), the impedance Z_o can be considered as purely ohmic:

$$Z_0 = \sqrt{\frac{L}{C}} \tag{A.20}$$

The damping factor α can then be calculated as:

$$\alpha = \frac{R}{2} \cdot \sqrt{\frac{C}{L}}.$$

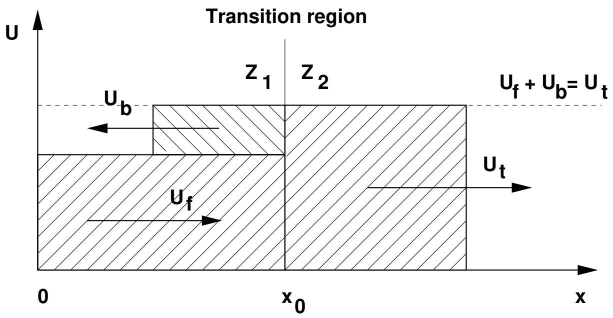

Figure A.2: Illustration of the transition region at which refraction takes place.

Transmittivity

In case of an impedance mismatch, introduced for example by a via, wave reflection and transmission take place. Such a case is illustrated in Figure A.2. The transmittivity t is:

$$t = \frac{2Z_2}{Z_1 + Z_2} \tag{A.21}$$

where Z_1 and Z_2 are the impedances of each side. In the case of a track termination, Z_2 is equal to Z_0 and Z_1 to Z_{load}. This leads to the relation $t = 1 + r$.

Appendix B

PPrD-MCM specification

Pin	**Pad Name**	**Description**	**Signal**
1	AIN1	analogue input FADC1 channel 1	analogue in
2	AIN2	analogue input FADC2 channel 2	analogue in
3	VDDA_FADC1	analogue power supply FADC 1	+5V
4	VDD_FADC1	digital power supply FADC 1	+5V
5	VEE_FADC1	analogue power supply FADC 1	–5V
6	VEE_FADC1	analogue power supply FADC 1	–5V
7	ENC1	clock FADC 1	40 MHz TTL in
8	VINT1	internal reference voltage FADC 1	analogue out
9	SERIALACK1	serial acknowledge FeAsic 1	TTL out
10	GNDA_FADC1	analogue ground FADC 1	ground
11	SERIALCLK1	serial clock FeAsic 1	TTL in
12	VDD_FEASIC1	digital power supply FeAsic 1	+5V
13	SERIALINP1	serial input FeAsic 1	TTL in
14	SERIALRDY1	serial ready FeAsic 1	TTL out
15	SERIALOUT1	serial output FeAsic 1	TTL out
16	SERIALCLK2	serial clock FeAsic 2	TTL in
17	SERIALACK2	serial acknowledge FeAsic 2	TTL out
18	SERIALINP2	serial input FeAsic 2	TTL in
19	SERIALRDY2	serial ready FeAsic 2	TTL out
20	SERIALOUT2	serial output FeAsic 2	TTL out
21	LISTANFLAG	listen flag all FeAsics	TTL in
22	LVL1ACCEPT	level-1 accept all FeAsics	TTL in
23	VCC_FINCO	power Finco TTL-PECL converter	+5V
24	VCC_FINCO	power Finco TTL-PECL converter	+5V
25	TMS	test mode select J-TAG interface	TTL in
26	VDD_FINCO	digital power Finco TTL pads	+5V
27	GNDA_FINCO	analogue ground Finco	ground
28	VDDA_FINCO	analogue power Finco	+5V
29	TRSTB	test mode reset Finco J-TAG interface	TTL in
30	ANALOGINP	positive analogue input Finco	analogue in

Table B.1: PPrD-MCM pin definition for connector J9.

Pin	Pad Name	Description	Signal
31	ANALOGINN	negative analogue input Finco	analogue in
32	DACRP	positive ref. voltage Finco DACs	analogue in (0-5V)
33	DACRN	negative ref. voltage Finco DACs	analogue in (0-5V)
34	DAC3	analogue output Finco DAC 3	analogue out
35	DAC2	analogue output Finco DAC 2	analogue out
36	DAC1	analogue output Finco DAC 1	analogue out
37	ANALOGOUT	analogue output Finco	analogue out
38	DAC0	analogue output Finco DAC 0	analogue out
39	STREGRESET	state register reset all FeAsics	TTL in
40	SYNCRESET	synchron reset all FeAsics	TTL in
41	SERIALRDY4	serial ready FeAsic 4	TTL out
42	SERIALOUT4	serial output FeAsic 4	TTL out
43	SERIALACK4	serial acknowledge FeAsic 4	TTL out
44	SERIALINP4	serial input FeAsic 4	TTL in
45	SERIALOUT3	serial output FeAsic 3	TTL out
46	SERIALCLK4	serial clock FeAsic 4	TTL in
47	SERIALINP3	serial input FeAsic 3	TTL in
48	SERIALRDY3	serial ready FeAsic 3	TTL out
49	SERIALCLK3	serial clock FeAsic 3	TTL in
50	VDD_FEASIC2	digital power supply FeAsic 2	+5V
51	SERIALACK3	serial acknowledge FeAsic 3	TTL out
52	GNDA_FADC2	analogue ground FADC 2	ground
53	ENC2	clock FADC 2	40 MHz TTL in
54	VINT2	internal ref. voltage FADC 2	analogue out
55	VEE_FADC2	analogue power supply FADC 2	–5V
56	VEE_FADC2	analogue power supply FADC 2	–5V
57	VDDA_FADC2	analogue power supply FADC 2	+5V
58	VDD_FADC2	digital power supply FADC 2	+5V
59	AIN3	analogue input FADC2 channel 1	analogue in
60	AIN4	analogue input FADC2 channel 2	analogue in

Table B.2: PPrD-MCM pin definition for connector J9 (continued).

Pin	**Pad Name**	**Description**	**Signal**
1	GND	ground plane	ground
2	VDD_GLINK2_GND	digital supply G-link 2	+5V
3	GND	ground plane	ground
4	VDD_GLINK2_HGND	digital supply high-speed G-link 2	+5V
5	GND	ground plane	ground
6	VDD_GLINK2_ECLGND	digital supply ECL pads G-link 2	+5V
7	-DOUT2	neg. high-speed data output G-link2	1600 MBd PECL out
8	DOUT2	pos. high-speed data output G-link2	1600 MBd PECL out
9	-STRBIN2	neg. clock G-link 2	40 MHz PECL in
10	STRBIN2	pos. clock G-link 2	40 MHz PECL in
11	STRBOUT2	Oscilloscope frame trigger G-link2	PECL out
12	LOCKED2	loop in-lock indication G-link 2	PECL out
13	TEMP2	test pin G-link 2	analogue out
14	SELIN	input select serial interface Finco	TTL in
15	TOSPYBUS0	SpyBus output bit 0	TTL out
16	TOSPYBUS1	SpyBus output bit 1	TTL out
17	TOSPYBUS2	SpyBus output bit 2	TTL out
18	TOSPYBUS3	SpyBus output bit 3	TTL out
19	TOSPYBUS4	SpyBus output bit 4	TTL out
20	TOSPYBUS5	SpyBus output bit 5	TTL out
21	TOSPYBUS6	SpyBus output bit 6	TTL out
22	TOSPYBUS7	SpyBus output bit 7	TTL out
23	CLRIN	input data serial interface Finco	TTL in
24	TCK	Test Access Port clock	TTL in
25	TDI	Test Data Input J-TAG interface	TTL in
26	TDO	Test Data Output J-TAG interface	TTL in
27	CLKIN	serial interface clock Finco	TTL in
28	DPD	active high power down Finco	TTL in
29	CLRDATA	reset serial interface Finco	TTL in
30	-RST	chip reset all G-links	PECL in

Table B.3: PPrD-MCM pin definition for connector J10 (continued).

Pin	Pad Name	Description	Signal
31	GND	ground plane	ground
32	-DAV	data available all G-links	PECL in
33	TEMP	current output themperature sensor	analogue out
34	GND	ground plane	ground
35	GND	ground plane	ground
36	SPYBUS_CTRL2	SpyBus control bit 2 all FeAsics	TTL in
37	MDFSEL	select double frame mode all G-links	PECL in
38	SPYBUS_CTRL1	SpyBus control bit 1 all FeAsics	TTL in
39	M20SEL	select 16 or 20 bit mode all G-links	PECL in
40	SPYBUS_CTRL0	SpyBus control bit 0 all FeAsics	TTL in
41	DIV0	VCO divider select	PECL in
42	GND	ground plane	ground
43	GND	ground plane	ground
44	BCCLK	Finco clock	40 MHz TTL in
45	FEASICCLK	clock for all FeAsics	TTL in
46	GND	ground plane	ground
47	GND	ground plane	ground
48	TEMP1	test pin G-link 1	analogue out
49	LOCKED1	loop in-lock indication G-link 1	PECL out
50	STRBOUT1	Oscilloscope frame trigger G-link1	PECL out
51	STRBIN1	pos. clock G-link 1	40 MHz PECL in
52	-STRBIN1	neg. clock G-link 1	40 MHz PECL in
53	DOUT1	pos. high-speed data output G-link1	1600 GBd PECL out
54	-DOUT1	neg. high-speed data output G-link1	1600 GBd PECL out
55	GND	ground plane	ground
56	VDD_GLINK1_ECLGND	digital supply ECL pads G-link 1	+5V
57	GND	ground plane	ground
58	VDD_GLINK1_HGND	digital supply high speed G-link 1	+5V
59	GND	ground plane	ground
60	VDD_GLINK1_GND	digital supply G-link 1	+5V

Table B.4: PPrD-MCM pin definition for connector J10 (continued).

Appendix C

Finco ASIC specification

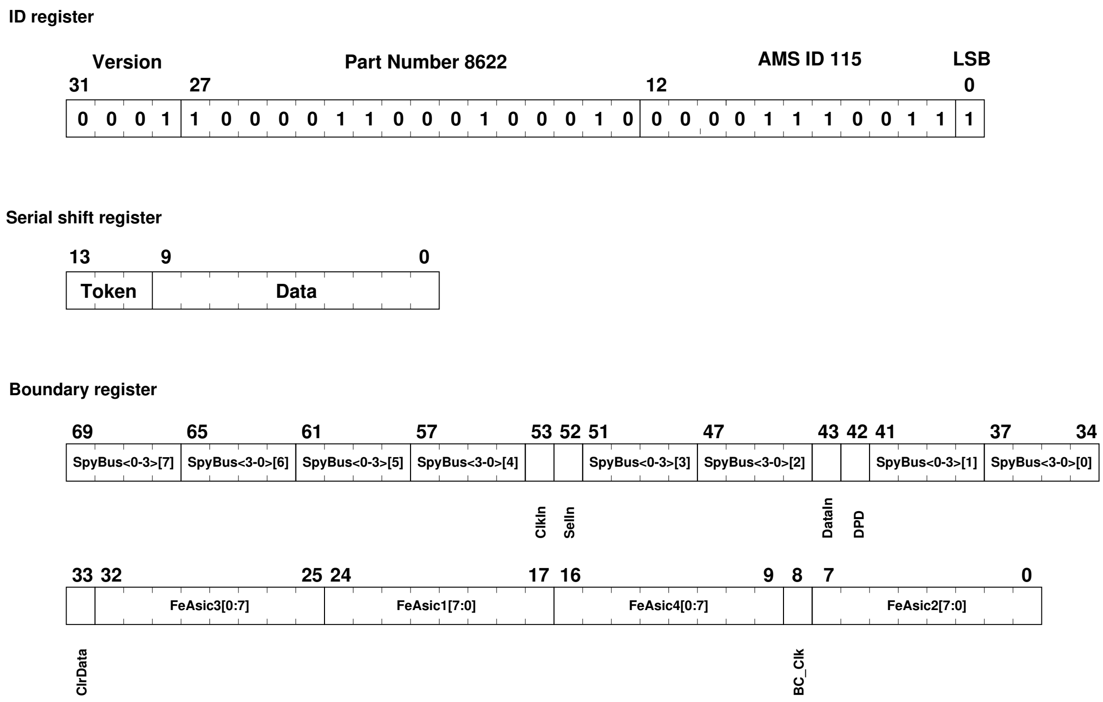

Figure C.1: Finco ASIC register definitions.

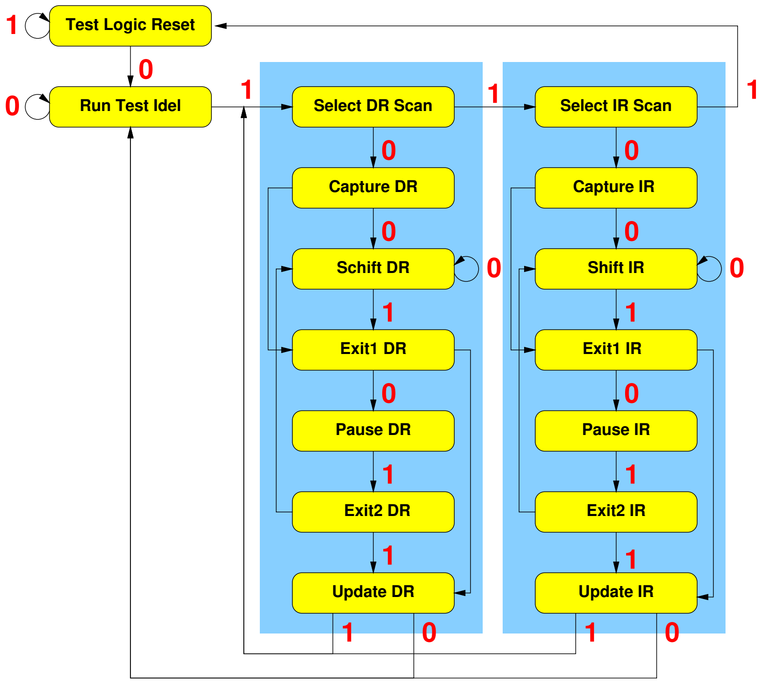

TMS **Test Mode Select**
TDI **Test Data Input**
TDO **Test Data Output**
TDRST **Test Data Reset**
TCK **Test Clock**

Figure C.2: State machine of the Test Access Port controler (TAP).

Pin	JEM	X Y [μm]	Pad Name	Description	Signal
A1	114	0 0	VCC3	power rail PECL pads	+5V
A2	111	450 0	VEE2	ground rail PECL pads	ground
A3	109	900 0	Glink2[0]	G-link 2 bit 0	PECL out
A4	106	1350 0	Glink2[1]	G-link 2 bit 1	PECL out
A5	102	1800 0	Glink2[2]	G-link 2 bit 2	PECL out
A6	98	2250 0	Glink2[3]	G-link 2 bit 3	PECL out
A7	94	2700 0	Glink2[4]	G-link 2 bit 4	PECL out
A8	90	3150 0	Glink2[5]	G-link 2 bit 5	PECL out
A9	88	3600 0	Glink2[6]	G-link 2 bit 6	PECL out
A10	84	4050 0	Glink2[7]	G-link 2 bit 7	PECL out
A11	80	4500 0	SpyBus[0]_pad[0]	SpyBus 0 bit 0	TTL in
A12	76	4950 0	SpyBus[1]_pad[0]	SpyBus 1 bit 0	TTL in
A13	72	5400 0	SpyBus[2]_pad[0]	SpyBus 2 bit 0	TTL in
A14	69	5850 0	SpyBus[3]_pad[0]	SpyBus 3 bit 0	TTL in
A15	66	6300 0	ClrData_pad	serial interface clear data	TTL in
A16	63	6750 0	ToSpyBus_pad[0]	SpyBus output bit 0	TTL out
B1	117	0 650	VDD	power rail TTL pads	+5V
B2	116	450 650	GND	ground rail TTL pads	ground
B3	110	900 650	FeAsic4[0]	FeAsic 4 bit 0	TTL in
B4	108	1350 650	FeAsic4[1]	FeAsic 4 bit 1	TTL in
B5	104	1800 650	FeAsic4[2]	FeAsic 4 bit 2	TTL in
B6	100	2250 650	FeAsic4[3]	FeAsic 4 bit 3	TTL in
B7	95	2700 650	FeAsic4[4]	FeAsic 4 bit 4	TTL in
B8	91	3150 650	FeAsic4[5]	FeAsic 4 bit 5	TTL in
B9	86	3600 650	FeAsic4[6]	FeAsic 4 bit 6	TTL in
B10	82	4050 650	FeAsic4[7]	FeAsic 4 bit 7	TTL in
B11	78	4500 650	SpyBus[0]_pad[1]	SpyBus 0 bit 1	TTL in
B12	74	4950 650	SpyBus[1]_pad[1]	SpyBus 1 bit 1	TTL in
B13	70	5400 650	SpyBus[2]_pad[1]	SpyBus 2 bit 1	TTL in
B14	67	5850 650	SpyBus[3]_pad[1]	SpyBus 3 bit 1	TTL in
B15	62	6300 650	DPD_pad	active high power down	TTL in
B16	61	6750 650	ToSpyBus_pad[1]	SpyBus output bit 1	TTL out
C1	120	0 1300	VCC2	power rail PECL pads	+5V
C2	119	450 1300	VSUB	PECL substrate contact	ground
C3	118	900 1300	Glink1[0]	G-link 1 bit 0	PECL out
C4	112	1350 1300	Glink1[1]	G-link 1 bit 1	PECL out
C5	107	1800 1300	Glink1[2]	G-link 1 bit 2	PECL out
C6	103	2250 1300	Glink1[3]	G-link 1 bit 3	PECL out
C7	97	2700 1300	Glink1[4]	G-link 1 bit 4	PECL out
C8	92	3150 1300	Glink1[5]	G-link 1 bit 5	PECL out

Table C.1: Finco ASIC pad definition.

Pin	JEM	X Y [μm]	Pad Name	Description	Signal
C9	89	3600 1300	Glink1[6]	G-link 1 bit 6	PECL out
C10	81	4050 1300	Glink1[7]	G-link 1 bit 7	PECL out
C11	75	4500 1300	SpyBus[0]_pad[2]	SpyBus 0 bit 2	TTL in
C12	71	4950 1300	SpyBus[1]_pad[2]	SpyBus 1 bit 2	TTL in
C13	65	5400 1300	SpyBus[2]_pad[2]	SpyBus 2 bit 2	TTL in
C14	60	5850 1300	SpyBus[3]_pad[2]	SpyBus 3 bit 2	TTL in
C15	59	6300 1300	Data_pad	serial interface input data	TTL in
C16	58	6750 1300	ToSpyBus_pad[2]	SpyBus output bit 2	TTL out
D1	127	0 1950	THEMP_OUT_pad	temp. sensor output pad	analogue out
D2	126	450 1950	BC_CLK_pad	bunch-crossing clock	TTL in
D3	125	900 1950	FeAsic2[0]	FeAsic 2 bit 0	TTL in
D4	124	1350 1950	FeAsic2[1]	FeAsic 2 bit 1	TTL in
D5	123	1800 1950	FeAsic2[2]	FeAsic 2 bit 2	TTL in
D6	105	2250 1950	FeAsic2[3]	FeAsic 2 bit 3	TTL in
D7	99	2700 1950	FeAsic2[4]	FeAsic 2 bit 4	TTL in
D8	93	3150 1950	FeAsic2[5]	FeAsic 2 bit 5	TTL in
D9	85	3600 1950	FeAsic2[6]	FeAsic 2 bit 6	TTL in
D10	79	4050 1950	FeAsic2[7]	FeAsic 2 bit 7	TTL in
D11	73	4500 1950	SpyBus[0]_pad[3]	SpyBus 0 bit 3	TTL in
D12	68	4950 1950	SpyBus[1]_pad[3]	SpyBus 1 bit 3	TTL in
D13	54	5400 1950	SpyBus[2]_pad[3]	SpyBus 2 bit 3	TTL in
D14	53	5850 1950	SpyBus[3]_pad[3]	SpyBus 3 bit 3	TTL in
D15	52	6300 1950	Sel_pad	serial interface select	TTL in
D16	51	6750 1950	ToSpyBus_pad[3]	SpyBus output bit 3	TTL out
E1	133	0 2600	VCC1	power rail PECL pads	+5V
E2	130	450 2600	VEE1	ground rail PECL pads	ground
E3	136	900 2600	Glink1[8]	G-link 1 bit 8	PECL out
E4	129	1350 2600	Glink1[9]	G-link 1 bit 9	PECL out
E5	137	1800 2600	Glink1[10]	G-link 1 bit 10	PECL out
E6	128	2250 2600	Glink1[11]	G-link 1 bit 11	PECL out
E7	101	2700 2600	Glink1[12]	G-link 1 bit 12	PECL out
E8	96	3150 2600	Glink1[13]	G-link 1 bit 13	PECL out
E9	83	3600 2600	Glink1[14]	G-link 1 bit 14	PECL out
E10	77	4050 2600	Glink1[15]	G-link 1 bit 15	PECL out
E11	50	4500 2600	SpyBus[0]_pad[4]	SpyBus 0 bit 4	TTL in
E12	41	4950 2600	SpyBus[1]_pad[4]	SpyBus 1 bit 4	TTL in
E13	49	5400 2600	SpyBus[2]_pad[4]	SpyBus 2 bit 4	TTL in
E14	42	5850 2600	SpyBus[3]_pad[4]	SpyBus 3 bit 4	TTL in
E15	48	6300 2600	Clk_pad	serial interface clock	TTL in
E16	45	6750 2600	ToSpyBus_pad[4]	SpyBus output bit 4	TTL out

Table C.2: Finco ASIC pad definition (continued).

Pin	JEM	X Y [um]	Pad Name	Description	Signal
F1	138	0 3250	VDD	power rail TTL pads	+5V
F2	139	450 3250	GND	ground rail TTL pads	ground
F3	140	900 3250	FeAsic1[0]	FeAsic 1 bit 0	TTL in
F4	141	1350 3250	FeAsic1[1]	FeAsic 1 bit 1	TTL in
F5	144	1800 3250	FeAsic1[2]	FeAsic 1 bit 2	TTL in
F6	165	2250 3250	FeAsic1[3]	FeAsic 1 bit 3	TTL in
F7	169	2700 3250	FeAsic1[4]	FeAsic 1 bit 4	TTL in
F8	174	3150 3250	FeAsic1[5]	FeAsic 1 bit 5	TTL in
F9	11	3600 3250	FeAsic1[6]	FeAsic 1 bit 6	TTL in
F10	17	4050 3250	FeAsic1[7]	FeAsic 1 bit 7	TTL in
F11	20	4500 3250	SpyBus[0]_pad[5]	SpyBus 0 bit 5	TTL in
F12	31	4950 3250	SpyBus[1]_pad[5]	SpyBus 1 bit 5	TTL in
F13	37	5400 3250	SpyBus[2]_pad[5]	SpyBus 2 bit 5	TTL in
F14	38	5850 3250	SpyBus[3]_pad[5]	SpyBus 3 bit 5	TTL in
F15	39	6300 3250	TDI_pad	test data input pad	TTL in
F16	40	6750 3250	ToSpyBus_pad[5]	SpyBus output bit 5	TTL out
G1	143	0 3900	VCC3	power rail PECL pads	+5V
G2	146	450 3900	VEE2	ground rail PECL pads	ground
G3	147	900 3900	Glink2[8]	G-link 2 bit 8	PECL out
G4	152	1350 3900	Glink2[9]	G-link 2 bit 9	PECL out
G5	162	1800 3900	Glink2[10]	G-link 2 bit 10	PECL out
G6	166	2250 3900	Glink2[11]	G-link 2 bit 11	PECL out
G7	170	2700 3900	Glink2[12]	G-link 2 bit 12	PECL out
G8	175	3150 3900	Glink2[13]	G-link 2 bit 13	PECL out
G9	7	3600 3900	Glink2[14]	G-link 2 bit 14	PECL out
G10	13	4050 3900	Glink2[15]	G-link 2 bit 15	PECL out
G11	18	4500 3900	SpyBus[0]_pad[6]	SpyBus 0 bit 6	TTL in
G12	21	4950 3900	SpyBus[1]_pad[6]	SpyBus 1 bit 6	TTL in
G13	26	5400 3900	SpyBus[2]_pad[6]	SpyBus 2 bit 6	TTL in
G14	30	5850 3900	SpyBus[3]_pad[6]	SpyBus 3 bit 6	TTL in
G15	32	6300 3900	TCK_pad	serial interface clock	TTL in
G16	33	6750 3900	ToSpyBus_pad[6]	SpyBus output bit 6	TTL out
H1	150	0 4550	VDD	power core cells	+5V
H2	151	450 4550	GND	ground core cells	ground
H3	154	900 4550	FeAsic3[0]	FeAsic 3 bit 0	TTL in
H4	159	1350 4550	FeAsic3[1]	FeAsic 3 bit 1	TTL in
H5	163	1800 4550	FeAsic3[2]	FeAsic 3 bit 2	TTL in
H6	167	2250 4550	FeAsic3[3]	FeAsic 3 bit 3	TTL in
H7	171	2700 4550	FeAsic3[4]	FeAsic 3 bit 4	TTL in
H8	176	3150 4550	FeAsic3[5]	FeAsic 3 bit 5	TTL in

Table C.3: Finco ASIC pad definition (continued).

Pin	JEM	X Y [um]	Pad Name	Description	Signal
H9	5	3600 4550	FeAsic3[6]	FeAsic 3 bit 6	TTL in
H10	9	4050 4550	FeAsic3[7]	FeAsic 3 bit 7	TTL in
H11	14	4500 4550	SpyBus[0]_pad[7]	SpyBus 0 bit 7	TTL in
H12	19	4950 4550	SpyBus[1]_pad[7]	SpyBus 1 bit 7	TTL in
H13	23	5400 4550	SpyBus[2]_pad[7]	SpyBus 2 bit 7	TTL in
H14	25	5850 4550	SpyBus[3]_pad[7]	SpyBus 3 bit 7	TTL in
H15	27	6300 4550	TDO_pad	test data output pad	TTL out
H16	28	6750 4550	ToSpyBus_pad[7]	SpyBus output bit 7	TTL out
I1	153	0 5200	GND	ground rail TTL pads	ground
I2	155	450 5200	VDD	power rail TTL pads	+5V
I3	156	900 5200	TRSTB_pad	test access port reset	TTL in
I4	160	1350 5200	TMS_pad	test mode select	TTL in
I5	164	1800 5200	VDDA	power analog cells	+5V
I6	168	2250 5200	GNDA	ground analog cells	ground
I7	172	2700 5200	DAC_RP_pad	positive DAC ref. voltage	analog in
I8	1	3150 5200	DAC_RN_pad	negative DAC ref. voltage	analog in
I9	2	3600 5200	DAC[0]_OUT_pad	DAC 0 outputs	analog out
I10	8	4050 5200	DAC[1]_OUT_pad	DAC 1 outputs	analog out
I11	10	4500 5200	DAC[2]_OUT_pad	DAC 2 outputs	analog out
I12	15	4950 5200	DAC[3]_OUT_pad	DAC 3 outputs	analog out
J7	173	2700 5850	AnalogInP_pad	positive differential input	analog in
J10	6	4050 5850	AnalogInN_pad	negative differential input	analog in
J12	12	4950 5850	AnalogOut_pad	analog output signal	analog out

Table C.4: Finco ASIC pad definition (continued).

Token	Binary value	Hex value
LoadDac0	0b000	0x0
LoadDac1	0b001	0x1
LoadDac2	0b010	0x2
LoadDac3	0b011	0x3
LoadDacTest	0b100	0x4
LoadMux	0b101	0x5
LoadFrame	0b110	0x6

Table C.5: Serial interface tokens.

Token	**Binary value**	**Hex value**	**Comment**
EXTEST	0b0000	0x0	external test
INTEST	0b0001	0x1	internal test
RUNBIST1	0b0010	0x2	built-in-self test
HIGHZ	0b0011	0x3	high impedence
unused	0b0100	0x4	—
ATPGTEST2	0b0101	0x5	int. Test (not used)
ATPGTEST3	0b0110	0x6	int. Test (not used)
ATPGTEST4	0b0111	0x7	int. Test (not used)
SAMPLE	0b1000	0x8	preload sample
IDDCLAMP	0b1001	0x9	IDD test
EXTEST2	0b1010	0xA	external test 2 (not used)
INTEST2	0b1011	0xB	internal test 2 (not used)
ATPGTEST1	0b1100	0xC	int. Test (not used)
IDCODE	0b1101	0xD	id register
SAMPLE2	0b1110	0xE	preload sample2
BYPASS	0b1111	0xF	bypass mode

Table C.6: Test Access Port (TAP) tokens.

Dimension x, y [μm]	**Pad count**	**Pad size**	**Orientation pad A1**
6980, 6175	143	85 μm	lower left (ASIC logo)

Table C.7: Finco ASIC general information.

Index

Bibliography

[AD] AD9058
Dual 8-Bit 50 MSPS A/D Converter
Analog Devices, Technical data sheed, Rev. A
`http://www.analog.com`

[AMS] AMS *Austria Mikro Systeme International AG*
`http://www.vertical-global.com`

[AMS95-1] Austria Micro Systems
0.8 μm CMOS Design Rules and Process Parameters
Austria Micro Systems, 1995

[AMS95-2] Austria Micro Systems
0.8 μm BiCMOS Design Rules and Process Parameters
Austria Micro Systems, 1995

[AMS96] Austria Micro Systems
AMS Hit-Kit 3.01 documentation
Austria Micro Systems, 10. December 1996

[Apd97] APD
Advanced Package Designer (APD) — Getting Started Guide
Cadence Product Version 2.0, Openbook documentation, February 1997
`http://www.cadence.com`

[Bes91] Best, R.
Digitale Messwertverarbeitung
R. Oldenbourg Verlag GmbH, ISBN 3-486-21573-6, Germany 1991

[Bra94] Brawn, I. P. et al.
Beam Test of a Single-Channel Bunch-Crossing Identification Module for the Level-1 Trigger
RD-27 note 31, October 1994
`http://www1.cern.ch/RD27/note31.ps`

[Bra96] Brawn, I. P.
Bunch Crossing Identification for the ATLAS Level-1 Calorimeter Trigger
Thesis, Faculty of Science University of Birmingham, England June 1996

[Cad] Cadence
Cadence Design Systems, Inc.
http://www.cadence.com

[Cha95] Chase, R. L. et al.
A fast monolithic shaper for ATLAS e.m. calorimeter
ATLAS internal Note, LARG-NO-10, 1 March 1995
http://preprints.cern.ch/archive/electronic/cern/others/atlnot/Note/larg/larg-010.ps.gz

[Cle98] Cleland, B.
PSPICE model of the ATLAS LAr. analogue trigger tower chain
private communication, University of Pittsburgh, July 1998

[CMC95] CMC
Draft Standard for a Common Mezzanine Card Family: CMC
IEEE P1386/Draft 2.0, April 1995

[Col98] Collot, J. et al.
*The LAr Tri-Gain Shape*r
ATLAS Internal Note, LArG-No-92, 11 March 1998
http://preprints.cern.ch/archive/electronic/cern/others/atlnot/Note/larg/larg-092.ps.gz

[Con96] Concept
Concept Schematic — User Guide
Cadence Product Version 3.0, Openbook documentation, February 1996
http://www.cadence.com

[Dav77] Davis, R. F.
Computerized Thermal Analysis
IEEE Trans. Parts, Hybrids, Packaging, vol. PHP-13 Sept. 1977

[Dil88] Dillinger, T. E.
VLSI Engineering
Prentice-Hall Inc., ISBN 0-13-942731-7, United States 1988

[Ell97] Ellis, J.
Physics at future colliders
Plenary session talk presented at Europhysics Conference on High Energy Physics, CERN-TH/97-367, Jerusalem, August 1997

[Emc97] DF/EMControl
DF/EMControl — User Giude
Cadence Product Version 12.0, Openbook documentation, February 1997
http://www.cadence.com

[Fea96] Front-End ASIC (FeAsic)
Front-End ASIC — User's Guide

Internal document, HD-ASIC-19-0896, Universität Heidelberg, Germany August 1996
`http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html`

[FEM97] Front End Module (FEM)
Front End Module — User's Guide
Internal document, HD-ASIC-36-1097, Universität Heidelberg, Germany October 1997
`http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html`

[Gar96] Garvey, J. et al.
Bunch Crossing Identification for the ATLAS Level-1 Calorimeter Trigger
ATLAS internal note, DAQ-NO-051, The UK Level-1 Trigger Group, England 20 May 1996

[Gre92] Greiner, W.
Spezielle Relativitätstheorie
Verlag Harri Deutsch, ISBN 3-8171-1205-X, Frankfurt am Main 1992

[Han95] Hanke, P. at al.
A flash analog-to-digital converter for trigger tests with the ATLAS calorimetry
Research and development note, RD27, Heidelberg 1995

[Har97] Harper, C. A.
Electronic Packageing & Interconnection Handbook - second edition
HcGraw-Hill Inc., ISBN 0-07-026694-8, United States 1997

[Hoe99] Hötzel, W.
Untersuchungen zur Zeitstruktur von Kalorimeterpulsen im ATLAS-Experiment am LHC
Diploma Thesis, IHEP-99-05, Universität Heidelberg, Germany 1999

[HP] HDMP-1012/14
Low Cost Gigabit Rate Transmitter/Receiver Chip set
Hewlett Packard, Technical data sheed
`http://www.hp.com`

[IBIS] IBIS
IBIS I/O Buffer Information
Specification ANSI/EIA-656
`http://www.eia.org/eig/ibis/ibis.html`

[IWR] Schimpf, U.
Personal support
IWR, Interdisciplinary Center for Scientific Computing, University of Heidelberg
`http://www.iwr.uni-heidelberg.de`

[IZM] Fraunhofer Institute
IZM, Institut für Zuverlässigkeit und Mikrointegration, Beriln
`http://www.izm.fhg.de`

[JTAG] JTAG
JTAG specification
JTAG Technologies IEEE Std 1149.1
http://www.jtag.com/homepage.html

[Ken94] Parker, K. P.
The Boundary-Scan Handbook
Kluwer Academic Publishers, ISBN 0-7923-9270-1, United States 1994

[KS99] Schmitt, K. et al.
Enthusiastic activities of the Heidelberg electronics department
The Heidelberg electronics group, IHEP Heidelberg 1996-1999

[LAr98] ATLAS Liquid Argon Group
ATLAS Liquid Argon Calorimeter Technical Design Report
ATLAS TDR-2, CERN/LHCC/96-41, CERN, Geneva 15 December 1996
http://atlasinfo.cern.ch/Atlas/GROUPS/LIQARGON/TDR/Welcome.html

[LEB95] Mass, A. et al.
Design and Test of an ASIC for Bunch Crossing Identification.
First Workshop on Electronics for LHC Experiments, CERN/LHCC/95-56, Lissabon, Spain 11-15 September 1995.

[LEB96] Mass, A. et al.
Front-End Digitization and readout System for the ATLAS Level-1 Calorimeter Trigger.
Second Workshop on Electronics for LHC Experiments, CERN/LHCC/96-39, Balatonfoerd, Hungary 23-27 September 1996.
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[LEB97] Pfeiffer, U. et al.
Performance of the Front-End Demonstrator System for the ATLAS Level-1 Calorimeter Trigger
Third Workshop on Electronics for LHC Experiments, CERN/LHCC/97-60, London, England 22-26 September 1997
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[LEB98] Pfeiffer, U. et al.
ATLAS Level-1 Calorimeter Trigger System Architecture
Fourth Workshop on Electronics for LHC Experiments, CERN/LHCC/98-36, Rome, Italy 21-25 September 1998.
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[LHC99] The LHC Study Group
The Large Hadron Collider Project
LHC homepage at CERN http://wwwlhc01.cern.ch/

[LHC] The LHC Study Group
The Large Hadron Collider — Conceptual Design Report

European Organization for Nucleear Research, CERN/AC/95-05(LHC), CERN, Geneva, 20 October 1995
http://www.cern.ch/CERN/LHC/YellowBook95/LHC95/LHC95.html

[LOI92] ATLAS collaboration
ATLAS Letter of Intent
CERN/LHCC/92-4, CERN, Geneva 1 October 1992

[Nie98] Niemann, B.
Datenkompression für die Auslese des ATLAS Level-1 Triggers
Diploma Thesis, IHEP-98-02, Universität Heidelberg, Germany 1998
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[Nol93] Nolting, W.
Grundkurs Theoretische Physik
Verlag Zimmermann-Neufang, 1993 Ulmen, ISBN 3-922410-20-0

[Pfe96] Pfeiffer, U.
Analoge optische Signalübertragung vom Kalorimeter zum ATLAS Level-1-Trigger
Diploma Thesis, IHEP-96-16, Universität Heidelberg, Germany 1996
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[Pfe99] Pfeiffer, U.
*Bunch-Crossing Identification for saturated calorimeter signal*s
ATLAS Trigger/DAQ note, ATL-DAQ-99-009, University of Heidelberg, Germany 17 May 1999.
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[Phos4] Phos4
PHOS4 - 4 Channel delay generation ASIC with 1 ns resolution
Datasheet, CERN — Microelectronics group

[Pov94] Povh, B. et al.
Teilche und Kerne — Einführung in die physikalischen Konzepte
Springer-Verlag, ISBN 3-540-56338-5, Berlin Heidelberg 199*4*

[PPR99] Pre-Processor ASIC (PPrAsic)
Specification of the Pre-Processor Asic
Preliminary Design Review, Universität Heidelberg, Germany 24. Juli 1999
http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html

[QT] Troll Tech AS
QT documentation
http://www.troll.no

[Ree97] Rees, D. L.
Bunch Crossing Identification Design
Discussion note, 16 May 1997

[Rem98] Readout Merger ASIC (RemAsic)
Readout Merger ASIC — User and Reference Manual
Internal document, Universität Heidelberg, Germany 18 March 1998
http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html

[Sch97] Schumacher, C.
Der ATLAS Level-1 Trigger, Auslese des Frontends
Diploma Thesis, HD-ASIC-31-0197, Universität Heidelberg, Germany 1997
http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html

[Sch98] Schumacher, C.
A flexible test and development system for the Pre-Processor
Talk held at the joint Level-1 Calorimeter Trigger meeting at Heidelberg, December 1998 http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html

[Sch99] Schumacher, C.
Status and future of the Heidelberg monitoring and control software
Talk held at the joit Level-1 Calorimeter Trigger meeting at RAL, June 1999
http://wwwasic.ihep.uni-heidelberg.de/atlas/docs.html

[Sco94] Scott, K. J.
Practical Simulation of Printed Circuit Boards and related structures
Research Studies Press LTD, ISBN 0-86380-161-7, England 1994

[Sig97] DF/SigNoise
DF/SigNoise - User Guide
Cadence Product Version 12.0, Openbook documentation, February 1997
http://www.cadence.com

[Sli] Boyle, O. et al.
The S-link Specification
CERN, HSI group, 27 March 1997
http://www.cern.ch/HSI/s-link/

[SPICE] SPICE
SPICE circuit simulator — User Guide
University of California at Berkley

[TDC99] RD12
Timing, Trigger and Control (TTC) Systems for LHC Detectors Homepage,21 May 1999 http://www.cern.ch/TTC/intro.html

[TDR98] ATLAS Level-1 Trigger Group
ATLAS First-Level Trigger Technical Design Report
ATLAS TDR-12, CERN/LHCC/98-14, CERN, Geniva 24 June 1998
http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html

[TDR99] ATLAS Collaboration
ATLAS Detector and Physics performance Technical Design Report

ATLAS TDR-14, CERN/LHCC/99-14, CERN, Geniva 25 May 1999
`http://atlasinfo.cern.ch/Atlas/GROUPS/PHYSICS/TDR/TDR.html`

[The97] DF/Thermax
DF/Thermax Expert for Package Design - User Guide
Cadence Product Version 12.0, Openbook documentation, February 1997
`http://www.cadence.com`

[Tho96] Thomas, D. E. et al.
The Veriolg Hardware Description Language
Kluwer, 1996

[TP94] ATLAS Collaboration
ATLAS Technical Proposal
CERN/LHCC/94-43, CERN, Geniva 15 December 1994
`http://www.cern.ch/pub/Atlas/TP/tp.html`

[URD98] ATLAS Level-1 Trigger Group
ATLAS Level-1 Calorimeter Trigger User Requirements Document
LVL1-Calo-URD-1.1.0, CERN, Geneva 10 April 1998
`http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/L1CalURD.ps`

[Wag96] Wagner, G.
Entwicklung einer Testumgebung für einen ASIC im Rahmen des ATLAS Level-1-Triggers
Diploma Thesis, IHEP-96-15, Universität Heidelberg, Germany 1996
`http://wwwasic.ihep.uni-heidelberg.de/atlas/publications.html`

[Wue] Würth Elektronik
`http://www.wuerth-elektronik.de`

Acknowledgments

At the beginning of this three year thesis I had the chance to work at the Particle Physics Department of the Rutherford Appleton Laboratory in England. During that six month stay I had the opportunity to meet our English colleagues and to learn about our prototype Level-1 Calorimeter Trigger electronics and data acquisition software. First of all I want to thank those who have made this very exciting visit possible. I owe many thanks to Prof. Karlheinz Meier, my supervisor, and to Prof. Franz Eisele, for their promotion during that period and Prof. G. E. Kalmus for the permission to stay. To all the members of the ATLAS UK-group I owe many thanks for their openness for discussion. Norman Gee, my person in charge, I want to thank for his tolerance. He always could find the time for each of my concerns. Tony Gillman and Viraj Perera I want to thank for the almost infinite number of useful suggestions and last but not least my great respect and thankfulness go to Eric Eisenhandler for his competence and his incredible discipline in proof reading of this thesis. On private grounds, I would like to say to all friends that I have made there, that I will never forget this time. Special thanks and my great respect go to my landlady, Jenny Colbourne, in Oxford for her atmosphere of friendliness and for all the evening suppers we have enjoyed with plenty of red wine.

Coming back to Heidelberg in July 1997 I immediately felt at home again. The fruitful discussions we have had within our ATLAS Heidelberg group was the foundation of the success of all the work I have done here. Many thanks go in particular to my supervisor, Prof. Karlheinz Meier, who has provided me with any type of support one could think of. He always took questions and suggestions seriously and it was easy to share his physics enthusiasm. Paul Hanke was my source of experience and guidance in the world of particle physics experiments. His humour has always made long meetings lighthearted and has raised the motivation. My colleague and friend Cornelius Schumacher has contributed to this work through his excellent software knowledge and many useful discussions. Prof. Norbert Herrmann I would like to thank for being my second referee.

The various joint meetings we have had were very informative and of great pleasure. Many thanks to the members of the ATLAS Level-1 Calorimeter Trigger collaboration, the people from Birmingham, Mainz, Queen Mary and Westfield College, RAL, and the people from Stockholm.

After the design of a prototype ASIC, I started with the Finco design at the beginning of 1998. The pleasant atmosphere within our Heidelberg ASIC laboratory and the experience of Michael Keller has guided me through the obstacles of the ASIC design process. Many thanks I owe to the ASIC group for the help and advice on the countless occasions when I needed assistance. During the Multi-Chip Module design, which was finished at the end of 1998, I have received many useful tips and support from Andreas Schilpp. Thanks also to Eric Jung for the collaboration and assistance during the Flip-Chip mounting process.

Many thanks go to the Electronic Department for constructing the numerous support and auxiliary electronics which was of utmost importance. In particular thanks to Klaus Schmitt who always has suggestions to a problem and his correctness is a warranty for electronics which indeed will work.

The balance between work and pleasure time would be long lost without numerous friends, thanks for making the last few years so much fun. I owe many thanks also to my uncle Rudolf Schäfer and my friend Michael Walter for their useful comments and the typing errors they have recovered. Last, but far from least, I would like to thank my parents and my sister for their encouragement, the love they devoted to me is far more than I could possibly describe here.