Dissertation
submitted to the
Combined Faculties of the Natural Sciences and Mathematics
of the Ruperto-Carola-University of Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

Put forward by
Jan Jongmanns
born in Nettetal

Oral examination on July 27, 2017

The Upgrade of the PreProcessor of the ATLAS Level-1 Calorimeter Trigger for LHC Run-2

Referees: Prof. Dr. Hans-Christian Schultz-Coulon
Prof. Dr. Norbert Herrmann

Zusammenfassung:

Die erste Stufe des Kalorimeter-Triggers des ATLAS Experiments am CERN untersucht innerhalb von ∼2 µs 7200 analoge Kalorimeter-Signale, um Objekte mit hoher transversaler Energie zu identifizieren und Energiesummen zu erstellen. Er besteht aus dem PreProcessor, welcher die Kalorimeter-Signale digitalisiert, sie der korrekten Strahlkreuzung des LHC zuordnet und sie bezüglich der transversalen Energie kalibriert, sowie zwei Prozessoren, welche zu jeder Strahlkreuzung die Ergebnisse des PreProcessors analysieren um Elektronen, Photonen, τ-Teilchen und Jets zu identifizieren und um die fehlende sowie die gesamte Transversale Energie zu berechnen. Um die für Run-2 erwareten höheren Raten und den erhöhten Pile-Up zu verarbeiten, wurden die ASIC-basierten Multichip Module, welche die Signalverarbeitungsalgorithmen implementieren, im Rahmen des ersten langen Shutdown des LHC durch verbesserte, FPGA-basierte Module ersetzt. Die Signalverarbeitung des neuen Multichip Moduls basiert auf der Run-1 Implementierung und erweitert diese durch verbesserte Rauschfilter, unabhängige Energiekalibrationen für elektromagnetische und hadronische Objekte, Verbesserungen bei der Prozessierung saturierter Signale und eine für jede Strahlkreuzung berechnete Pile-Up Korrektur. In dieser Arbeit wird das neue Multichip Modul beschrieben, mit einem Fokus auf die Verbesserungen der Signalverarbeitungsalgorithmen und der daraus resultierenden verbesserten Performance des Level-1 Kalorimeter-Triggers in Run-2.

Abstract:

The Level-1 Calorimeter Trigger of the ATLAS experiment at CERN analyses 7200 analogue signals from the calorimeters within ∼2 µs to identify objects with high transverse energy and to determine energy sums. It consists of a PreProcessor, which digitises the signals, assigns them to the correct LHC bunch crossing and calibrates them to transverse energy, and two object finding processors. During the first long shutdown of the LHC, the main processing unit of the PreProcessor, the ASIC-based Multichip Module, was replaced by an improved FPGA-based version to adapt to the increased rates and pile-up that are expected for Run-2. The new Multichip Module builds upon the Run-1 signal processing implementation and expands it with enhanced noise filtering techniques, independent energy calibration for electromagnetic and hadronic objects, improved processing of saturated signals and bunch-crossing-wise pile-up subtraction. This thesis describes the new Multichip Module, focussing on the improvements to the signal processing algorithms and the increased performance gained by their implementation.

Contents

1 Introduction **11**

2 The Standard Model of Particle Physics **14**
2.1 Standard Model of Particle Physics 15
2.2 Beyond the Standard Model 17

3 The ATLAS Experiment at Large Hadron Collider **18**
3.1 The Large Hadron Collider 18
3.1.1 LHC Machine 18
3.1.2 Proton Beam Structure 21
3.1.3 Aspects of Proton-Proton Collisions 22
3.2 ATLAS Detector 22
3.2.1 Coordinate System 23
3.2.2 Inner Detector 24
3.2.3 Calorimetry 25
3.2.4 Muon Spectrometer 30
3.2.5 Magnets 31
3.2.6 Further Subdetectors 32
3.3 ATLAS Detector Control System 33
3.4 ATLAS Trigger and Data Aquisition System 33
3.4.1 Level-1 Trigger 35
3.4.2 High Level Trigger and Data Aquisition 37
3.4.3 TDAQ Upgrade in LS1 37
3.5 ATLAS Operation 38

4 The ATLAS Level-1 Calorimeter Trigger **39**
4.1 Functional Overview 39
4.2 Analogue Input Signals 40
4.2.1 Input Signal Path 40
4.2.2 Signal Shape 42
4.3 PreProcessor 44
4.3.1 Functionality 44
4.3.2 Hardware Implementation 46
4.4 Cluster Processor 52
4.4.1 Functionality 52
4.4.2 Hardware Implementation 53

4.5 Jet-Energy Processor . . . 54
4.5.1 Functionality . . . 54
4.5.2 Hardware Implementation . . . 55

5 The new Multichip Module 56
5.1 Hardware Implementation . . . 57
5.2 Digital Signal Processing . . . 59
5.2.1 ADC Input . . . 60
5.2.2 Input Timing . . . 62
5.2.3 Bunch-Crossing Identification Logic . . . 64
5.2.4 LVDS Transmission . . . 69
5.2.5 Event Data Readout . . . 70
5.2.6 Monitoring . . . 72
5.2.7 Testing Facilities . . . 73
5.2.8 Clock Management . . . 74
5.2.9 Configuration . . . 76

6 Functional Tests of the new Multichip Modules 79
6.1 Test Overview . . . 79
6.2 Functionality Tests . . . 80
6.2.1 Test Setups . . . 81
6.2.2 Configuration Test . . . 84
6.2.3 Rate-Metering Test . . . 85
6.2.4 LVDS Transmission Tests . . . 86
6.2.5 Event Readout Test . . . 87
6.2.6 Analogue Input Tests . . . 89
6.2.7 Input Timing Tests . . . 95
6.2.8 Histogramming Test . . . 101
6.3 Summary of Test Results . . . 101
6.4 Installation of nMCMs at CERN . . . 102

7 Improved Bunch-Crossing Identification on the nMCM 103
7.1 Digital Filter Extension . . . 104
7.2 Dynamic Pedestal Correction . . . 106
7.2.1 Pile-Up Effects on L1Calo Input Signals . . . 106
7.2.2 Pedestal Correction Algorithm . . . 108
7.2.3 Pedestal Correction Algorithm Implementation . . . 110
7.2.4 Synchronisation of the Pedestal Correction . . . 111
7.2.5 Readback of Pedestal Average Memory . . . 112
7.3 Separate LUTs for CP and JEP Systems . . . 113
7.4 Improved BCID of Saturated Signals . . . 114
7.4.1 Updated BCID Decision Logic . . . 116
7.4.2 80 MHz Saturated BCID Algorithm . . . 117

8 Additional Monitoring and Testing Facilities on the nMCM **119**
8.1 Monitoring Algorithms . . . 119
8.2 Event Data Readout . . . 119
8.2.1 Additional Readout on the nMCM . . . 120
8.2.2 Readout of 80 MHz ADC Data . . . 121
8.3 Updates to Playback . . . 122
8.3.1 Playback of ADC Data at 80 MHz . . . 122
8.3.2 Separation of Histogramming and Playback . . . 123

9 Performance of the nMCM **124**
9.1 Input Noise . . . 124
9.2 Autocorrelation Filters . . . 127
9.3 Pedestal Correction . . . 128
9.4 Separate Energy Calibration for CP and JEP . . . 130
9.5 BCID of Saturated Signals . . . 130

10 Summary **133**

A Lists **134**
A.1 List of Figures . . . 134
A.2 List of Tables . . . 135

B Bibliography **137**

1 Introduction

Particle physics concerns itself with the understanding of nature at the most fundamental level. By analysing the most fundamental and indivisible constituents of matter and their interactions, understanding of the behaviour of physical systems at larger scales can be gained.

To probe the small distance scales of particle physics, high energy scales have to be reached. To do so, modern experiments utilise strong particle accelerators, which use electromagnetic fields to accelerate particles and to subsequently collide them. The prime example for this is the Large Hadron Collider located at the CERN laboratory in Geneva, Switzerland. It is a synchrotron using high frequency cavities to accelerate two beams of protons to energies of up to 13 TeV per particle, with superconducting magnets bending the beams onto a circular path.

The high energy protons are collided in four interaction points along the ring, causing interactions which can produce a plethora of possible final state particles. By measuring these particles and their properties, insights into the processes involved in the high-energy collision is gained. To this end, each of the collision points is equipped with a large particle detector, which detects the particles that are produced in the collision.

One of these is the ATLAS experiment, a general purpose detector that is designed to measure all particles produced in the collisions as inclusively as possible. An integral part of the experiment is the trigger system. Its purpose is to select those events which should be stored for later data analysis, reducing the input rate of 40 MHz to 1 kHz. Within this system the Level-1 Calorimeter Trigger has the task to analyse data from the ATLAS calorimeters to find candidates of highly energetic jets, taus, electrons and photons as well as the missing and total transverse energy. It is organized into several processing stages, with a dedicated PreProcessor finding large energy depositions and calibrating them and two downstream object finding processors which execute sliding window algorithms to identify particle candidates.

The Level-1 Calorimeter Trigger system has been upgraded during the first long shutdown of LHC in 2013 and 2014, both by replacing hardware components and by upgrading the signal processing implementation. On the PreProcessor side this upgrade consists of the replacement of the central processing platform, the Multichip Module, by a new version based on modern and flexible components. Most prominently the ASIC performing the main signal processing has been replaced by an FPGA, offering enough resources to implement the original ASIC algorithms as well as new advanced features.

Within this thesis the upgraded, new Multichip Module is described, with a focus on the signal processing implementation and production tests. The second chapter

of this document gives an overview of the Standard Model of Particle Physics. In the third chapter both the LHC and ATLAS machines are described, while chapter four focuses on the Level-1 Calorimeter Trigger system. The fifth chapter describes the hardware implementation of the new Multichip modules and the initial version of the signal processing algorithms. Chapter six describes the production test procedure used to verify the functionality of the new hardware. Chapters seven and eight provide detailed descriptions of the updates of the signal processing algorithms. The performance of the upgraded system during the first two years of operation is described in chapter nine. Finally, chapter ten forms a conclusion and overall summary of the presented material.

Author's Contributions

The development of a system like the one presented in this thesis requires expertise in hardware, firmware and software development, as well as an understanding of the physical processes involved in the production, decay and detection of physics signatures. As such, most of the topics presented here are results of a collaborative effort within a larger group. Nevertheless there are specific parts which the author explicitely contributed to the overall effort.

The author contributed significantly to the planning, design and implementation of the new trigger algorithms included in the new Multichip Module. By improving the digital filter, the author made it possible to use more efficient noise filtering schemes that assure a high efficiency for low energy signals. The author separated the energy calibration for the electromagnetic and hadronic objects, increasing the flexibility for the respective calibration procedures. The author was central to the development of the pedestal correction algorithm at all stages, from the first conceptual design to the details of the final implementation. Furthermore the author implemented the new 80 MHz based bunch-crossing identification logic for saturated calorimeter signals. To further improve the processing of signals in the saturated regime, the author developed and applied the updates of the BCID decision logic.

The author also provided several improvements of the monitoring and testing algorithms on the new Multichip Module. The separation of the energy calibration requires additional monitoring procedures to observe the respective results independently, which the author implemented. To validate the additional functionality in the saturated regime, the author developed and implemented a readout technique for 80 MHz ADC samples. For the same reason, the playback functionality was expanded by the author so that it can be executed at the higher frequency of 80 MHz.

The author also contributed to the production tests of the new Multichip Modules. The author developed several of the software tools described in this thesis, specifically the tests of the monitoring algorithms and the Signal Generator scan and the tests of the input stage of the CALIPPR FPGA. The author also adapted a standalone version of the L1Calo online simulation software to the test environment, improving the flexibility for the input patterns that are used in the test procedure.

The author performed a substantial part of the production tests himself, including both the execution of the tests and the upload of the results to an online database. The author was also involved in one of the installation campaigns at CERN, where part of the system was equipped with the new modules.

Though it is not described in this thesis, the author also implemented the neccessary changes in the ReM FPGA to enable the features of the new Multichip Module described above. The author also provided advice and technical support to the various studies and analyses of the PreProcessor performance. Last but not least the author was involved in the daily operations of the PreProcessor system for most of the time of the thesis, most strongly during a one year stay at CERN.

2 The Standard Model of Particle Physics

Particle physics aims to describe nature by analysing the most fundamental constituents of matter and their interactions with each other.

While the idea that matter consists of indivisible particles at the smallest scale had already been proposed by philosophers in ancient cultures like Greece, it took until the early 19th century for a quantifyable theory to emerge. Between 1805 and 1808, the chemist John Dalton, noticing how chemical elements form compounds only in certain ratios, proposed that all elements are made up of indivisible particles calles atoms, with different elements consisting of different kinds of atoms [1]. With the discovery of the electron by Thomson in 1897 [2] and experiments by Rutherford and Royds [3, 4] in the first decade of the 20th century, the structure of atoms as a small, positively charged massive nucleus surrounded by negatively charged electrons was established. Insight into the structure of the nucleus was gained through the investigations of isotopes. Proven through experiments by Soddy [5, 6], Aston [7, 8] and Thomson [9, 10], their existance lead to the model of a nucleus formed of positively charged protons and neutral neutrons. The development of the atomic theory culminated in the proofs for nuclear fission by Hahn [11] and Meitner [12] and nuclear fusion by Oliphant [13] and Bethe [14].

In the second half of the 20th century, the development of particle accelerators and the observation of cosmic rays led to the discovery of many new particles beyond the ones that make up our everyday world (e.g. [15, 16, 17, 18]). With dozens of mesons and baryons discovered in the 1950s, the term 'particle zoo' was brought up to describe the large number of different, apparently fundamental particles. In 1961, Gell-Mann managed to bring order to the zoo, by ordering the particles into octets and decuplets according to their charge and the newly introduced 'strangeness' quantum number [19]. Noticing this symmetry ultimately led to the introduction of the quark model, which describes mesons and baryons as composite particles, with their constituents, the quarks, being truly elementary, indivisible particles [20, 21, 22].

This process of understanding the large number of observed particles as a consequence of the underlying, simpler structure is mirrored by the unification process of the forces that govern the interactions of particles. In 1873, Maxwell published a treatise in which he formulated the electric and magnetic phenomena as consequences of a single electromagnetic force. Glashow [23], Salam [24] and Weinberg [25] managed to unify the electromagnetic force and the weak force, that governs radioactive decay, in a consistent way in the late 1960s. This model has been combined with the theory of Quantum Chromodynamics that describes the strong interactions between quarks to arrive at what is known as the Standard Model of Particle Physics in the following years. Since then, the Standard Model has been successful in accu-

rately predicting the outcome of the majority of particle physics experiments. The most prominent example of the recent past is the discovery of the Higgs Boson at the Large Hadron Collider in 2012 [26, 27]. Nevertheless, some unanswered questions remain that are not solved in the Standard Model, making the search for 'new physics' one of the main drivers for many particle physics experiments.

This chapter gives a brief overview of the Standard Model in section 2.1, and discusses some of the remaining open questions not answered by the Standard Model in Section 2.2.

2.1 Standard Model of Particle Physics

The *Standard Model of Particle Physics* is a relativistic quantum field theory that describes the elementary particles and their interactions. In a quantum field theory, both matter and interactions are modeled by fields, whose quantised excitations are interpreted as particles. For example, all electrons in the universe are excitations of a single electron field, and the photons are similarly represented by the electromagnetic field. The interaction between particles in this picture is realised by the exchange of intermediate particles, explaining e.g. the electromagnetic interaction as an exchange of photons between charged particles.

On a more fundamental level, interactions are described using gauge symmetries. Given that a theory is symmetric under a *global transformation*[1] of the matter fields, the principle of locality suggests that the fields should be symmetric under *local tranformations*[2] as well. This invariance under local transformations, also called gauge symmetry, can be achieved by introducing additional *gauge fields* in the theory, whose transformations cancel the effects of the transformations of the matter fields. The resulting theory is invariant under gauge transformations and the gauge fields realise interactions between the matter particles.

A key difference between matter particles and interaction particles is their statistical behaviour. Matter particles are *fermions*, i.e. they obey Fermi-Dirac statistics and no two of them can occupy the same state. Interaction particles are *bosons*, i.e. they obey Bose-Einstein statistics and an arbitrary number of them can occupy the same state. Following the spin-statistics-theorem (REF!!!), all fermions have half-integer spin ($\frac{1}{2}$, $\frac{3}{2}$, ...) while bosons have integer spin (0, 1, ...). The bosons corresponding to a gauge field are called *gauge bosons*.

The SM is built from three such gauge theories:

1. Quantum Electrodynamics is a gauge theory using the $U(1)$ symmetry group. It describes the interaction of charged particles via exchange of photons, γ.

[1] A transformation is referred to as global if the transformation parameter is the same at all points in spacetime. A simple example would be a constant phase shift in the field.

[2] Following the previous example, in a local transformation the phase shift is different at each point in space and time.

2. The weak interaction is a gauge theory of the symmetry group $SU(2)$. All particles take part in the weak interaction, transmitted via the W^+, W^- and Z bosons.

3. The strong interaction is described by Quantum Chromodynamics, an $SU(3)$ gauge theory. It describes the strong interaction between quarks, realised as exchange of gluons.

		Name	Symbol	Mass $[\frac{\text{MeV}}{c^2}]$	Charge [e]	Spin
Fermions	Leptonen	Electron	e^-	$5.11 \cdot 10^{-1}$	-1	1/2
		Muon	μ^-	$1.06 \cdot 10^2$	-1	1/2
		Tau	τ^-	$1.78 \cdot 10^3$	-1	1/2
		Electron Neutrino	ν_e	$< 2 \cdot 10^{-6}$	0	1/2
		Muon Neutrino	ν_μ	< 0.19	0	1/2
		Tau Neutrino	ν_τ	$< 1.82 \cdot 10^1$	0	1/2
	Quarks	Up Quark	u	2.2	$+2/3$	1/2
		Charm Quark	c	$1.27 \cdot 10^3$	$+2/3$	1/2
		Top Quark	t	$1.73 \cdot 10^5$	$+2/3$	1/2
		Down Quark	d	4.7	$-1/3$	1/2
		Strange Quark	s	$9.6 \cdot 10^1$	$-1/3$	1/2
		Bottom Quark	b	$4.18 \cdot 10^3$	$-1/3$	1/2
Gauge Bosons		Photon	γ	0	0	1
		Gluon	g	0	0	1
		W-Boson	$W^\pm$	$8.04 \cdot 10^4$	± 1	1
		Z-Boson	Z	$9.12 \cdot 10^4$	0	1
		Higgs-Boson	H	$1.25 \cdot 10^5$	0	0

Table 2.1: List of particles included in the Standard Model. The antiparticles of the fermions are not listed explicitely. Values adapted from [28].

Table 2.1 gives an overview of the elementary particles in the SM.

The fermions are divided into two groups, the leptons and the quarks. The different particle types differ in their behaviour, as only the quarks take part in the strong interaction.

Both the leptons and the quarks are divided into three generations, with each generation consisting of a pair of particles. The particles in the higher generations have the same properties as the ones in the first generation, except for the mass, that rises with the generations.

Each lepton generation consists of a massive, charged particle and an approximately massless neutrino. While the massive partner takes part in both the electromagnetic and the weak interaction, the neutrinos only interact weakly.

The quark generations each consist of one *up-* and one *down-type* quark. These have different electromagnitic coupling strengths, with charges of $+\frac{2}{3}$ and $-\frac{1}{3}$, respectively. In contrast to the leptons, all quarks take part in all three interactions.

The properties of the gauge bosons result from the structure of the underlying gauge group. Photons are massless particles that do not couple directly to each other, because the underlying $U(1)$ group is abelian. The gluons, while also massless, do couple to each other, due to the non-abelian nature of the $SU(3)$ gauge group.

The $SU(2)$ symmetry of the SM is spontaneously broken: while the symmetry is present in the theory, the ground state violates the symmetry. This is caused by an additional field, the Higgs field, whose potential energy leads to the asymmetric ground state. As a consequence of the spontaneous symmetry breaking, the $SU(2)$ bosons aquire a non-zero mass and an additional particle is introduced to the model: the Higgs boson H.

2.2 Beyond the Standard Model

While all other particles in the standard model had been known to exist for a long time already, the Higgs boson remained the last missing piece until its discovery in 2012. With its discovery, all particles in the Standard Model have been tested experimentally, further solidifying it as the central theory of particle physics.

Nevertheless, open questions remain which the Standard Model does not provide an answer to. These include:

- Astronomical observations show that the *baryonic matter*, i.e. ordinary matter consisting of atoms, only makes up about 5% of the observed universe, with the remainder consisting of *dark matter* and *dark energy* [29]. Neither of these are described by the Standard Model.

- While the Standard Model formulates the electromagnetic, weak and strong interactions in a cohesive framework, no microscopic formulation for the *gravitational force* exists as of yet.

- The asymmetry between matter and antimatter observed in the universe is not reflected in the existing asymmetries in the Standard Model [30]. Thus, additional mechanisms that further this asymmetry have to exist.

To answer these questions conclusively, further particle physics experiments are required. One of these is the ATLAS experiment at the Large Hadron Collider, both of which are described in the next chapter.

3 The ATLAS Experiment at Large Hadron Collider

One way to probe particle physics processes is colliding particles at high energies. For this purpose, the Large Hadron Collider (LHC) has been constructed to accelerate two beams of protons and collide them at four interaction points. It is the flagship project of the European Organization for Nuclear Research, CERN[1], in Geneva, Switzerland.

ATLAS is one of four large experiments situated at the LHC. Installed at one of the four pp interaction points, the ATLAS detector records the particles that are produced in the collisions as inclusively as possible.

This chapter provides an overview of the LHC machine and its typical operation parameters (Section 3.1) and gives a detailed description of the ATLAS detector, (Section 3.2).

3.1 The Large Hadron Collider

The LHC [31] is a circular particle accelerator with a circumference of 26.7 km. It is installed in the tunnel that previously housed the Large Electron Positron Collider (LEP) between 45 and 170 m underground. It is optimised to collide protons, but can also be operated with heavy ions.

The particles are collided in four interaction points, each of which is equipped with one of four major experiments. The ATLAS [32] and CMS [33] experiments are general purpose detectors that are designed to record pp collisions in a wide phase space. The LHCb [34] experiment is a forward spectrometer optimised for the detection of B mesons. The ALICE [35] experiment focusses on heavy ion collisions to probe the quark-gluon-plasma that forms at the high energy densities produced in these processes.

3.1.1 LHC Machine

Figure 3.1 shows the CERN accelerator complex. It consists of a chain of pre-accelerators that prepare the proton beams for injection into the LHC. The protons are produced by stripping hydrogen atoms of their electrons in a strong electric field. They are then accelerated in a linear accelerator (LINAC2), followed by a series of increasingly stronger circular accelerators: the Proton Synchrotron

[1] Conseil européen pour la recherche nucléaire

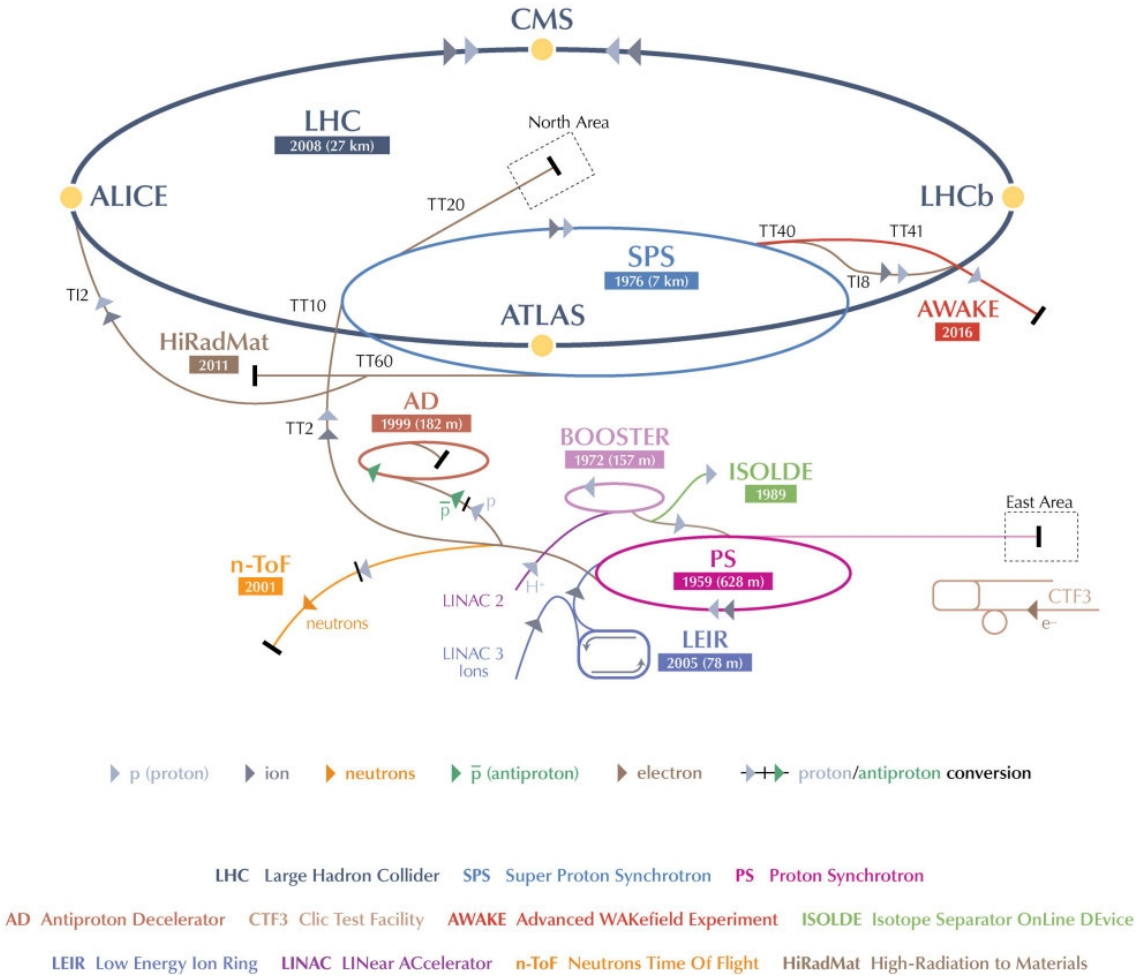

Figure 3.1: The CERN accelerator complex. [36]

Booster (BOOSTER), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS). At injection into the LHC, the protons have an energy of about 450 GeV

The LHC itself is a circular synchrotron that consists of eight straight sections that are connected by eight curved sections. One of the straight sections is equipped with superconducting radiofrequency cavities that are operated at 400.8 MHz to accelerate the proton beams. The curved sections are equipped with superconducting dipole magnets that produce a magnetic field of up to 8.34 T to bend the proton trajectories and keep them in the ring. At the remaining sections, quadrupole and octopole magnets are installed that focus the proton beams, most strongly at the four interaction points.

The central performance parameters of a particle accelerator are the centre-of-mass energy $\sqrt{s}$, that determines which particle physics processes can be effectively probed in a collision, and the luminosity $\mathcal{L}$, that describes the number of collisions that are produced. The number of collision events N that result in a final state X

is given by

$$N = \mathcal{L} \cdot \sigma_{pp \to X}, \tag{3.1}$$

where the cross-section $\sigma_{pp \to X}$ describes the probability of the process $pp \to X$ to occur in a collision. The cross-section depends on the underlying physical laws governing the process and the available phase space, both of which depend strongly on the centre-of-mass energy. Achieving a large number of events for rare final states thus depends on both maximizing the luminosity and reaching a sufficiently large centre-of-mass energy.

As it is defined above, the luminosity $\mathcal{L}$ corresponds to the total amount of events that are accumulated over a defined period of time. For a given point in time, the *instantaneous luminosity* $\mathcal{L}_{inst}$ gives the rate at which events are produced. It is defined by $\mathcal{L}_{inst} = \frac{d}{dt}\mathcal{L}$. In this context, $\mathcal{L}$ is also referred to as the *integrated luminosity*.

The LHC is designed to collide protons at a centre-of-mass energy of 14 TeV, corresponding to 7 TeV per beam. The design luminosity is $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

The LHC is operated in a series of *fills*, in which it contains proton beams and collides them. At the start of each fill, protons pass through the accelerator chain described above and are injected into the LHC. Once the LHC is filled, it starts accelerating the protons, until they reach their target energy. The focussing magnets around the collision points are configured such that the beams are brought into collision, to start the luminosity production.

During a fill, the instantaneous luminosity falls as a function of time, as protons are lost from the beams with each collision that occurs. Once it falls below a certain threshold, it is more efficient to restart the operational cycle, by *dumping* the beams and injecting new beams into the LHC. One fill can last many hours, e.g. the longest fill of 2016 operations contains 37 continous hours of luminosity production[2]. Interfill periods usually last for about one to two hours.

After the LHC first started operation in September 2008, a faulty electrical connection between two of the superconducting magnets caused a gas explosion, damaging a section of the accelerator [37]. This necessitated extensive repairs, so that the LHC was only restarted in November 2009. In a conservative approach, the beam energy was slowly increased from 450 GeV up to 3.5 TeV until March 2010. Data was recorded at this energy until November 2011, collecting a dataset of $5.6\,\mathrm{fb}^{-1}$[3]. In March 2012, the LHC operation resumed at a centre-of-mass energy of 8 TeV. With luminosities up to $7.73 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, a total dataset of $23.3\,\mathrm{fb}^{-1}$ was recorded until December 2012. This whole period of data taking is referred to as Run-1 of the LHC [38].

Run-1 was followed by the first long shutdown (LS1) of the LHC, that lasted until early 2015. During this time, the accelerator complex was prepared to provide both higher energies and higher luminosities. In the years 2015 and 2016, LHC success-

[2] Fill 5045

[3] The barn (b) is used as a unit for cross sections in particle physics. It is defined as $1\,\mathrm{b} := 10^{-24}\,\mathrm{cm}^2$. From Equation 3.1 it follows that the unit for luminosity is the inverse of this, b^{-1}.

fully provided *pp* collisions at a centre-of-mass energy of 13 TeV. In 2016, the luminosity surpassed the design goal, with peak luminosities up to $1.38 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ having been achieved. The total integrated luminosity is $4.2\,\mathrm{fb}^{-1}$ in 2015 and $38.5\,\mathrm{fb}^{-1}$ in 2016 [39].

The period of operation following LS1 is referred to as Run-2, and is planned to last until the end of 2018. It is to be followed by the second long shutdown (LS2), which will see further improvements to the accelator chain. This is expected to increase the luminosity beyond $2 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

3.1.2 Proton Beam Structure

The LHC proton beams consist of bunches of protons that typically contain of the order of 10^{11} particles. Each proton beam consists of 3564 possible bunch slots. Due to operational constraints in the filling of the LHC and the dumping of the beams, not all of these slots actually contain bunches. By design, a maximum of 2808 bunches can be *filled*, while the remainder stays *empty*.

The 'collision' of two bunch slots at one of the interaction points is referred to as *Bunch Crossing* (BC). BCs occur at a frequency of 40.08 MHz, i.e. are spaced approximately 25 ns apart. One revolution, or *orbit*, of the LHC corresponds to 3564 BCs, after which the same bunches collide again. Each BC in the orbit is thus assigned a unique BC Number[4] between 0 and 3563.

In a typical data taking run, the protons are organised into *trains* of filled bunches, with *gaps* of empty bunches between them. This results in a similar structure for the BCs. The orbit contains several trains of *paired* BCs, where the bunches are filled in both beams, that are separated by gaps of *unpaired* or *empty* BCs, where either only one or none of the beams contains a filled bunch.

A further distinction is made in the bunch spacing within the trains. In Run-1, bunch trains typically consisted of alternating filled and empty bunches, resulting in a 50 ns spacing between collisions. In Run-2, all bunches in a bunch-train are filled, resulting in a bunch spacing of 25 ns. This has consequences for the detectors that record the collisions, as it changes how events overlap with each other.

This overlap between events is referred to as *pile-up*. It is quantified by the mean number of *pp* interactions that take place in a collision of two filled bunches, μ. It depends on the machine operation and the luminosity: if more protons interact, at the same time, the luminosity increases. While maximising the luminosity is one of the keys to enhance the sensitivity to rare processes, it has to be balanced against the increasing pile-up, which deteriorates the detector performance. In 2016, the average pile-up was measured to be $< \mu >= 24.9$.

The magnitude of the pile-up also depends on the bunch spacing. By decreasing the bunch spacing, more filled bunches can be fit into the LHC at the same time. This allows to reach the same luminosity as for larger bunch spacing with less

[4]This is also often referred to as BC Identifier, or BCID. To avoid confusion with the BC Identification algorithms discussed later, the term BC Number will be used throughout this document.

focussed beams, and thus reduced μ. While this reduces the pile-up originating from the same BC (*in-time pile-up*), the newly introduced collisions in the preceeding BC lead to an additional, *out-of-time pile-up*. Depending on the specifics of the detector, either of these can have a larger effect.

With the increasing luminosity, pile-up is thus one of the main challenges for the detectors in their respective upgrade programs.

3.1.3 Aspects of Proton-Proton Collisions

Proton-proton (pp) collisions offer distinct advantages compared to other experiments. As protons have a higher mass than electrons, $m_p = 938\,\mathrm{MeV}$ compared to $m_e = 511\,\mathrm{keV}$, the energy loss due to Bremsstrahlung, which scales as $P \sim m^{-4}$, is far less for proton beams than for electron beams. This allows pp colliders to reach much higher centre-of-mass energies than e^+e^- colliders using the same geometry. Comparing pp to $p\bar{p}$ colliders, protons are easier to mass produce than antiprotons, allowing for higher statistics in particle collisions.

The main drawback in colliding protons lies in their nature as composite particles. In the low energy limit, each proton consists of a triplet of *valence quarks*, two up quarks and one down quark. At higher energy scales, the deeper structure of the proton becomes apparent, with the gluons that bind the valence quarks together taking part in the interaction, as well as *sea quarks*, that are produced by gluons undergoing pair conversion within the proton. All these constituents of the proton are referred to as *partons*.

The interaction of two protons at high energies is equivalent to the interaction of partons. Because each parton only carries a fraction of the total momentum of the proton, the centre-of-mass energy of the parton-parton interaction is not known a-priori. Furthermore, even in a symmetric pp collision where both protons have the same momentum, the parton-parton collision is typically asymmetric. This has important consquences for the design of particle detectors at pp colliders, because the initial state momentum along the beam axis is not known. Due to this reason, final state particles are usually analysed with regard to their *transverse momentum*, p_T, i.e. the momentum component in the plane transverse to the beam axis, which is zero in the initial state.

The type of parton that participates in the interaction depends on the energy scale of the pp collision. At LHC energies, gluons typically have the highest probability of interacting.

3.2 ATLAS Detector

ATLAS [32] is one of four large experiments at the LHC. It is a general-purpose particle detector optimised on recording pp collisions as inclusively as possible.

Figure 3.2 shows the ATLAS detector. It consists of three major parts: a cylindrical *barrel* around the pp interaction point in the centre and two *endcaps* attached

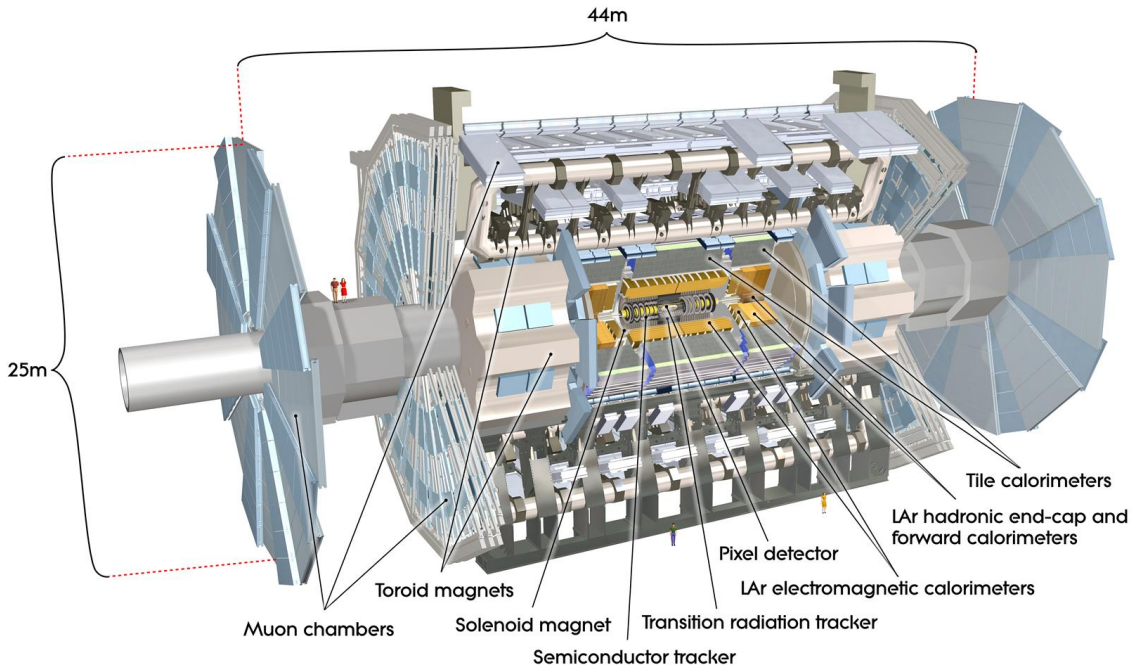

Figure 3.2: The ATLAS detector. [32]

to the sides of the barrel. It is 44 m long, 25 m high and weighs 7000 t in total.

The detector consists of several layers of sub-detectors. Closest to the interaction point is the *Inner Detector*, which is surrounded by a solenoid magnet. It measures the tracks of traversing charged particles. This allows to reconstruct the momentum of the particles from the track curvature in the magnetic field, and to identify the point of origin of the tracks. The next layer is the *Electromagnetic Calorimeter*. It is built to contain electromagnetic showers that are initiated by incident electrons and photons, thus measuring the energy of these particles. While hadronically decaying particles also initiate showers in the electromagnetic calorimeter, these typically extend much further than electromagnetic showers. The *Hadronic Calorimeter* is thus required to fully contain these showers. The outermost sub-detector is the *Muon Spectrometer*, that detects any charged particles that are not stopped within the calorimeter and measures their momentum in a magnetic field generated by large toroid magnets.

3.2.1 Coordinate System

ATLAS uses a right-handed coordinate system with its origin in the nominal interaction point. The x-axis points towards the centre of the LHC ring, the y-axis points upwards towards the surface and the z-axis points along the beam pipe. The azimuth angle ϕ is defined in the x-y-plane from 0 to 2π starting from the x axis direction. The polar angle θ is defined from 0 to π starting from the z-axis direction.

Typically the polar angle is expressed in terms of the pseudorapidity,

$$\eta = -\ln \tan(\theta/2). \tag{3.2}$$

A pseudorapidity of zero corresponds to the transverse direction, while it approaches $+/- \infty$ for the positive/negative z-axis direction.

In the high-energy limit, the pseudorapidity is equal to the rapidity,

$$y = \frac{1}{2} \ln \left(\frac{E + p_L}{E - p_L} \right), \tag{3.3}$$

where p_L is the momentum component along the beam pipe. Differences in rapidity are invariant under lorentz boosts along the longitudinal axis, making it a useful quantity for physics analyses.

3.2.2 Inner Detector

The Inner Detector detects traversing charged particles and is used to reconstruct their trajectories. By measuring the bending radius of these tracks in the solenoid magnetic field, the momentum of the particles is obtained. By extrapolating the tracks, the primary collision vertex as well as decay vertices of secondary particles are identified.

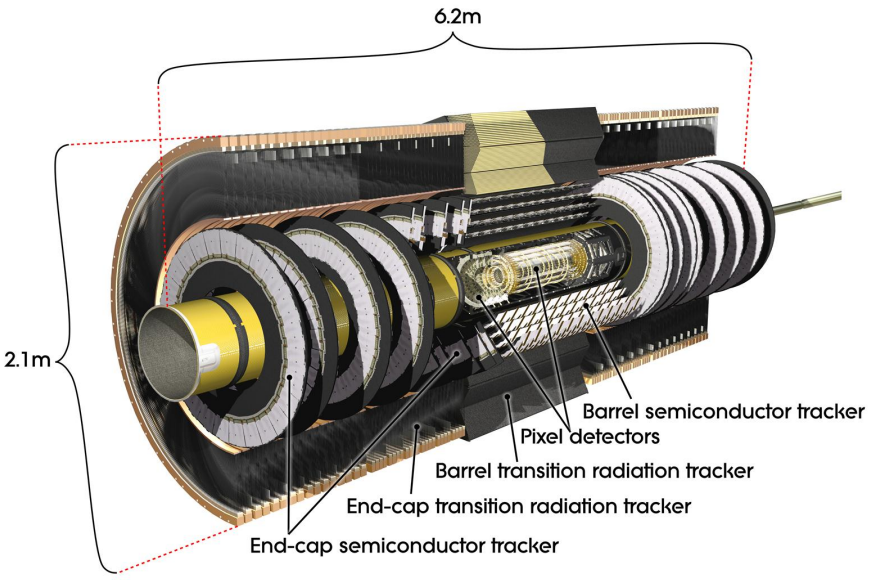

Figure 3.3: Overview of the Inner Detector of the ATLAS experiment. [32]

Figure 3.3 shows an overview of the Inner Detector. It covers the region $|\eta| < 2.5$ and consists of three major components. The innermost of these is the Pixel Detector. In Run-1 it consisted of three layers of silicon pixel detectors, in both the barrel and the endcap. The pixel size is $50 \times 400\,\mu\text{m}$ in $\phi \times z$ (R) in the barrel

(endcap). This results in a position resolution of 10 µm in the $R - \phi$ plane and 115 µm along the z-axis. In total, the three original pixel detector layers contain 80.4 million readout channels.

In LS1, an additional fourth pixel layer has been installed in the centre of the Pixel Detector to counteract ageing effects in the original innermost layer. This so called *Insertable B-Layer* [40] adds a further 12 million channels to the readout volume. A reduced pixel size of 50×250 µm and a shorter distance to the beam pipe - 3.27 instead of 5.05 cm - also lead to a better vertex resolution.

Outside the Pixel Detector, the Semiconductor Tracker is installed. It consists of four barrel layers and nine layers in either of the endcaps. Each layer in turn consists of two layers of silicon strip detectors. In the barrel, one layer is arranged parallel to the z-axis and the second layer is tilted by 40 mrad with respect to the first. In this way, both the ϕ and the z coordinate can be measured precisely. In the endcap, one layer is arranged radially, and the two layers are tilted by the same angle. The Semiconductor Tracker provides a resolution of 17 µm in the $R-\phi$ and 580 µm along the z (R) axis in the barrel (endcap) and contributes 6.3 million channels to the readout volume.

The part of the Inner Detector furthest from the interaction point is the Transition Radiation Tracker. It consists of drift tubes with a diameter of 4 mm. The tubes are filled with a xenon-based gas mixture which is ionised by traversing particles. The produced charges are collected on gold-plated tungsten wires that serve as anodes. The barrel section consists of 73 layers of 144 cm long tubes. They are arranged parallel to the z-axis with an average distance of 7 mm. Each endcap section is eqipped with 160 layers of radially oriented tubes of 37 cm length. Due to this setup, the Transition Radiation Tracker is only sensitive to the $R - \phi$ plane, where it reaches a resolution of 130 µm.

In addition to the ionisation caused by charged particles, the Transition Radiation Tracker also detects transition radiation photons. These are emitted by charged particles when they pass the boundary between two materials with different dielectric constants. The photons are absorbed in the gas, which induces large signals in the tubes. The energy of the emitted photons is proportional to the Lorentz factor, $\gamma = E/m$. The Transisition Radiation Tracker thus contributes to the particle identification by providing information on the mass of charged particles. In particular, it allows to distinguish electrons from heavier particles like charged pions.

3.2.3 Calorimetry

The calorimeter absorbs incident particles and determines the energy and position of the resulting particle showers.

Figure 3.4 shows an overview of the ATLAS calorimeter. It consists of two layers, each of which is further subdivided into a barrel and two endcap regions, and additional forward calorimetry. The inner layer is the electromagnetic calorimeter. It is optimised for the measurement of electrons and photons, which produce electromagnetic showers. The outer layer is the hadronic calorimeter. It contains

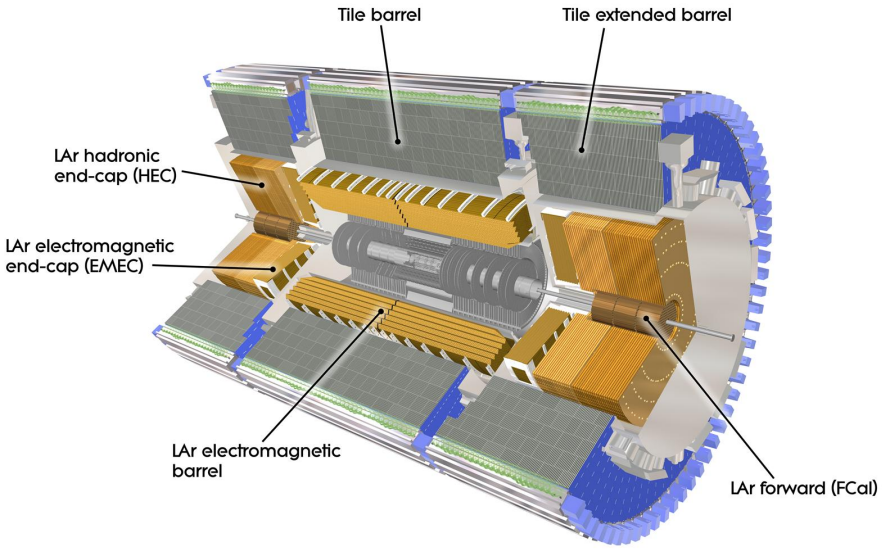

Figure 3.4: Overview of the ATLAS Calorimeter. [32]

hadronic showers (jets), that typically extend beyond the electromagnetic calorimeter, and contributes to the measurement of the missing transverse energy. The forward calorimeters also contribute to the hadronic energy measurement. In total, the ATLAS calorimeter covers the region $|\eta| < 4.9$.

All ATLAS calorimeters are sampling calorimeters, constructed from alternating layers of absorbers and active material. The absorber layers are constructed from dense material, so that the incident particles interact with it more strongly. This way, the particle showers are contained more effectively in the calorimeter, allowing for a more compact and cheaper construction of the detector. The deposited energy of the particles is measured in the active layers.

Only the part of the total shower energy that is deposited in the active medium is detected. As the shower shape varies statistically, the fraction of energy that is deposited in the active layers also fluctuates. This reduces the energy resolution of sampling calorimeters compared to homogenous calorimeters, where the full detector consists of active material.

The energy resolution achieved by the different layers of the ATLAS calorimeter is shown in table 3.1.

Calorimeter	σ_E/E
Electromagnetic Calorimeter	$10\%/\sqrt{E} \oplus 0.7\%$
Hadronic Calorimeter	$50\%/\sqrt{E} \oplus 3\%$
Forward Calorimeter	$100\%/\sqrt{E} \oplus 10\%$

Table 3.1: Energy resolution of the different ATLAS calorimeter layers. [32]

The different calorimeter partitions are described in the following.

Electromagnetic Barrel Calorimeter

The Electromagnetic Barrel Calorimeter (EMB) consists of two half barrels, one covering the positive and one covering the negative z direction. With a total length of about 6.4 m and a thickness of 0.6 m, it covers the region $|\eta| < 1.475$. Depending on $|\eta|$, the thickness corresponds to 22 to 33 radiation lengths[5] X_0, as seen from the interaction point.

The EMB utilises lead plates as absorbers, that are formed in a radially oriented accordion shape. The volume between the plates is filled with liquid argon (LAr) as active medium that is ionised by the secondary particles produced in the electromagnetic showers. Kapton electrodes are placed in the centre between two layers of lead to collect ionisation charges. They follow the accordion shape of the absorber plates at a constant distance of 2.1 mm. This leads to a typical drift time of about 450 ns.

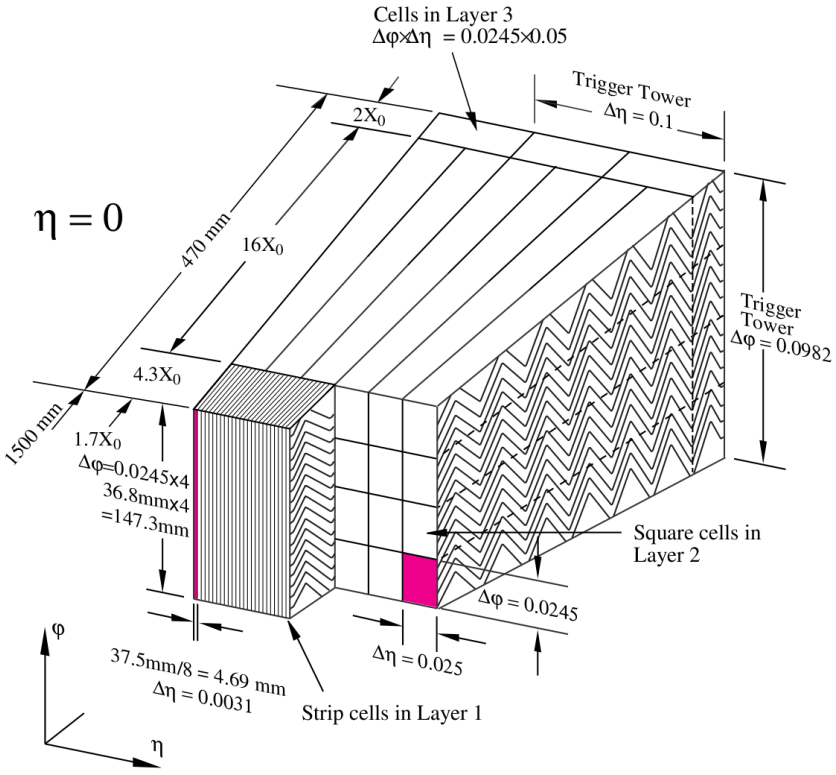

Figure 3.5: Structure of the EMB modules. [32]

Each half barrel consists of 16 modules. Figure 3.5 illustrates the structure of the EMB modules. They consist of three layers of different granularity. The first layer

[5]The radiation length is defined as the average distance after which the energy of an incident particle is reduced by a factor $\frac{1}{e}$ due to electromagnetic interactions with the detector material.

consists of strip cells with a size of (0.003×0.1) in $(\Delta\eta \times \Delta\phi)$, with a thickness of 4.3 X_0. The fine segmentation in η allows for a good discrimination between photons and neutral pions. The second layer is intended to absorb most of the energy deposited by incident particles and to provide a good position measurement. This is achieved by using square cells with a size of (0.025×0.0245) in $(\Delta\eta \times \Delta\phi)$ and a depth of 16 X_0. The third layer consists of larger cells, (0.05×0.0245) in $(\Delta\eta \times \Delta\phi)$, that measure the tails of the electromagnetic showers. It has a thickness of 2 X_0.

An additional Presampler layer is installed in front of the EMB modules. It is used to correct for the energy loss in the inactive material between the interaction point and the calorimeter, i.e. the ID, the cryostat wall and the solenoid magnet. It consists of 64 cells of 1.1 cm thickness.

Electromagnetic Endcap Calorimeter

The two Electromagnetic Endcap Calorimeters (EMECs) are placed symetrically around the z-axis on either side of the EMB. Each EMEC extends from 0.33 m to 2.01 m from the beam pipe and has a thickness of 0.63 m. This corresponds to a coverage of $|\eta| < 3.2$ and a depth of 24 to 38 X_0, depending on η.

The EMEC uses the same LAr-lead sampling technology with kapton electrodes as the EMB. The accordion geometry is arranged parallel to the z-axis.

Each EMEC consists of two co-axial wheels: the outer wheel covers the region $1.375 < |\eta| < 2.5$, while the inner wheel covers $2.5 < |\eta| < 3.2$. Each of the wheels is in turn subdivided into eight modules. These are organised similar to the EMB, with several layers if differing granularity. For the region $1.5 < |\eta| < 2.5$, three layers with the same granularity as in the EMB are used. For $1.375 < |\eta| < 1.5$ and $2.5 < |\eta| < 3.2$ only two layers are installed, which have a coarser granularity of up to (0.1×0.1) in $(\Delta\eta \times \Delta\phi)$.

Similar to the EMB, the EMEC is equipped with a presampler that covers the region $1.5 < |\eta| < 1.8$. It is thinner compared to its EMB counterpart, with a thickness of only 0.5 cm.

Hadronic Tile Calorimeter

The hadronic Tile calorimeter consists of a 5.8 m long barrel (TileLB) section that is placed symetrically around the interaction point and a 2.6 m long extension (TileEB) along both the positive and the negative z direction. The thickness of about 1.97 m corresponds to 7.4 interaction lengths[6] λ in radial direction. The Tile calorimeter covers the region $|\eta| < 1.7$.

Each of the parts of the Tile calorimeter is divided into 64 modules. The design of one module is shown in Figure 3.6. It is built from stainless steel plates, that

[6]The interaction length is the average distance after which the energy of a particle traversing the material is reduced by a factor $\frac{1}{e}$. Contrary to the similarly defined radiation length X_0, the interaction length includes both electromagnetic and strong interactions.

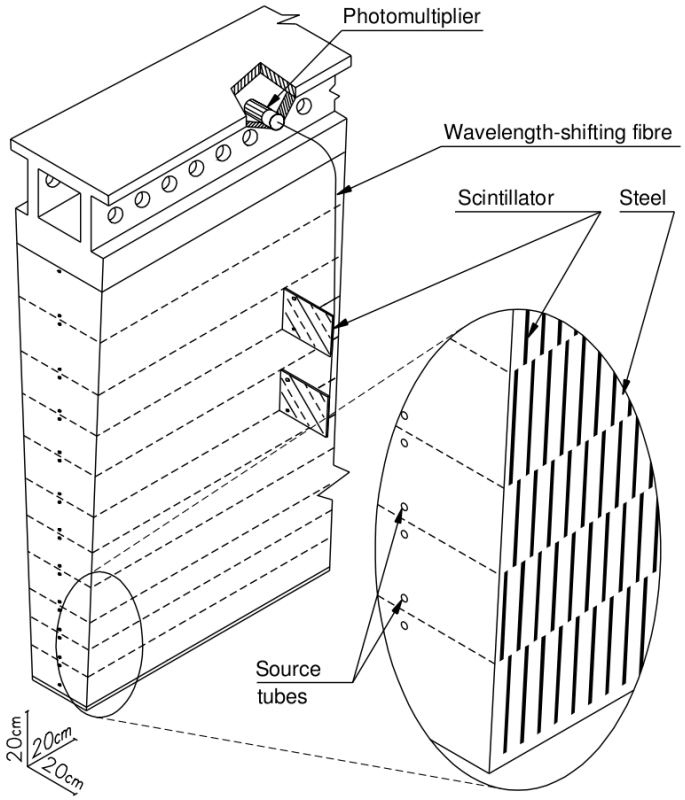

Figure 3.6: Structure of the Tile modules. [32]

function as absorbers, and scintillating tiles as active medium. The thickness of the steel plates is 5 mm, the tiles are 3 mm thick. The scintillation light produced by incident particles is read out by wavelength-shifting fibres that are attached to photomultiplier tubes. The module is organised into three layers of cells, that are formed by grouping the readout fibres of multiple tiles on the same photomultiplier tube. The granularity in $(\Delta\eta \times \Delta\phi)$ is (0.1×0.1) in the inner two layers, and (0.2×0.2) in the outer layer.

Hadronic Endcap Calorimeter

A Hadronic Endcap Calorimeter (HEC) is installed on both sides of the detector. Each HEC consists of two wheels with an outer radius of 2.03 m and a combined depth of 12 λ. They cover the region $1.5 < |\eta| < 3.2$.

The HEC is constructed from copper plate absorbers and uses LAr as active medium. The plates have a thickness of 25 mm in the wheel closer to the interaction point and 50 mm in the wheel further out. The 8.5 mm gap between the plates is equipped with three electrodes, dividing it into four equally sized drift zones of 1.8 mm. The central electrode is used to read out the ionisation charge, while the outer two supply the high voltage. The drift time amounts to 430 ns. The cell granularity is given by the electrode segmentation. It is (0.1×0.1) in $(\Delta\eta \times \Delta\phi)$

up to $|\eta| < 2.5$ and (0.2×0.2) for larger $|\eta|$.

Forward Calorimeter

The Forward Calorimeters (FCals) are installed between the beam pipe and the endcaps. The FCal consists of three modules of 45 cm thickness on each side of the detector. In total, the FCal covers the region $3.1 < |\eta| < 4.9$ and a depth of about 10 λ.

The first module (FCal1) is optimised for electromagnetic interactions. It consists of copper plate absorbers that are penetrated by equidistant holes. The electrodes are placed within these holes and take the form of cylindrical pipes that run parallel to the z-axis. Each pipe is filled with LAr and contains an axially oriented copper staff. The drift distance between the pipe wall and the staff is about 0.27 cm.

The second and third FCal modules (FCal2 and FCal3) function as hadronic calorimeters. Each of them consists of two copper plates of 2.35 cm thickness, between which the electrodes are arranged. These are very similar to the ones used in FCal1, the only differences being the material for the central staff - tungsten - and the drift distance - 0.38 cm in FCal2 and 0.51 cm in FCal3. The remaining volume between electrodes and the plates is filled with tungsten.

Due to the relatively small drift distance, the drift time in the FCal is much smaller than in the other LAr calorimeters used in ATLAS. It is the shortest in FCal1, where it is of the order of 60 ns.

3.2.4 Muon Spectrometer

The muon system detects charged particles that are not stopped in the calorimeters. It consists of different gaseous detectors that are divided into two types: highly position sensitive detectors in order to measure the tracks of charged particles in the toroid magnetic field; and very fast detectors used for the trigger. In total, the Muon Spectrometer covers the region $|\eta| < 2.7$.

Figure 3.7 shows an overview of the ATLAS muon systems. In the barrel region, three cylindrical layers of Monitored Drift Tube chambers placed in the toroid magnets are used to measure the particle tracks. Each chamber contains three to eight layers of 30 mm diameter drift tubes. The tubes are filled with an argon-based gas mixture, that is ionised by traversing charged particles. In each tube, the charge is collected by a centrally positioned, coaxial tungsten-rhenium wire of 50 µm diameter.

In each endcap, four layers of Monitored Drift Tubes cover the region up to $|\eta| < 2$. in the most forward region, $2 < |\eta| < 2.7$, the inner-most layer is instead equipped with Cathode Strip Chambers. These are multi-wire proportional chambers with two cathode layers, each of which is segmented into strips. In one cathode, the strips are oriented parallel to the wires, while they are orthogonal on the other plate. This allows for a high position resolution in two dimensions. The main reason to use this type of detector in the forward region is the higher granularity compared to the drift

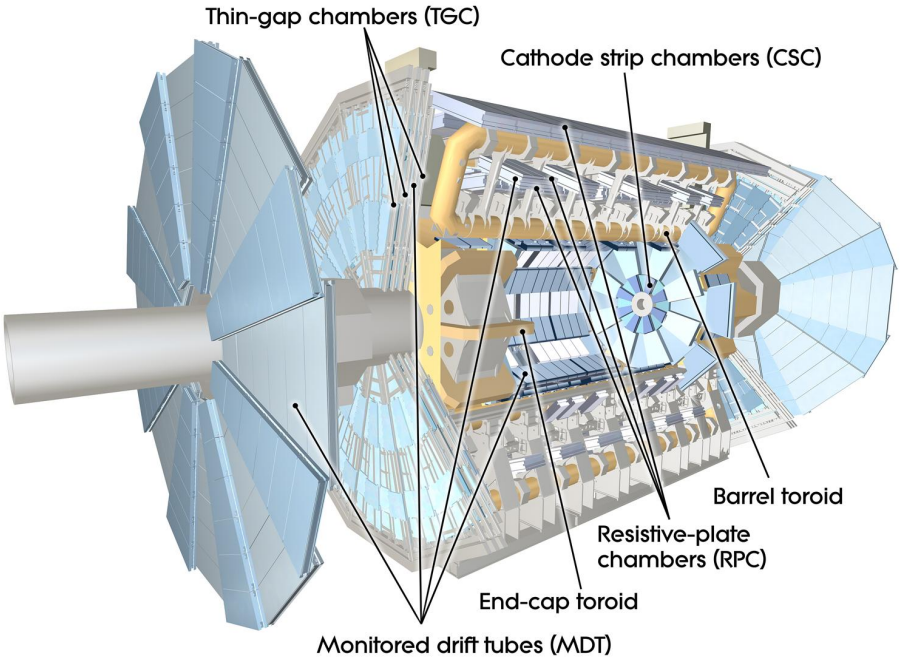

Figure 3.7: Overview of the ATLAS Muon Spectrometer. [32]

tubes. This alleviates the challenge imposed by the higher rates and backgrounds at larger $|\eta|$.

The Monitored Drift Tubes attain a position resolution of about 35 µm per chamber, with a total drift time of 700 ns. The Cathode Strip Chambers achieve a position resolution of 40 µm orthogonal to the magnetic field and 5 mm parallel to it. Their timing resolution if 7 ns, with a drift time of 40 ns.

The muon trigger information for the barrel is obtained from Resistive Plate Chambers. These consist of two parallel resistive plates, placed 2 mm apart, with a high voltage applied between them. The volume between the plates is filled with gas that is ionised when a particle passes through it. The electric field then causes the production of an avalanche from these primary charges. The signal is read out via metal strips that are placed on the outside of the chamber. They achieve a position resolution of 10 mm and a timing resolution of 1.4 ns.

In the endcaps, Thin Gap Chambers are used for the trigger. These are also multi-wire proportional chambers, but the distance between the wires, 1.8 mm, is larger than the distance from the wires to the plates, 1.4 mm. This leads to short drift times and a good timing resolution of 4 ns. The position resolution is 2-6 mm in radial direction and 3-7 mm in ϕ.

3.2.5 Magnets

The ATLAS detector contains four large superconducting magnets. They provide the magnetic fields that are required by the tracking and muon systems.

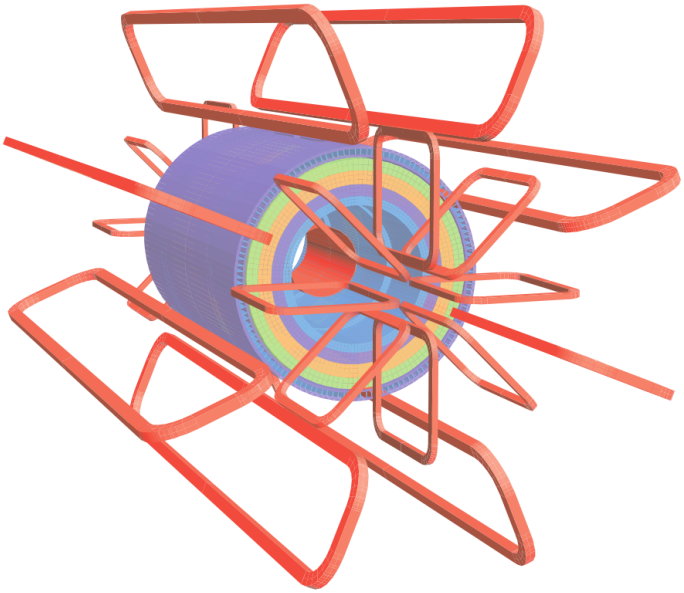

Figure 3.8: Overview of the ATLAS magnet systems. The different magnets are shown in red, with the solenoid in the centre and the toroids on the outside of the detector. [32]

Figure 3.8 shows an overview of the ATLAS magnets. In the centre of the barrel part of the detector, a solenoid magnet is installed between the Inner Detector and the Electromagnetic Calorimeter. It consists of a single-layer NbTi coil that produces a 2 T field in the Inner Detector that is oriented parallel to the z-axis. The solenoid extends over a length of 5.8 m. As it is positioned before the Electromagnetic Calorimeter, the material contribution of the solenoid has to be minimised to reduce the probability of particles starting to shower before the calorimeter. To achieve this, the solenoid is formed of a single coil and is installed in the same cryostat as the electromagnetic calorimeter. This results in a total of only 0.66 radiation lengths of material in front of the calorimeter.

One barrel toroid and two endcap toroid magnets make up the remainder of the ATLAS magnet systems. All three toroids consist of eight coils each, which are placed symetrically around the beam pipe, producing a cylindrical magnetic field in the Muon Spectrometer. The field strength is 0.5 T in the barrel and 1 T in the endcaps.

3.2.6 Further Subdetectors

Next to the major components described above, ATLAS contains several additional subdetectors for various specialised purposes.

- The Minimum-Bias Trigger Scintillator (MBTS) consists of two rings of 2 cm thick scintillating tiles on each side of the detector. It is used to detect and trigger on non-diffractive pp interactions.

- The Luminosity Measurement with Cherenkov Integrated Detector (LUCID) monitors the luminosity seen by the ATLAS experiment. It consists of gas filled aluminium tubes that are installed close to the beam pipe in the muon endcap region. Particles traversing the gas produce Cherenkov radiation that is detected by photomultiplier tubes.

- Three forward detectors are installed at larger distance from the interaction point to analyse *pp* interactions in the very forward region: The Zero Degree Calorimeter (ZDC) is situated at 140 m from the interaction point and is optimised to detect neutrons in heavy ion collisions; Absolute Luminosity for ATLAS (ALFA) is located at 240 m from the interaction point and measures elastic *pp* scattering; ATLAS Forward Protons (AFP) has been newly installed during Run 2 at 210 m distance from the interaction point to complement the ALFA program by measuring diffractive processes.

3.3 ATLAS Detector Control System

The *Detector Control System* (DCS) collects and monitors environmental data from the different ATLAS subsystems. This includes the applied voltages, the temperature of both devices and the environment, gas pressure in the detector etc. . These quantities are continously compared against thresholds which define the safe operating conditions of the different detectors. If a deviation beyond a threshold is detected, appropriate warnings or alarms are raised. For strong deviations, automatic countermeasures are induced, in the worst case shutting down the affected component.

3.4 ATLAS Trigger and Data Aquisition System

The signals induced in the detectors are stored in dedicated data buffers in the detector front-end electronics. The ATLAS Trigger and Data Aquisition (DAQ) systems, commonly referred to as TDAQ, extract the event data from the front-end and transmit it to mass storage.

As the event rate of 40 MHz, coupled with an event size of several MB, is far too large to store all events, the trigger is used to reduce the rate by selecting only those events that are of potential interest to physics analysis. If an event is selected by the trigger, the DAQ system collects the data from all ATLAS subdetectors, assembles it into a common event format and stores it on disk.

An overview of the TDAQ system is shown in Figure 3.9. It is broadly divided into four parts: the Readout of the Detector, the Level-1 (L1) Trigger, the High Level Trigger (HLT) and the Data Flow system.

The L1 trigger receives reduced granularity data from the calorimeters and the muon systems at 40 MHz. If an event is selected, an L1 Accept (L1A) signal is sent to all ATLAS subdetectors, which initiates the readout of the detector data.

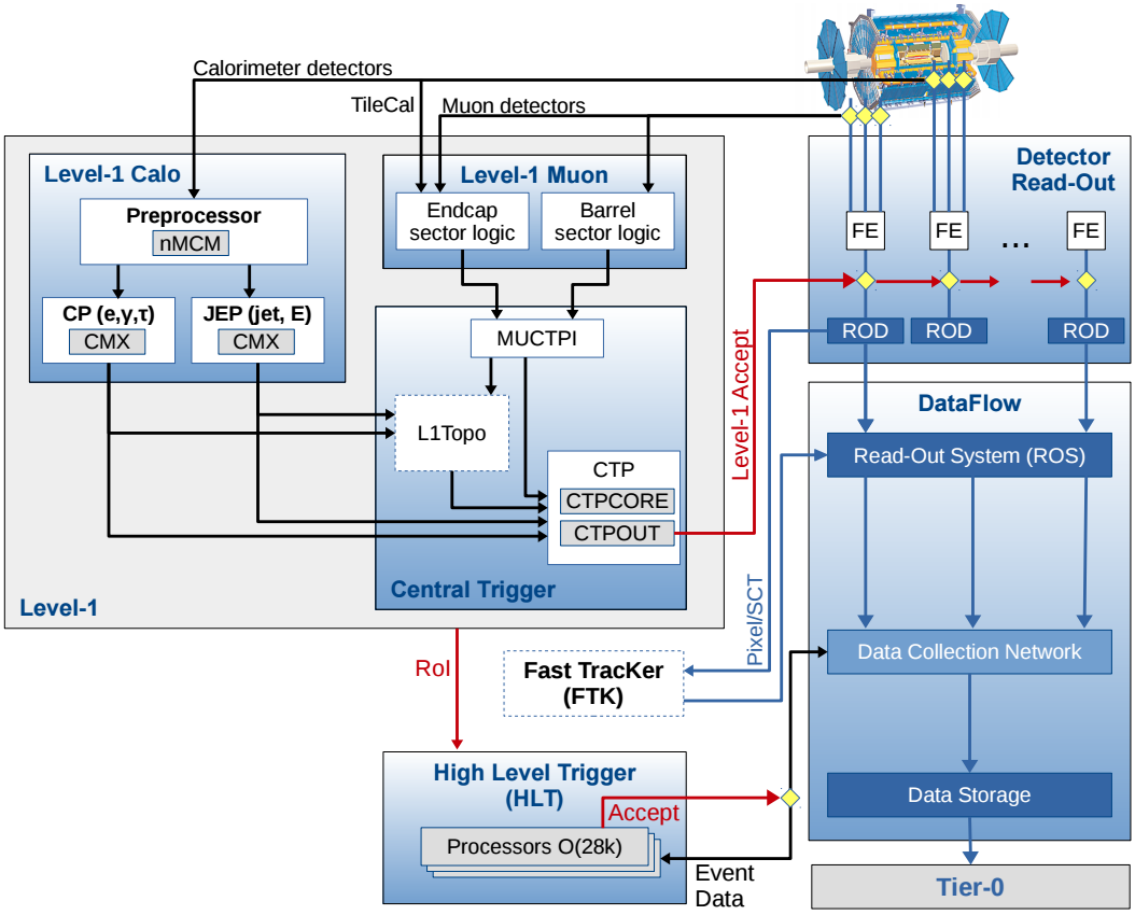

Figure 3.9: Overview of the ATLAS TDAQ system. At the time of writing, the L1Topo and FTK system are in a commissioning state and are planned to be used in 2017 data taking. [41]

The time between the transmission of data off the detector and the reception of the L1A is limited to 2.5 µs by the size of the buffers in the detector front end. Due to this constraint, the L1 trigger is implemented entirely in custom hardware. The target L1A rate is 100 kHz, corresponding to a rate reduction by a factor of 400.

Upon receiving the L1A signal, the full granularity data is read out from the detector front-end to a further buffering stage in the DAQ sytem. The data is then analysed by the HLT, using also Regions of Interest (RoIs) received from the L1 trigger that identify the locations of particle candidates found by that system. If an event is selected by the HLT, it is fully reconstructed and transmitted to mass storage. The target event rate at the HLT output is 1 kHz.

The different components of the TDAQ system are briefly discussed in the following.

3.4.1 Level-1 Trigger

The L1 trigger analyses reduced granularity data from the calorimeter and muon detectors to identify events that contain highly energetic muons, electrons, photons, jets and hadronically decaying tau leptons as well as large missing transverse energy (E_T^{miss}) and total transverse energy (E_T^{tot}).

As seen in Figure 3.9, it is divided into several major components: the L1 Calorimeter Trigger (L1Calo), the L1 Muon Trigger (L1Muon), the L1 Topological Trigger (L1Topo) and the Central Trigger Processor (CTP).

Level-1 Calorimeter Trigger

The L1Calo trigger is a pipelined processing system that analyses about 7200 analogue signals from the calorimeters. It digitises and synchronises these signals, calibrates them to E_T and identifies particle candidates in the resulting E_T distribution across the calorimeter. The particle candidates are transmitted to the L1Topo system, and the multiplicity of particle candidates above certain E_T thresholds is forwarded to the CTP. If a given event is selected by the L1 trigger, the L1Calo trigger provides event data to the readout system and supplies RoI information to the HLT. A detailed description of L1Calo is given in Chapter 4.

Level-1 Muon Trigger

The L1Muon trigger analyses data from the trigger chambers, i.e. the RPCs in the barrel and the TGCs in the endcaps. It detects muon candidates and estimates their p_T. The measured p_T is compared to six programmable thresholds, and the number of muons passing each threshold is provided to the CTP. Upon receiving an L1A from the CTP, the L1Muon trigger also supplies RoI information to the HLT and event data to the DAQ.

The L1Muon algorithm uses coincidences between detector stations in different layers to detect traversing particles. For each p_T threshold, pairs or triplets of coinciding trigger chambers are defined that indicate a muon candidate above the threshold. These coincidence windows are derived by assuming a muon coming from the nominal interaction point and passing a station in one of the detector layers, referred to as the pivot layer. In the barrel, this is the middle layer of RPCs, while in the endcap it is the back layer of TGCs. Depending on which trigger chamber in the pivot plane is passed by the muon, one or two further stations have to be hit as well in order to satisfy the threshold condition.

Level-1 Topological Trigger

The L1Topo trigger receives particle candidates from the L1Calo and L1Muon system and applies selection criteria that involve the topology of the event. Among others, this includes cuts on the angular separation between particle candidates or the invariant mass of systems of particles.

Central Trigger Processor

The CTP receives particle multiplicities from the L1Calo and L1Muon systems as well as the results of the L1Topo trigger. It also receives trigger inputs from the forward detectors and from the MBTS. The inputs are checked against a programmable L1 trigger menu, that contains a list of trigger items. Each trigger item defines certain conditions that have to be satisfied in order to trigger the event. For example, the `L1_J400` trigger item is satisfied if at least one jet above an E_T threshold of 400 GeV is detected. If at least one item in the trigger menu is met for a given event, the event is accepted and the CTP issues an L1A. In total, the CTP can form 512 possible trigger items from up to 320 inputs.

In order to keep the rates, expecially of low energy trigger items, under control, the CTP applies programmable prescales to each trigger item. The prescale acts as a reduction factor for the trigger rate: with a prescale of n, only every n'th event satisfying the prescaled trigger item is accepted. The prescales can be adapted online during a data taking run, allowing to adjust the L1 trigger output as the luminosity decreases.

As the application of a prescale introduces a statistical bias to the trigger selection, *unprescaled* trigger items are of special importance. For this reason, the *lowest unprescaled trigger items* are among the most important ones included in the trigger menu, and usually used as primary triggers for physics analyses.

The CTP also vetoes events according to two deadtime criteria. These are applied to prevent the readout buffers in the detector front-end from overflowing. The simple deadtime vetoes all triggers in a programmable number of BCs after each L1A. The complex deadtime applies a leaky bucket algorithm to model the amount of events stored in the front-end buffers. Each bucket is defined by two programmable numbers: the number of events that can be stored in one derandomiser at the same time and the number of BCs required to transmit one event out of the derandomiser. The bucket contents are incremented by one at each L1A, and reduced by one whenever the programmed number of BCs has passed. As long as the content of the bucket is equal to the programmed limit, all L1As are vetoed. The CTP supports up to four buckets in parallel, providing enough flexibility to cover the wide range of detector readout implementations.

Lastly, the CTP distributes several LHC protocol signals to the ATLAS subdetectors using the Timing and Trigger Control (TTC) system. This inludes the 40.08 MHz BC clock and the ORBIT signal, which indicates the beginning of a new LHC revolution. The latter is also referred to as the Bunch Counter Reset (BCR), as it is used to reset the internal BC number counters of various subsystems. A further TTC signal is the Event Counter Reset (ECR), that is used to reset the internal event counters in all ATLAS subsystems.

3.4.2 High Level Trigger and Data Aquisition

If the L1 trigger selects an event, subdetector-specific Readout Drivers (RODs) collect the data and transmit it to the Readout System (ROS), a network of commercial PCs. From there, it is made available to the HLT via the Data Collection Network.

The HLT is a software based trigger implemented in a farm of commercial PCs. It operates in two stages, with a partial event reconstruction in the first stage to reject the majority of the events and a second stage for more detailed selections. In the first stage, full granularity event data for the RoIs found by the L1 trigger is used to reconstruct the particle candidates with a higher precision . If an event passes this stage, the event is fully reconstructed by also retrieving the remaining data. At this stage, more precise selections are applied, that are similar to those used in offline data analysis. Events that also pass this stage are then stored to disk.

3.4.3 TDAQ Upgrade in LS1

The ATLAS TDAQ system has been upgraded during LS1 to adapt to the increased centre-of-mass energy and luminosity in Run-2 [42]. Both of these changes increase the rate of many types of physics signatures that are detected by the TDAQ system. In addition, the increased pile-up has a higher probability of *faking* these signatures, leading to an increase in fake rates as well.

The goal of the upgrade is to increase the performance of the TDAQ system such that the high trigger efficiency and low thresholds for unprescaled trigger items achieved in Run-1 are maintained despite these challenges. While the maximum trigger rate is increased from 75 kHz to 100 kHz, this is not sufficient to operate the same Run-1 system.

The upgrade of the different TDAQ systems is summarized in the following:

- The **L1Topo** trigger was added to the L1 trigger system. By changing form single-object E_T thresholds to multi-object topological selections, the rates of triggers required for several physics signatures can be strongly reduced.

- The upgrade of the L1Calo trigger is described in more detail in the following chapters. In summary, the central processing units of the PreProcessor subsystem was replaced by an upgraded version to enhance the signal calibration and the interface to the central trigger was upgraded to support the new L1Topo system.

- The connectivity of the L1Muon system to the central trigger was upgraded to support the L1Topo system.

- The CTP was similarly upgraded to include the additional inputs provided by the L1Topo system. In total, the number of L1 trigger items was increased from 256 to 512.

- During Run-1, the HLT consisted of two separate computer farms - the RoI based Level-2 trigger and the Event Filter that uses fully reconstrcuted events. These two systems were merged for Run-2, to allow for optimal resource sharing and reduced complexity in the connectivity between the systems.
- The readout system and data collection network were upgraded to improve the detector readout rate and the output rate to the HLT and the data storage system.

3.5 ATLAS Operation

The operation of the ATLAS detector is controlled by the TDAQ and DCS systems. A typical *run* of the experiment starts before protons are colliding in the LHC. It begins with the configuration of all subdetectors and their readout devices with the parameters appropriate to the LHC filling conditions. In this stage, the majority of trigger items are deactivated, with only a minimal set of triggers that are required for monitoring the detector condition being active. Once the beams in the LHC reach their stable collision configuration, the physics triggers are activated and the DAQ system begins to gather the recorded data.

Besides these *data taking runs*, ATLAS also performs regular *calibration runs*. These are used to gather data that is required to calibrate the detector and readout electronics. The are usually performed in interfill periods, with no protons in the LHC.

Each run is divided into *luminosity blocks*, or lumi blocks. These are periods of time during which the operating conditions like the detector configuration, the instantaneous luminosity or the trigger prescales are constant. A lumi block typically lasts for 60 s, but can be shorter if e.g. adjustments of the configuration are performed.

4 The ATLAS Level-1 Calorimeter Trigger

The ATLAS L1Calo trigger [43] is a pipelined, digital system of custom electronics that is operated at 40.08 MHz. It receives analogue input from 7168 *trigger towers*, each of which is comprised of calorimeter cells in a region of the ATLAS calorimeter, with a typical size of (0.1×0.1) in $(\Delta\eta \times \Delta\phi)$. For each BC, these input signals are analysed to identify candidates for electrons, photons, τ particles and jets in the calorimeters, and to compute global energy sums like the missing transverse energy E_T^{miss} or the total transverse energy $\sum E_T$. The identified particle candidates are forwarded to the L1Topo system, and the number of candidates that pass programmable thresholds is transmitted to the CTP.

The latency of the L1Calo processing per event is below 1 µs. The signal transmission delay from the detector to L1Calo and further on to the CTP is of the same order, resulting in a total latency of about 2.1 µs. It is thus within the bound of 2.5 µs required from the L1 trigger.

4.1 Functional Overview

Figure 4.1 shows an overview of the L1Calo trigger. It consists of three subsystems:

- The *PreProcessor* digitises the analogue input signals and assigns, at each BC, a digital E_T value to each trigger tower.
- The *Cluster Processor* (CP) identifies candidates for electrons, photons and hadronically decaying τ particles in the η-ϕ distribution of the E_T provided by the PreProcessor.
- The *Jet/Energy Processor* (JEP) similarly identifies jet candidates and computes the E_T^{miss} and E_T^{sum} variables.

The particle candidates identified by the CP and JEP systems are collected in merging modules (CMX). These perform the transmission of the identified *trigger objects* (TOBs) to the L1Topo system, apply E_T thresholds and send the *multiplicity* of passing particles to the CTP.

Upon receiving an L1A signal, the L1Calo trigger transmits event data to the DAQ system, which contains the digitised trigger tower signal as well as results from various stages of processing. This data allows to monitor the performance of the system and to verify its functionality. As described in Section 3.4, the L1Calo trigger also transmits RoI information regarding the identified particle candidates to the HLT.

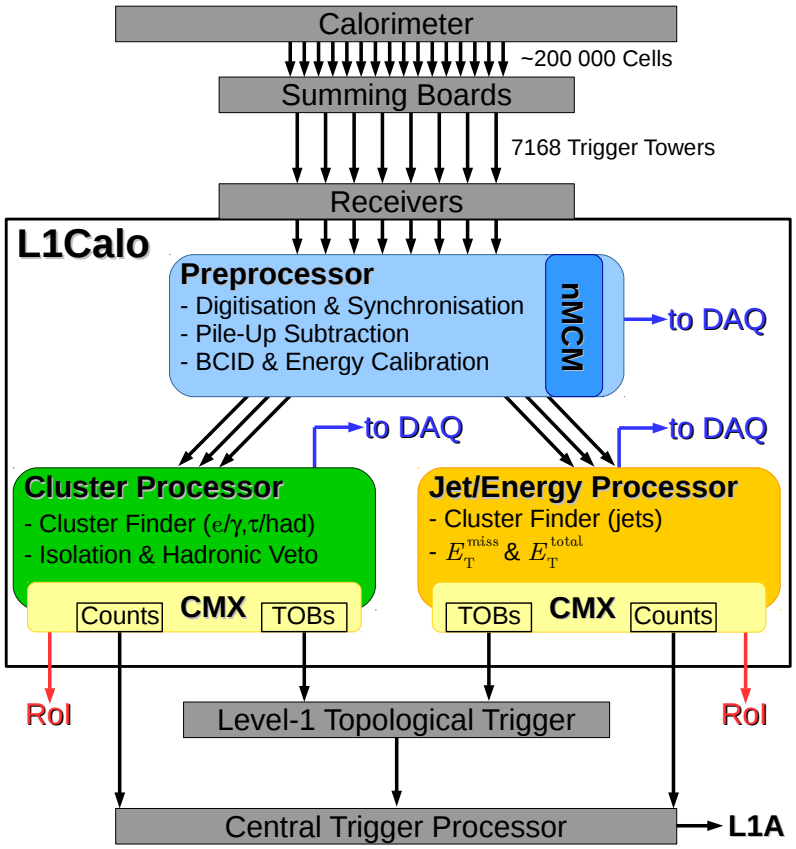

Figure 4.1: Overview of the ATLAS L1Calo trigger.

The following chapters describe the formation of the trigger tower signals and the functionality and implementation of the different L1Calo subsystems.

4.2 Analogue Input Signals

The ATLAS calorimeters consist of a total of about 200000 individual cells. Processing the calorimeter signals at this full, cell-level granularity in the L1Calo trigger is not feasible given the required amount of cabling as well as cost and space limitations. The L1Calo trigger thus receives input signals at a reduced granularity.

4.2.1 Input Signal Path

The L1Calo input signals are formed in the front-end electronics of the calorimeters, independently for the electromagnetic and the hadronic layer. They are formed as analogue sums of signals from calorimeter cells in regions of η and ϕ that are projective towards the nominal interaction point and extend over the whole depth

of the calorimeter layer[1].

The typical size of these trigger towers is (0.1×0.1) in $(\Delta\eta \times \Delta\phi)$ (see Figure 3.5 for an EMB trigger tower), with larger towers at $|\eta| > 2.5$. The number of cells contained in a trigger tower varies by detector partition, ranging from a few in the endcaps to up to 60 in the EMB. Figure 4.2 shows the trigger tower granularity for positive η values and one quadrant in ϕ in the electromagnetic calorimeter. The remaining ϕ quadrants follow the same structure, and the negative η side of the detector is covered symmetrically.

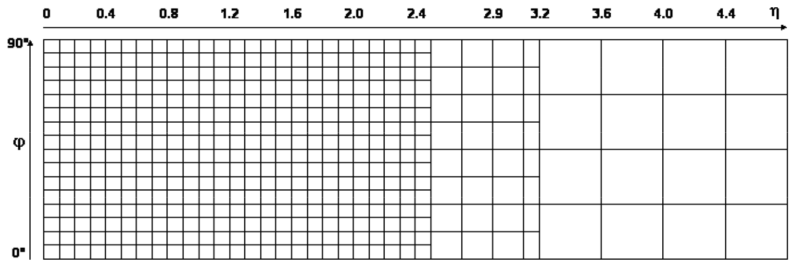

Figure 4.2: Trigger towers in the electromagnetic calorimeter for $\eta > 0$ and one quadrant in ϕ. [43]

For trigger towers in the EMB and EMEC, the amplitude of the signals is proportional to the total E_T deposited within them, with a guaranteed linear range up to 256 GeV. For the hadronic and forward calorimeters, the trigger tower signals are proportional to the raw energy instead.

The trigger tower signals are transmitted differentially over 616 16-way twisted-pair copper cables from the detector to the ATLAS electronics cavern, as shown in Figure 4.3. The cable path is optimised to achieve a short cable length, to minimise the transmission delay. The resulting cable lengths range from 30 m for the closest EMB tower up to 70 m for the furthest TileEB tower.

The cables from the Tile calorimeter also carry signals from the rear sampling layer, which are used by the L1Muon system to suppress backgrounds. These are disentangled from the trigger tower signals in the Tile Calorimeter Patch Panels (TCPPs). The Tile trigger towers are then sent to the Receiver system, which also receives the signals from the LAr calorimeters.

The main function of the Receiver system is to adjust the amplitude of the analogue signals using variable gain amplifiers (VGAs), each of which is controlled via a 12-bit DAC. In this way, the voltage scale is set as required by the L1Calo system, i.e. 10 mV per 1 GeV E_T. For trigger towers from the EMB and EMEC, the applied gain mostly compensates for signal attenuation in the long cables from the calorimeters. For trigger towers from the hadronic and forward calorimeters, the gain includes an additional factor $\sin\theta$ to convert the raw energy signals to E_T.

[1]Here, FCal1 is counted among the electromagnetic calorimeters, while FCal2&3 are counted as hadronic calorimeters.

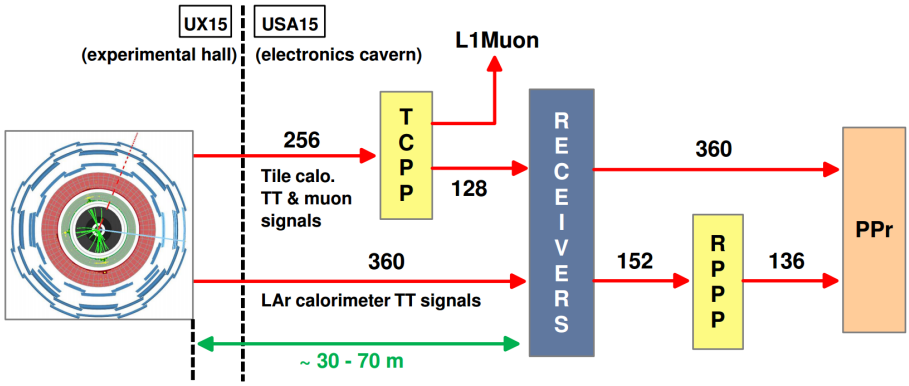

Figure 4.3: Path of the trigger tower signals from the calorimeter to the inputs of the PreProcessor. The numbers in the figure give the amount of cables installed along the respective paths. [44]

Additionally, the Receiver system re-orders the signals between the input and output cables to accomodate for differences in module coverage between the calorimeter front-end and the PreProcessor modules. At the Receiver output, each cable typically carries 16 trigger tower signals corresponding to a 0.4×0.4 area in η-ϕ.

The Receiver system also performs additional analogue summing in two calorimeter regions:

- In the overlap region between the EMB and EMEC partitions ($1.375 < |\eta| < 1.475$) cells from different calorimeter partitions lie in the same η-ϕ region. These signals have to be combined.
- In the hadronic FCal regions (FCal2&3) the η granularity is reduced by adding the signals from neighbouring channels (see Figure 4.7). This is done to achieve the same number of input channels for the electromagnetic and hadronic calorimeter layers.

Finally, the Receiver provides programmable monitoring outputs that provide access to the trigger tower signals. When the detector is not accessible, this is the only way to measure the analogue trigger tower signals.

The output signals of the Receiver are routed to the input of the PreProcessor. Some detector regions use a non-standard mapping of trigger towers to the PreProcessor inputs [45]. For these regions, additional *Receiver to PreProcessor Patch Panels* (RPPPs) are installed to re-order the signals accordingly.

4.2.2 Signal Shape

The basic shape of the trigger tower signals results from the shaping applied to the cell signals in the front-end electronics of the detectors. For the LAr calorimeters, the cell signal is of triangular shape, with a rise time below 1 ns and a decay time that corresponds to the drift time of the charges in the LAr gap, typically 450 ns. In

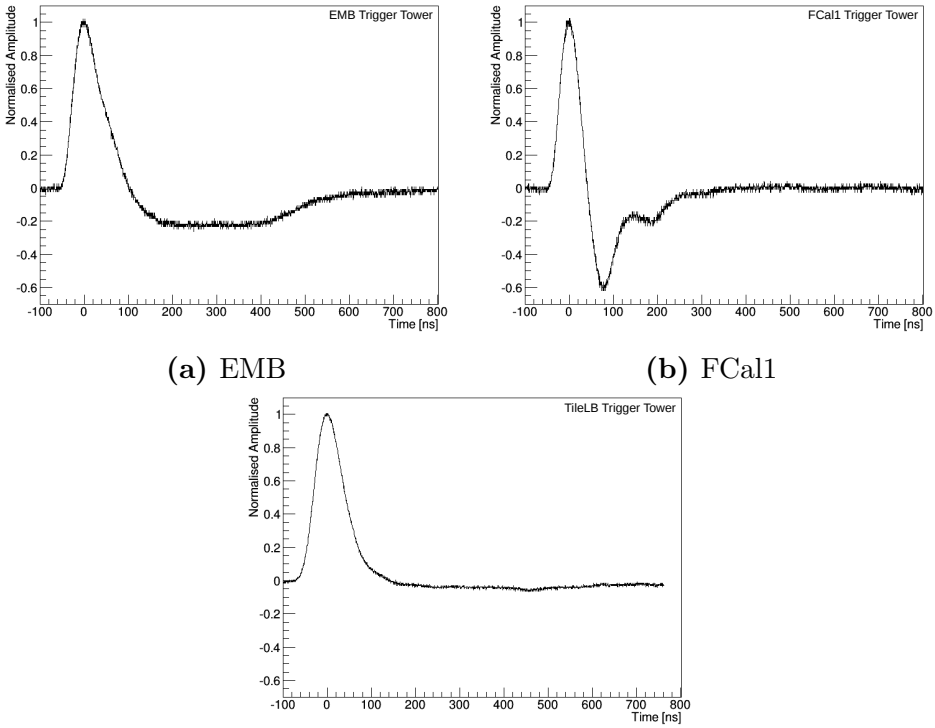

(a) EMB **(b)** FCal1

(c) Tile

Figure 4.4: Pulse shapes for trigger towers from the EMB (a), FCal1 (b) and Tile (c) calorimeters, as measured at the Receiver monitoring output.

the front-end electronics [46], the signal is amplified and shaped to a bipolar form, to optimise the signal to noise ratio. The positive part of the resulting signal extends over 5 BCs, with the negative undershoot covering the remainder of the duration of the unshaped signal. The shaped signal is sent to the summing circuit that builds the trigger tower signal.

For the Tile calorimeter, the cells produce short signal pulses with a rise time of about 5.5 ns and FWHM of around 17 ns [47]. These are shaped to unipolar[2] signals with FWHM of 50 ns in the electronics, resulting in a similar shape as the positive part of the LAr signals.

Figure 4.4 shows example pulse shapes of trigger towers in different calorimeter partitions. The pulses have been recorded with an oscilloscope attached to the monitoring output of the Receivers, and are normalised to the amplitude. Of note is the difference in the undershoot between the EMB (Figure 4.4 (a)) and the FCal1 (Figure 4.4 (b)). Due to the lower drift time, the FCal1 signal extends over a shorter time, resulting in a more strongly pronounced undershoot after the positive part of the bipolar pulse.

[2]The Tile signals actually do have an undershoot. Because it is very long and thus of low amplitude compared to the LAr undershoot, it is usually ignored in the signal processing.

4.3 PreProcessor

4.3.1 Functionality

The main functionality of the PreProcessor is to analyse the analogue trigger tower input signals to find significant energy depositions, assign them to the correct BC, calibrate them to E_T and send the E_T results to the subsequent object finding processors. This process is performed in several steps, as shown in Figure 4.5 and summarised below.

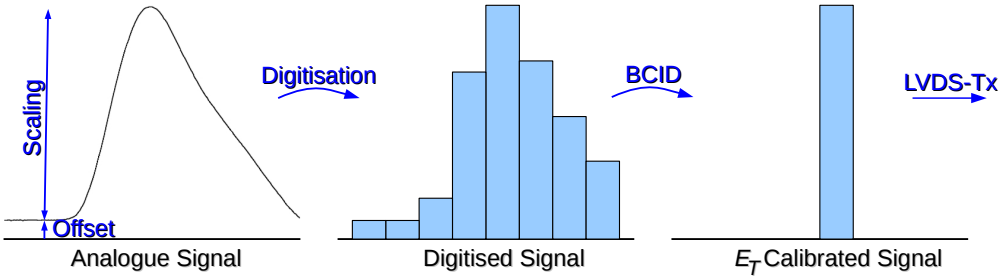

Figure 4.5: Functionality of the PreProcessor.

- The **analogue input signals are conditioned** to prepare them for digitisation.
- The signals are then **digitised** based on the LHC bunch-crossing frequency of 40.08 MHz. The digitisation is performed with a 10-bit resolution within the linear range of 0 to 256 GeV, such that the least significant bit corresponds to 250 MeV. The digitisation strobe is adjustable to ensure sampling as close as possible to the maximum of the pulse, to optimise the performance of the subsequent digital processing.
- Due to differences in time-of-flight of particles within the detector and different cable length along the signal transmission path, the trigger tower signals do not arrive at the PreProcessor input synchronously. To assure that signals originating from the same event are processed in the same BC, the PreProcessor **synchronises** the signals after digitisation.
- While the trigger tower pulses extend over many BCs, the PreProcessor has to associate the energy of the pulse to the single BC it originates from. This process is called **Bunch-Crossing Identification** (BCID). The BCID algorithms measure the pulse E_T and assign it to the peak BC of the pulse[3]. They also comprise noise filtering, baseline subtraction and a minimal energy cut,

[3]The BC in which the event actually occured lies at the start of the pulse and not the peak. Using the peak of the pulse instead is a design choice for the BCID algorithms. The delay induced by this is accounted for in the CTP processing.

to remove signals compatible with noise. The E_T is measured between 0 and 256 GeV in steps of 1 GeV, corresponding to an 8-bit BCID output value. The resolution is changed compared to the 10-bit input to reduce the data that has to be sent to the CP and JEP systems.

- The BCID result are **transmitted to the CP and JEP systems** as Low Voltage Differential Signals (LVDS). The data format differs between the two systems. The CP uses the full trigger tower resolution of (0.1×0.1) in $(\eta \times \phi)$. To reduce the amount of links required between the systems, the PreProcessor multiplexes the data for two channels onto a single link. The JEP operates instead on JetSums with a size of (0.2×0.2), which the PreProcessor prepares accordingly. In both cases, the data is then packed into a specified format and sent along LVDS links at a data rate of 480 MBit/s.

Besides performing these main data processing steps, the PreProcessor also has to provide event data to the DAQ upon reception of an L1A signal from the CTP. The data is mainly comprised of the 10-bit digitised inputs and the 8-bit energy values obtained from them. The PreProcessor also provides rates and spectra of both the 10-bit and the 8-bit values in a non-event-based way, by collecting the corresponding data and making it avaiblable for readout independent of L1A signals. Finally, the PreProcessor is connected to the DCS system and has to supply it with environmental data like the temperatures and supply voltages of its components.

Consequences of Wrong BCID Decisions

The proper performance of the BCID algorithms is crucial to the functionality of the L1Calo algorithms. If a trigger tower signal is assigned to the wrong BC, the object identification algorithms in the CP and JEP are impacted, as the E_T of the trigger tower is set to zero in the actual BC. The effect differs between *late* and *early* mis-identified trigger towers:

- Because a typical event contains more than one trigger tower with a significant energy deposition for a particle with high E_T, a late mis-identification of a trigger tower usually does not result to the CP and JEP algorithms missing the particle. Because the E_T of the identified particle is reduced, the mis-identification can lead to the particle candidate missing an E_T threshold, but the event is usually still triggered by the L1Calo system. Furthermore, events usually contain more than one particle, and these can also lead to the event being triggered. If the event is triggered, the detector read out occurs normally and the event data that is used for physics analyses is preserved.

- If the E_T of a trigger tower is assigned to an early BC, and the E_T is sufficiently large for the CP and JEP algorithms to identify it as an object, the early BC can be triggered by the L1Calo system. If this happens, the actual event is not triggered, because of the dead-time imposed by the CTP system after

each triggered event (see Section 3.4.1). Thus, only the detector information for the early BC is retrieved, while the information from the actual BC is lost. While a partial reconstruction of the actual BC is possible also in these cases, the energy resolution of the calorimeters is strongly reduced in these cases. For the inner detector only a single BC is read out for each trigger, so the tracking information is lost in case of wrong triggers. Because of this, early mis-identifications impact physics analyses more strongly than late mis-identifications.

4.3.2 Hardware Implementation

The PreProcessor is implemented as a system of VME[4] modules organised into eight 9U VME64xP crates that are situated in the ATLAS electronics cavern.

Each crate contains several types of modules, as listed below:

- The main processing components are the *PreProcessor Modules* (PPMs). They are 9U VME modules that provide the functionality described above. In total, 124 PPMs are installed in the PreProcessor system. Two of the eight PreProcessor crates are equipped with 14 PPMs, while the other six crates hold 16 PPMs each. All PPMs are identical, designed to process 64 trigger tower signals each that are received over four input cables.

- For each PPM, one *Rear G-Link Transmitter Module - Optical Tx* (RGTM-O) is installed on the back side of the crate. It transfers the event data from the PPM and transmits it to the DAQ system via the corresponding ROD.

- Each crate holds one 9U *Timing Control Module* (TCM), which implements the interfaces to the TTC and DCS systems. From the TTC side, it receives various protocol signals, like the 40.08 MHz LHC bunch-crossing clock, the L1A signal and the bunch- and event-counter resets, and supplies them to all PPMs in the crate. On the DCS side, it collects the environmental data from all PPMs in the crate and transmits it to the DCS.

- One 6U *Single Board Computer* (SBC) functions as crate controller. It is used to control the other modules in the crate via the VMEbus, by sending configuration data and operational commands.

Figure 4.6 shows a PPM. The main signal processing path flows from the left to the right of the module. The trigger tower signals are first received in four input connectors, each of which accepts up to 16 signals. From each connector, the signals are routed to one *Analogue Input* (AnIn) board, that performs the required signal conditioning to prepare the input signals for digitisation. The conditioned analogue signals are then routed to 16 *Multichip Modules* (MCMs), with each MCM processing four trigger tower signals. These are the main processing units of the

[4]Versa Module Europa

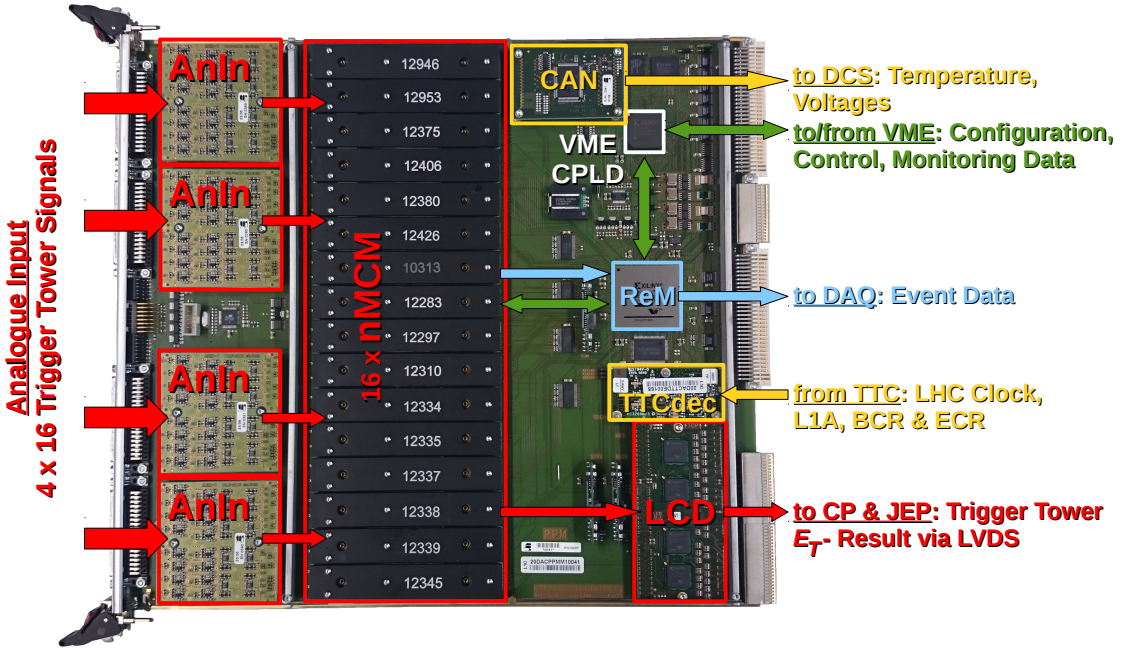

Figure 4.6: The PreProcessor module in Run-2. The various submodules and their connections are indicated, as are the in- and outputs. The MCMs have been replaced with nMCMs (see Chapter 5)

PreProcessor and implement the majority of the functions described before: signal digitisation and synchronisation, Bunch-Crossing Identification and serialisation and LVDS transmission of the E_T results. The LVDS streams from the MCMs are routed to the *LVDS Cable Driver* (LCD), which performs the fan-out required to correctly map the trigger tower signals to the CP and JEP modules and applies pre-compensation to the signals to ensure error-free transmission.

The MCMs store the digitised input data and E_T results in decicated readout buffers. Upon reception of an L1A, the data regarding the triggered BC is collected in the *Readout Manager* (ReM) FPGA[5]. The ReM FPGA then formats the data and sends it to the DAQ. The ReM FPGA also serves as the main configuration platform of the PPM, by receiving configuration and control commands from the SBC via the VME interface and distributing them to the addressed PPM components. It also collects monitoring data from the MCMs and makes it accessible to VME.

The TTC signals are received on the *TTC Decoder* (TTCdec) card, where they are decoded by the *TTC Receiver* (TTCrx) chip. The signals are then routed to the ReM FPGA, which distributes them to the other submodules. Data on the temperature and supply voltages of the PPM components is collected in a microcontroller device that is connected to the DCS system via a CANbus and the TCM.

Figure 4.7 shows the η-ϕ coverage of the PPM and its components for the hadronic calorimeter layer. In the central region, each PPM covers a (0.4×1.6) area in

[5]Xilinx XCV1000-E

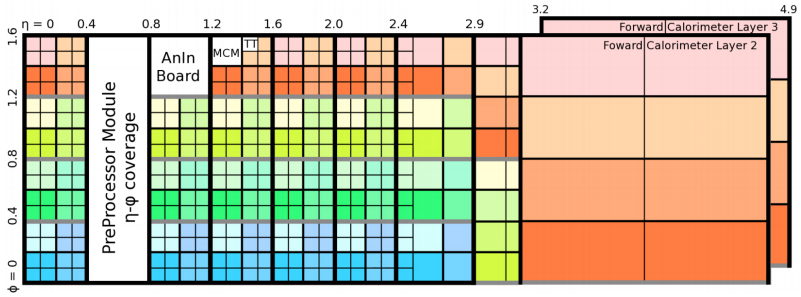

Figure 4.7: Coverage of the different PreProcessor module components in the hadronic calorimeter for $\eta > 0$ and one quadrant in ϕ. [48]

($\Delta\eta \times \Delta\phi$), with each AnIn board processing a group of 4×4 trigger towers. The four trigger towers processed by one MCM are arranged in a 2×2 pattern. In the forward region, where the trigger towers become larger, the mapping differs from this, as indicated by the colour coding in the figure. In the extreme case of the FCal, the entirety of FCal2&3 is covered by two PPMs, one each for the positive and the negative η direction.

The coverage in the electromagnetic calorimeter is largely identical to the hadronic calorimeter, with the exception of the FCal. As the FCal1 towers are more granular than their FCal2&3 counterparts, they require twice as many PPM inputs (see Figure 4.2). This factor of two is balanced by the fact that only a single layer (FCal1) is used compared to the two layers in the hadronic calorimeter, resulting in the same coverage of the full FCal1 by two PPMs.

The following sections provide a more detailed description of the most important PPM components.

Analogue Input Board

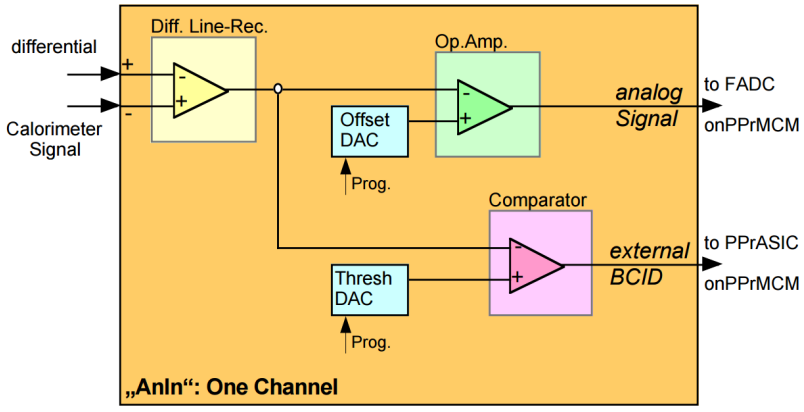

Figure 4.8: Analogue signal conditioning circuit on the AnIn board. [49]

Each AnIn board receives and conditions signals from 16 trigger towers. Figure 4.8 shows the signal conditioning circuit for one trigger tower. In the first step,

the differential signals are converted to single-ended signals by a differential line-receiver. The output of the line-receiver is split into two copies, one of which is routed to an operational amplifier that fits the signal into the digitisation window. For this purpose, it applies a fixed gain factor of $G = 0.43$, reducing the maximum amplitude of the trigger tower signals of 2.5 V such that they fit the 1 V digitisation window. Additionally, the output of an 8-bit programmable *Digital-to-Analogue Converter* (DAC) is added to the signal, with 2.4 mV per DAC step. This allows to configure the baseline of the single ended signal. The DAC is used for two purposes: to set the baseline such that the signals fit entirely within the digitisation window, including the undershoot of the LAr signals; and to guarantee a uniform baseline for all PreProcessor channels. The resulting signal is then routed to the corresponding MCM input.

The other copy of the line-receiver output is routed to a voltage comparator. It checks the signal voltage against a programmable threshold that is defined by a further 8-bit DAC in 10 mV steps. The comparator output is a single bit digital output that indicates whether the input signal exceeds the threshold. It is forwarded to the MCM, where it is used in the External BCID algorithm (see Section 5.2.3).

Multichip Module

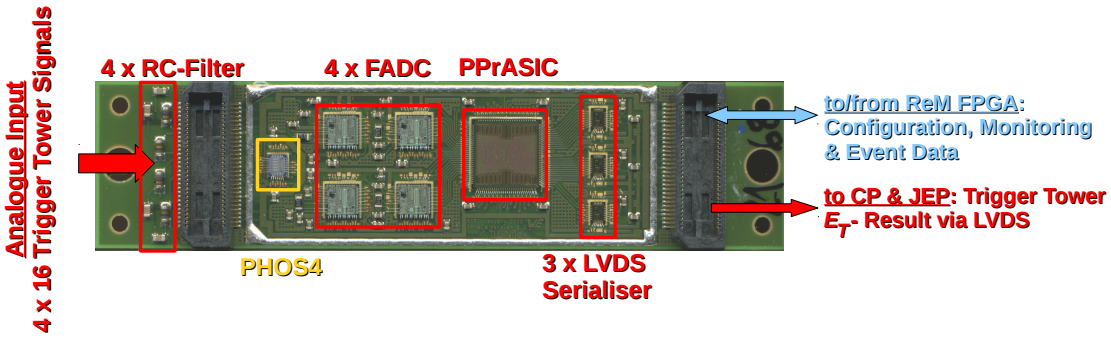

Figure 4.9: The Multichip Module installed on the PPMs during Run-1.

The MCM is shown in Figure 4.9. Several components are installed on the module to perform the signal processing.

From the input connector, the analogue signals are first routed through a low-pass filter in order to reduce high-frequency noise. As the filter also reduces the signal amplitude, the gain factor applied on the AnIn boards is increased accordingly.

The signals are then digitised in four *Flash Analogue-to-Digital Converters* (FADCs), one for each of the MCM inputs. The sampling occurs at the LHC bunch-crossing frequency of 40.08 MHz with a 10-bit resolution, with one ADC count corresponding to 250 MeV. The digitisation strobes are provided by a *PHOS4 Timing Chip* [50]. It allows to displace the strobe with respect to the 40.08 MHz in steps of 1 ns. The strobes can be set independently for the four FADCs, and are configured from VME

via the ReM FPGA, which then transfers the values to the PHOS4 chip over an Inter-Integrated Circuit (I^2C) bus.

The digitised signals are routed to a custom-built, 4-channel *Application Specific Integrated Circuit* (ASIC), the PPrASIC. It implements all the main functionalities of the PreProcessor described in the previous section, i.e. signal synchronisation and Bunch-Crossing Identification. The internal processing is described in more detail in Section 5.2.

The PPrASIC provides three 10-bit parallel streams of E_T data, two for the CP and one for the JEP processor. Each stream is routed to one of three *LVDS Serialisers* that are installed on the MCM. They serialise the 10-bit input data and transmit it to the LCD board at 400 MBit/s. The PPrASIC also implements the event data readout and monitoring functionalities described above.

During the LS1 upgrade [42], the MCMs were replaced by an enhanced version. This upgrade is described in detail in Chapter 5.

LVDS Cable Driver

The LCD board receives a total of 48 LVDS streams from the 16 MCMs, which it drives along 11 m long cables to the CP and JEP systems. To accomplish this the LCD fulfills two tasks. First, the different processor modules require data from the same trigger tower inputs. These are situated along the azimuthal borders of the processor modules. The signals of these trigger towers are thus duplicated in four FPGAs[6] on the LCD. The second function of the LCD is to pre-compensate for losses along the cables to the processors. This is achieved by an RC circuit, through which the signals are routed to the output connector.

Readout Manager FPGA

The ReM FPGA has two main tasks: the configuration of the PPM components; and the collection of event data from the MCMs and subsequent transmission to the DAQ system. For these purposes, it implements several interfaces that connect it to the various PPM components.

Figure 4.10 shows an overview of the functional blocks of the ReM FPGA:

- A *VMEbus* connects the ReM FPGA to the crate controller. The ReM FPGA receives configuration data and control commands from the crate controller, and transmits monitoring data to it.

- Four *Serial Peripheral Interface* (SPI) buses allow the ReM FPGA to configure the DACs on the AnIn boards.

- 32 *Serial Interface* links connect the ReM FPGA with the 16 MCMs, allowing for bidirectional data transfer. They are used to send configuration data and

[6] Xilinx X2CV_250

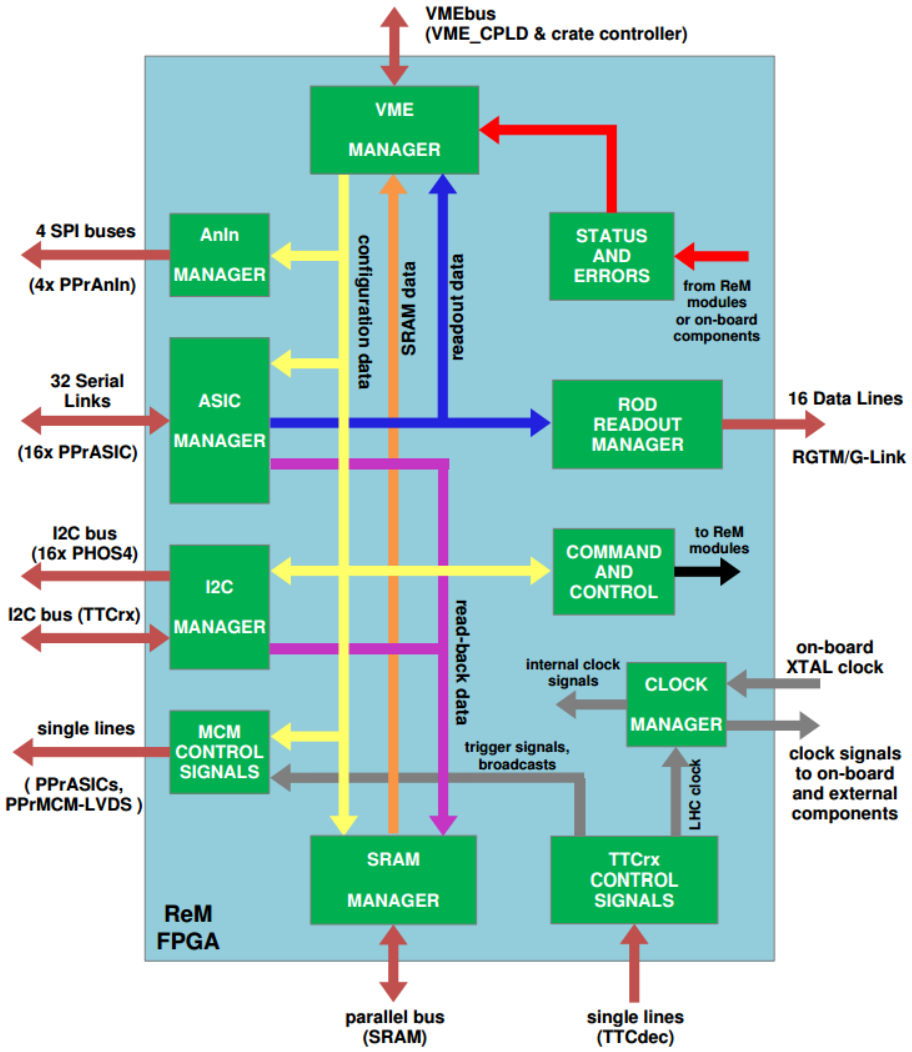

Figure 4.10: Block diagram of ReM FPGA functionality. [44]

requests for monitoring data to the MCMs, and to transmit event, monitoring and status data in the other direction.

- 16 *I^2C data buses* are used by the ReM FPGA to configure the PHOS4 chips on the MCMs. An additional I^2C bus provides the ReM FPGA with access to the TTCrx chip.

- A bidirectional parallel bus connects the ReM FPGA to an on-board *Static Random-Access-Memory* (SRAM) device, that is used to extend the memory resources of the ReM FPGA.

- The ReM FPGA receives *clock signals* from both an on-board crystal oscillator (XTAL) at 40.00 MHz and from the TTCdec card at the LHC bunch-crossing

frequency of 40.08 MHz[7]. The latter is used to drive the majority of the internal processing of the ReM FPGA and is distributed to most of the other PPM components. The 40.00 MHz XTAL clock is used in the data transmission to DAQ on the RGTM-O.

- Additional single data lines connect the ReM FPGA to different on-board devices and are used to collect and distribute status and control data.

4.4 Cluster Processor

4.4.1 Functionality

The CP system [43] identifies energy clusters in the calorimeters that correspond to isolated electrons (e), photons (γ) or hadronically decaying τ particles (τ/had). For this purpose it receives the 8-bit PreProcessor E_T output of trigger towers within $|\eta| < 2.5$, where the finest η-ϕ resolution is provided.

The CP implements two algorithms - one optimised for e/γ candidates and one for τ/had candidates. Both algorithms operate on a window of 4×4 trigger towers over both the electromagnetic and the hadronic calorimeter layer. The algorithm scans the entire region defined above by moving the window in steps of one trigger tower in either η or ϕ. This is referred to as a *sliding window*.

As shown in Figure 4.11 (a), the window is divided into four regions: an electromagnetic core of 2×2 trigger towers, an electromagnetic isolation ring of 12 trigger towers surrounding it, and a corresponding hadronic core and isolation ring. Within the electromagnetic core region, the four possible E_T sums of neighbouring trigger towers are calculated as indicated in the figure.

An e/γ candidate is identified if one of the four core trigger tower sums is larger than a programmable threshold. Additional isolation criteria can be required, that are satisfied if the E_T sum in the two isolation rings are below programmable thresholds. A hadronic veto can also be applied, that places a similar threshold on the E_T sum of the hadronic core area.

The τ/had algorithm considers the same 2×1 trigger tower E_T sums as above and adds the hadronic core energy to each of them. This sums are then compared to programmable thresholds. Isolation criteria can be applied exactly as in the e/γ case.

An additional criterion is applied to avoid multiple identifications of the same object by the sliding window algorithm: the total E_T sum of both the electromagnetic and the hadronic core areas is calculated, and then compared the the equivalent sums of eight surrounding areas as shown in Figure 4.11 (b). If a local maximum is detected, the object is identified and the η-ϕ-coordinate of the object is sent to the HLT with the RoI information.

[7] If no TTC connection is provided, the TTCdec sends a 40.00 MHz clock instead, that is generated from a local crystal oscillator.

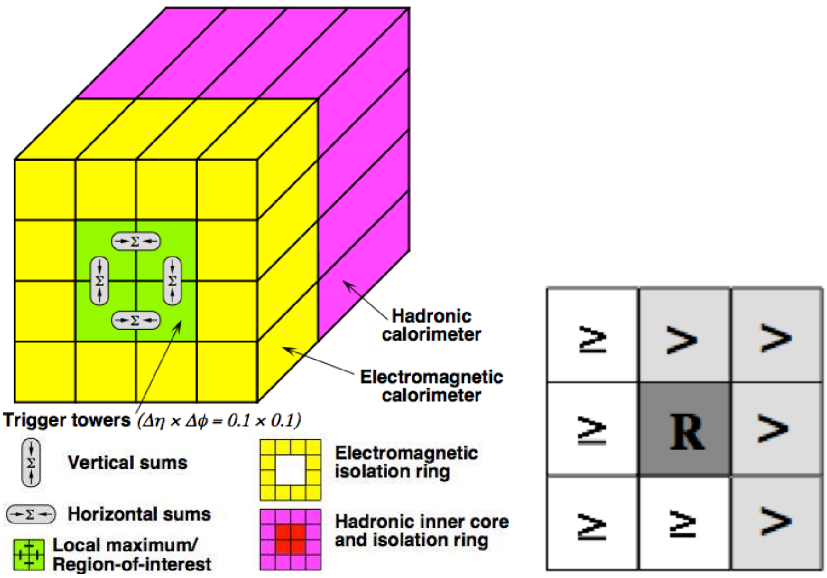

(a) CP Clustering **(b)** RoI Definition

Figure 4.11: Illustration of the CP algorithm. (a) shows the definition of the clustering algorithm, (b) shows the RoI definition. [43].

4.4.2 Hardware Implementation

The CP consists of four VME crates, each of which is equipped with 14 9U *Cluster Processor Modules* (CPMs) that perform the algorithms detailed above, two merging modules that collect the data from the CPMs and transmit it to the central trigger systems and one TCM that distributes the LHC clock and protocol signals to the other modules and collects environmental data for the DCS.

Each CPM covers an η-ϕ area equivalent to one PPM, i.e. 90° in ϕ and 0.4 in η. This corresponds to 64 overlapping 4×4 windows.

The trigger tower E_T data from the PreProcessor is received on the crate backplane. A total of 160 serial streams are routed to each CPM, 80 each for trigger towers from the electromagnetic and the hadronic layer. The streams are first deserialised and then processed in 20 *Serialiser FPGAs*. These devices re-serialise the data streams and then transmit them to eight *CP Chips*, that perform the CP algorithms. The Serialiser FPGAs also fan-out 120 of the re-serialised streams and transmit the copies to other CPMs in the crate over the backplane, taking care of overlaps in the processed areas.

The outputs of the CP Chips are collected in two Merging FPGAs, that gather the data of the CPM and transmit it to the merging modules, that build crate-wide result summaries. A single merging module in one of the CP crates gathers the results from all CP mergers, to provide the system-wide object multiplicity to the CTP.

Two *Readout Controller* (ROC) FPGAs collect data from the processing components that are transmitted as event data to the DAQ and as RoI information to the

HLT upon reception of an L1A. The LHC protocol signals like the L1A are received via a TTCdec card, identical to the one installed on the PPM.

The two merging modules gather the results of the CPMs in the crate to construct crate- and system-wide results. In Run-1, the *Common Merger Modules* (CMMs) [43] received the multiplicity of objects above programmable E_T thresholds from the CPMs, which were then summed to produce the total result that was transmitted to the CTP. For Run-2, the CMMs have been replaced by the *Common Merger Modules Extended* (CMXs) [42]. Instead of threshold multiplicities, these receive identified object candidates in the form of *Trigger Objects* (TOB). The TOBs describe the type, position and energy of the object candidates. The CMX then transmits the TOBs to the L1Topo system and compares the TOB energies to programmable thresholds to produce the multiplicities that are then transferred to the CTP.

The CPM hardware has been left unchanged during the LS1 upgrade, the firmware implementation of the FPGAs has been updated. As mentioned above, the format of data sent to the merger module changed from threshold multiplicities to TOBs. This change is supported by increasing the data rate over the crate backplane from 40 to 160 MBit/s. The trigger algorithms have also been improved by changing the isolation threshold from a fixed value to an E_T dependent variable [41].

4.5 Jet-Energy Processor

4.5.1 Functionality

The JEP system [43] locates jet candidates in the calorimeters and computes global energy sums. Instead of trigger towers, it operates on (0.2×0.2) JetSums, which it receives from the PreProcessor for the full calorimeter coverage of $|\eta| < 4.9$. The JEP computes 10-bit *Jet Elements*, by combining the two 9-bit JetSums from the electromagnetic and hadronic calorimeters that share the same η-ϕ coordinate.

The jet algorithm scans the calorimeters with three sliding windows different sizes of 2×2, 3×3 and 4×4 Jet Elements. Similar to the CP algorithm, the windows are slided by one step into either the η or the ϕ direction, to cover the entire area. At each window position, the E_T contained in the window is computed by summing the Jet Elements and compared to programmable thresholds. To uniquely identify local maxima, the RoI for the 3×3 algorithm is defined as the window that gives the largest E_T sum, as shown in Figure 4.12. For the other two window sizes, the RoI is identified by the central four Jet Elements.

The Jet Elements are also used to compute the global energy sums as follows:

- The total transverse energy ($\sum E_T$) is calculated as the *scalar* sum of the E_T of all Jet Elements.
- The E_T^{miss} is calculated as the *vectorial* sum of the Jet Element E_T. For each Jet Element, the components along the x and y directions, E_x and E_y, are computed by multiplying the E_T with a factor of $\cos\phi$ and $\sin\phi$ respectively.

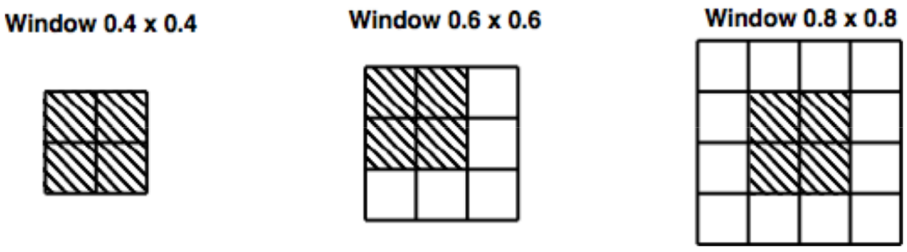

Figure 4.12: Illustration of sliding windows used in the JEP algorithm and respective RoI positions. [43]

The E_T^{miss} then follows from the sums of E_x and E_y across the full JEP system as

$$E_T^{miss} = \sqrt{(\sum E_x)^2 + (\sum E_y)^2}, \tag{4.1}$$

4.5.2 Hardware Implementation

The JEP consists of two crates, each housing 16 *Jet/Energy Processor Modules* (JEMs) that implement the jet and energy sum algorithms, two merging modules to collect the data and one TCM.

Each JEM receives JetSum data from four PPMs, covering an area of 4×8 Jet Elements over both the electromagnetic and the hadronic calorimeter. A total of 88 LVDS streams per JEM are received over the crate backplane and de-serialised. The resulting parallel streams are routed to four *Input FPGAs* that sum the 9-bit JetSums from the two calorimeter layers to 10-bit Jet Elements. Similar to the CP, the Input FPGAs also fan-out the signals required by neighbouring JEMs and distribute them via the crate backplane.

The Jet Elements are then sent to the *Jet FPGA*, which performs the jet algorithm. The Input FPGAs also compute the E_x and E_y components for the Jet Elements which, together with the E_T, are sent to the *Sum FPGA* that computes sums of the three quantities for the 4×8 core area of the JEM.

One of the merger modules in the JEP crate collects the jet algorithm results, while the other receives the local energy sums and performs the $\sum E_T$ and E_T^{miss} calculations. The latter is performed using a look-up table for the final adding-in-quadrature operation (Equation 4.1).

Similar as for the CP, the Run-1 CMMs have been replaced by CMXs for Run-2, changing from object multiplicities to TOBs for the jet information transferred over the crate backplane. The TOBs are transmitted to the L1Topo system and compared to energy thresholds to obtain threshold multiplicities that are sent to the CTP. The energy sums are similarly compared to thresholds and the result is sent to the CTP.

5 The new Multichip Module

The higher centre of mass energy and the increased luminosity delivered by the LHC after LS1 impose several challenges for the ATLAS detector and TDAQ systems. For the L1Calo system, the main consequences are the following:

- The pile-up noise on the trigger tower signals increases significantly, surpassing the thermal noise in almost all detector regions. A high detection efficiency for low energy signals thus requires an adapted noise reduction strategy compared to Run-1.
- The rates of several trigger items - most notably the missing energy triggers - showed a strong non-linearity as a function of luminosity in Run-1. Due to the increased luminosity, these high fake-rates are even more severe in Run-2, and thus have to be corrected for.
- The higher centre-of-mass energy leads to a higher probability for very large signals, which makes strong saturation effects more likely to occur. The processing of saturated signals in Run-1 was not validated beyond centre-of-mass energies of 8 TeV. Improvements of the handling of saturated signals are thus required for Run-2.

The PreProcessor system was upgraded to accomodate for the new requirements at an early stage in the analysis of the trigger tower signals. An upgrade of the PreProcessor also makes further improvements possible, like separate energy calibrations for the CP and JEP systems. For this upgrade, an improved replacement module for the MCM has been developed: the new Multichip Module (nMCM).

The nMCM design makes use of the improvements in technology since the development of the original MCM. Most importantly, the nMCM no longer uses the custom PPrASIC to implement the majority of the trigger algorithms. Instead, a commercial Field-Programmable Gate Array (FPGA) has been chosen. These reprogrammable chips offer a flexibility that is not achieved with the hard-coded ASIC implementation of the MCM. Modern FPGAs also offer sufficient resources to both implement the Run-1 PreProcessor algorithms and to improve on them, with a similar form factor as the PPrASIC.

As changes to the PPM motherboard should be kept to a minimum, strong requirements are imposed on the design of the nMCM. Firstly, the form factor of the nMCM is constrained to a maximum width equal to that of the MCM, so that the PPM still fits 16 of the modules. Furthermore, the position and the pin-out of the connectors has to be identical between the module types, to ensure the proper connectivity to the other PPM components.

Like the MCM, each nMCM processes four trigger tower channels. Next to improving on the signal processing in the ways mentioned above, it has to perform all the tasks previously fulfilled by the MCM:

- The analogue trigger tower signals have to be digitised;
- The input signals have to be synchronised to account for different arrival times;
- Significant energy deposits in the trigger tower have to be identified and their transverse energy extracted from the signal shape;
- Transverse energy results have to be transmitted as LVDS signals to the CP and JEP systems;
- Event processing data has to be provided to the DAQ system upon reception of an L1A.

This chapter describes the nMCM implementation. The hardware is described in Section 5.1, while the initial implementation of the signal processing algorithms is given in Section 5.2.

5.1 Hardware Implementation

Figure 5.1 shows the production version of the nMCM, highlighting the major components on both the top and the bottom side of the module.

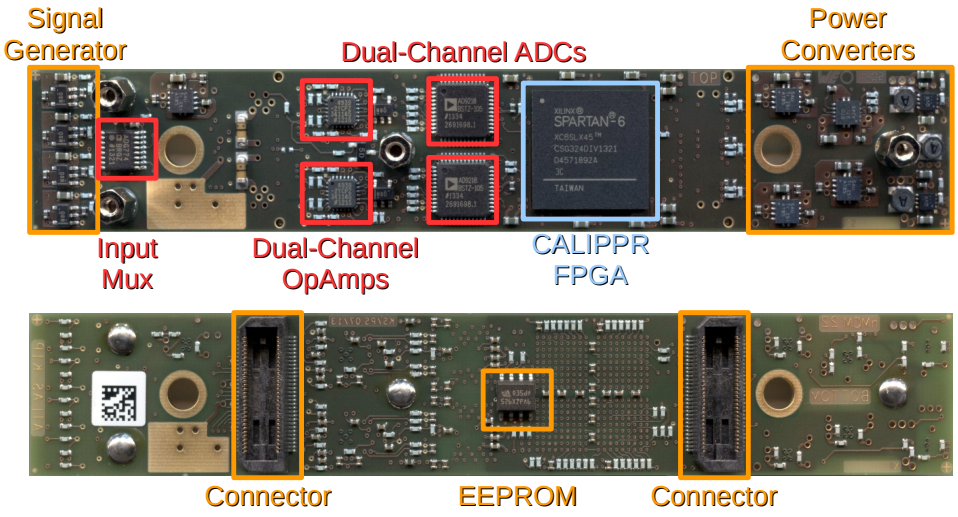

Figure 5.1: The new Multichip Module. The top side is shown on top and the bottom side is shown on the below.

Compared to the MCM, the new module is slightly longer, while keeping the same distance between the connectors. This allows the placement of additional

components without violating the space requirements of the individual modules on the PPM.

The main signal processing components are placed on the top side of the module. The four analogue trigger tower signals are routed from the input connector to two dual channel ADCs[1]. Since the ADCs require differential inputs, two dual channel Operational Amplifiers (OpAmps)[2] convert the signals from single ended to differential mode.

The signals are digitised with 10-bit resolution and a sampling frequency of 80 MHz. This frequency is faster than that available on the MCM by a factor of two. The main reason for this increase is to keep the digitisation latency identical between the two systems: The ADCs require five clock cycles to propagate the digitisation results through the internal processing pipeline [51]. Using the 80 MHz frequency, this corresponds to 2.5 BCs which is equal to the delay of the ADCs used on the MCM. In addition, the higher digitisation frequency provides new possibilities in the digital signal processing (see Sections 7.4.2 and 8.2.2).

The digitisation results are routed to a Xilinx Spartan-6 FPGA device[3]. It is referred to as the CALorimeter Information PrePRocessor (CALIPPR) FPGA. Its programming implementation combines the functionalities of the PPrASIC, the PHOS4 chip and the LVDS drivers present on the MCM. This is described in detail in Section 5.2.

The FPGA is of volatile nature, i.e. its progamming is only sustained as long as the device remains powered. To provide an efficient way to reprogram the CALIPPR FPGA after the power of the nMCM has been cycled, an Electically Erasable Programmable Read-Only Memory (EEPROM) is installed on the module. It has a memory size of 2 MB, which is sufficient to store a single programming bit-file of about 1.5 MB. It is used to automatically program the CALIPPR FPGA every time the nMCM is powered up. In addition, the EEPROM holds an identification number, the nMCM ID. This number is unique to each module and allows the tracking of all the modules in the system.

Because the configuration path for the EEPROM is via the VME bus over the ReM and CALIPPR FPGAs, the EEPROM can not be directly reconfigured if the CALIPPR FPGA is not programmed. For cases where the EEPROM is empty, like is the case after the production of the nMCM, an alternative way of loading the firmware onto the module is thus required. In these cases, the CALIPPR FPGA is programmed over the JTAG chain of the PPM, which allows data transfer over the VME bus. This configuration process is much slower than the loading from the EEPROM and is thus reserved for special cases. Once the CALIPPR FPGA is programmed, the EEPROM can be accessed and the bit-file can be stored.

The nMCM also contains an on-board Signal Generator circuit. It is controlled from the CALIPPR FPGA and allows the generation of analogue signals that can be

[1] Analogue Devices AD9218
[2] Analogue Devices ADA4939-2
[3] Xilinx Spartan6 XC6SLX45-3 (CSG324)

used as input to the ADCs. This makes it possible to perform standalone functionality tests of the nMCM, without the need for an external signal. For a description of the Signal Generator functionality see Section 5.2.7. A quad-2:1 multiplexer device[4] allows to select channel-wise the digitisation input, i.e. the trigger tower signal or the Signal Generator Output. Its configuration is realised from VME via the CALIPPR FPGA, with the trigger tower signal being the default.

The PPM supplies the nMCM with the same power voltages as the MCM, +3.3 V and +5 V. A set of power conversion circuits derives from these the voltages required by the active components on the nMCM. The +3.3 V is used to derive the different voltages for the CALIPPR FPGA, while the +5 V is used to power the ADCs and the remaining active components, as shown in Figure 5.2.

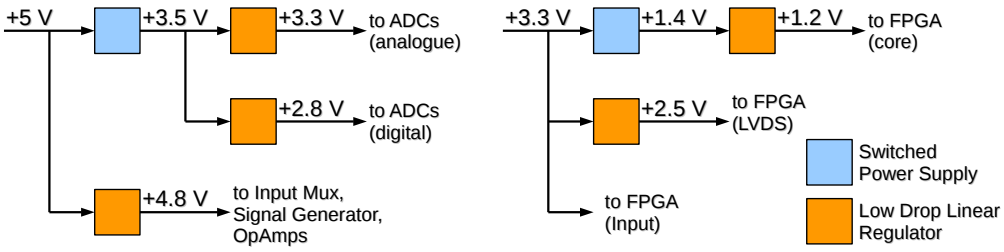

Figure 5.2: Power conversion scheme on the nMCM.

5.2 Digital Signal Processing

The main digital processing unit on the nMCM is the CALIPPR FPGA, which implements the majority of the PreProcessor algorithms.

Figure 5.3 shows an overview of the functional blocks implemented in the initial version of the CALIPPR FPGA implementation. This version focussed on implementing the core functionality, upon which the additional processing algorithms for Run-2 are based. The initial implementation is thus largely based on the PPrASIC functionality [52], with only minor changes in order to adapt to the new hardware. It was primarily used in the nMCM production tests, described in Chapter 6. The additional algorithms were introduced only after the tests, and are described in Chapters 7 and 8.

The CALIPPR FPGA implementation is organised in four channels. Each channel receives and analyses the digitised data for one of the four trigger towers processed by the nMCM. The data is first latched with an internal clock signal in the Input block. The latched signals are sent to a FIFO block that is used to synchronise the signals from different trigger towers across the whole PreProcessor system. The synchronised signals are analysed by the BCID logic. These algorithms identify BCs for which a significant energy deposition is present in the trigger tower and calibrate

[4] Analogue Devices ADG774

the input signals for these BCs to transverse energy values. The E_T results are then formatted and sent to the CP and JEP systems using an LVDS transmitter.

In addition to these main trigger algorithms, the CALIPPR FPGA also implements event readout, monitoring and testing facilities. The Readout block provides event data to the DAQ system upon reception of an L1A. The Rate Metering and Histogramming algorithms are used to monitor both the input ADC data and the output of the BCID algorithms. The Playback block makes it possible to use programmable patterns as input to the processing logic instead of the latched ADC data, to perform tests of the algorithm logic. The CALIPPR FPGA also controls the Signal Generator circuit mentioned before.

The CALIPPR FPGA is configured from VME via the ReM FPGA. The configuration data is transmitted using a custom Serial Interface, with groups of two channels being served by the same interface. The Serial Interface is also used to transmit event readout and monitoring data from the nMCM to the ReM FPGA. The ReM FPGA also forwards the LHC bunch-crossing clock to the CALIPPR FPGA. All clock signals required by the design are derived from this input in the Clock Manager block. This block also generates the ADC strobes used to digitise the trigger tower signals.

These functionalities are described in more detail in the following sections.

5.2.1 ADC Input

The ADCs digitise the analogue input signals to the nMCM at twice the LHC bunch-crossing frequency, i.e. at 80.16 MHz, and the results are routed to the input stage of the CALIPPR FPGA.

Figure 5.4 shows a timing diagram of the involved signals. The top box in the figure shows the signals in the ADC, while the lower box shows the signals in the input stage of the CALIPPR FPGA.

In the ADC the analogue input signal is digitised every time the 80 MHz ADC strobe signal has a rising edge, leading to a new value being transmitted as ADC output to the input stage of the CALIPPR FPGA.

The input stage is divided into two steps. In the first step, the ADC data is latched with the internal 80 MHz clock of the CALIPPR FPGA. Both this internal clock and the ADC strobe are aligned with the 40 MHz LHC clock, but the latter is shifted by the fine timing delay (see Section 5.2.2). To avoid metastability, where the ADC data is latched while it is being updated, the ADC data is latched with both the rising and the falling edge of the internal clock. Depending on the fine timing setting, only one of these two signals is retained for further processing. In the example shown in Figure 5.4 the rising edge is used, .

The majority of the signal processing algorithms in the CALIPPR FPGA operate at 40 MHz. Because of this, the second step of the input stage reduces the input data stream from 80 MHz to 40 MHz by dropping every other sample from the data stream. This is done by performing an edge selection with a 40 MHz clock. Depending on the edge, either the 'first' or the 'second' 80 MHz sample in the 40 MHz clock

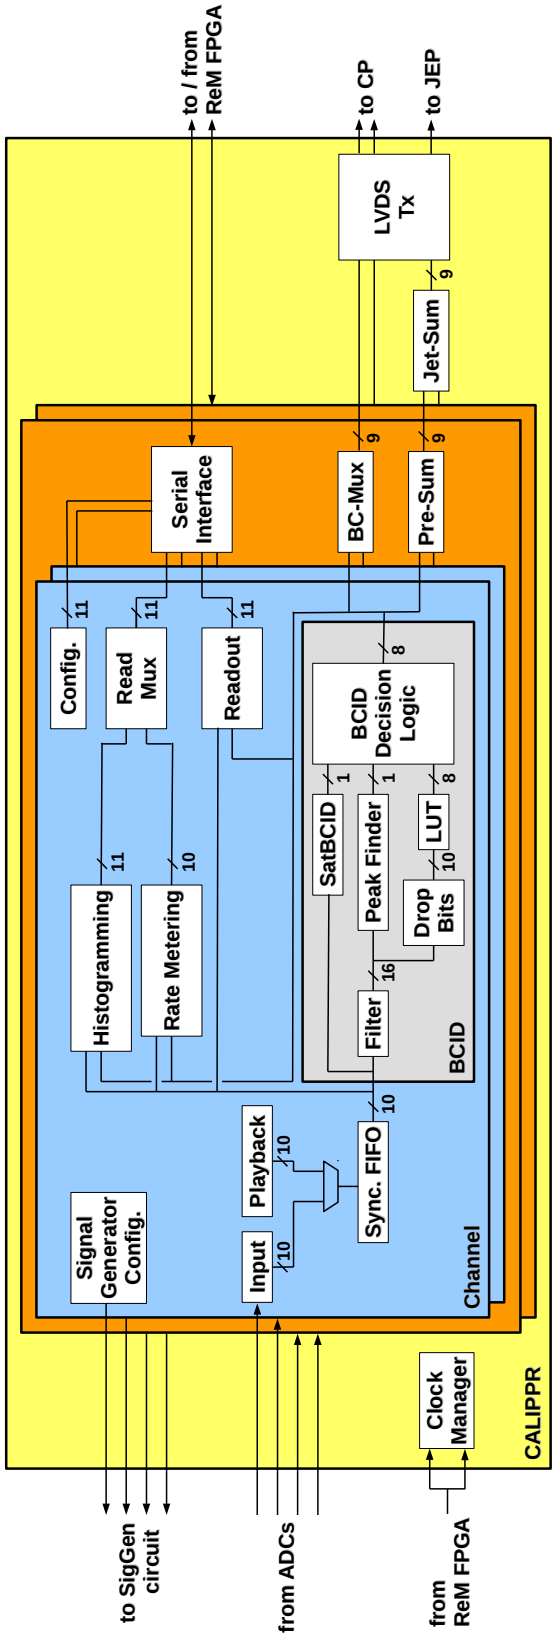

Figure 5.3: Overview of the initial implementation of the CALIPPR FPGA.

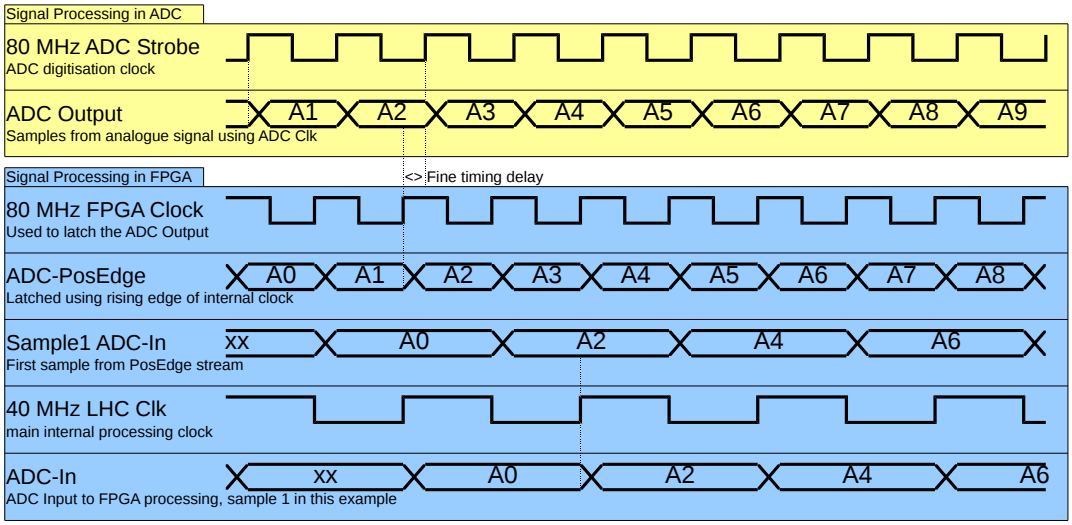

Figure 5.4: Timing diagram of the ADC input for the case of using the rising edge of the 80 MHz clock and the first 80 MHz sample in the latching configuration.

period is selected. Figure 5.4 shows a case where the first sample is selected.

The two steps of the input stage are also referred to as 80 MHz latching and 40 MHz latching.

5.2.2 Input Timing

The analogue signals of different trigger towers do not arrive at the PreProcessor inputs synchronously. The major contribution to the timing difference between trigger towers are different cable lengths: Depending on the position of the trigger tower, the cables from the calorimeters to the ATLAS trigger cavern have a length between 30 and 70 m. The maximum difference of 40 m corresponds to a delay of about 200 ns, i.e. 8 BCs. Further timing differences are caused by the time-of-flight differences between the different detector regions. Correcting for these timing differences of the trigger tower inputs is a crucial aspect to the PreProcessor functionality.

Two timing mechanisms are implemented in the PreProcessor. The *coarse timing* corrects for the bulk of the timing difference. It makes sure that signals originating from the same BC are also processed in the same BC.

While the remaining timing difference after the coarse timing is less than one BC, it can still lead to analogue signals being sampled many ns away from the peak by the ADCs. Because the BCID algorithms require a sampling close to the peak for their optimal performance, the *fine timing* corrects this difference.

Coarse Timing: Synchronisation FIFO

The coarse timing is implemented by means of a First-In-First-Out (FIFO) pipeline [52]. This is a series of 10-bit registers that store the ADC digitisation result for a configurable number of BCs. The maximum FIFO depth is 16 BCs, corresponding to a delay of 400 ns. The configuration also allows to bypass the FIFO entirely, leading to a 'delay' of 0 BC.

Fine Timing: ADC Strobe Generation

The fine timing mechanism delays the ADC strobe signal with respect to the 40 MHz LHC BC clock. On the MCM, this task was performed using the PHOS4 chip [50]. It allowed to delay all four ADC clocks on the MCM individually in 25 steps of 1 ns. On the nMCM, the CALIPPR FPGA takes over this functionality.

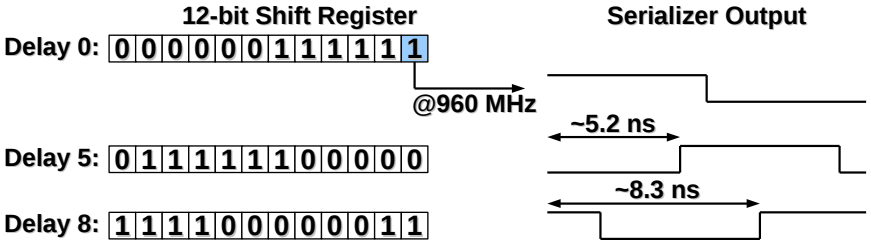

Figure 5.5: The fine timing mechanism in the CALIPPR FPGA. The upper row illustrates the generation of the ADC strobe for a fine timing delay of 0. The lower two rows show the fine timing delays 5 and 8.

Figure 5.5 shows an illustration of the fine timing mechanism. Instead of delaying existing clock signals, the FPGA generates the ADC strobes directly. This is done using high-speed output serializers to generate a clock-like signal. They are operated at 960 MHz, corresponding to one bit per 1.042 ns. By transmitting a pattern of 6 consecutive bits set to '1' and six bits set to '0', the output of such a serialiser behaves like a 80 MHz clock signal.

The pattern is defined by a configurable 12-bit register, where each pattern generates a clock with a different fine timing delay: By displacing the pattern by a number of bits n, the strobe is delayed by $n \times 1.042$ ns relative to the 40 MHz BC clock. In total, 12 different fine timing delays can thus be configured, covering the 12.5 ns period of the ADC clock in equidistant steps.

It should be noted that, in order to be as close as posible to the original MCM implementation, the configuration allows to set the delay in 24 steps. The additional 12 steps use the same delay settings as the first 12 steps, resulting in a double-coverage of the samples. In order to fully cover the 25 ns LHC period, the different 40 MHz latching configurations described before have to be used.

5.2.3 Bunch-Crossing Identification Logic

The Bunch-Crossing Identification (BCID) is one of the central tasks of the PreProcessor in the L1Calo system. The BCID algorithms analyse the digitised calorimeter signals to identify those BCs for which trigger towers contain a significant energy deposition. For each BC, the output of the BCID logic for a given trigger tower channel is the energy deposited in the trigger tower in that BC.

The different functional blocks of the initial version of the BCID logic are shown in Figure 5.3. The main BCID algorithm is a Finite Impulse Response (FIR) Filter that increases the signal-to-noise ratio, followed by a Peak Finder algorithm that finds local maxima in the filter output. If a peak is found, the respective BC is marked for the BCID Decision Logic block. A Look-Up-Table (LUT) memory is used to calibrate the filter output to the correct energy scale expected by the CP and JEP systems.

If the energy deposition in the trigger tower exceeds approximately 256 GeV, the respective analogue signal saturates the ADC. In these circumstances, the combination of FIR Filter and Peak Finder is no longer a reliable BCID mechanism. A dedicated algorithm for saturated pulses (SatBCID) is thus implemented in the BCID logic.

A third BCID method is provided by a voltage comparator installed on the AnIn board (see Section 4.3.2).

The results of the different algorithms are combined in the BCID Decision Logic, which produces the final energy output of the BCID logic for each BC.

The individual algorithms are discussed in the following.

Digital Filter

The FIR Filter [52] operates by integrating over five 40 MHz ADC samples, assigning a separate weight to each of them:

$$f_i = \sum_{k=0}^{5} c_k A_{i-2+k}, \tag{5.1}$$

where f_i is the output of the FIR filter for BC i, c_k is the k'th filter coefficient and A_n is the ADC sample at BC n. The filtering is performed for each BC, producing a stream of 16-bit filtered data.

The filter coefficients are 4-bit long, and are configurable individually for each channel. The typical configuration used in Run-2 was the *Matched Filter* scheme, where the coefficients resemble the expected pulse shape. This way, the input experiences the strongest amplification when a signal peak is centered in the filter window, while noise peaks are smoothed out. The Matched Filter scheme is optimal in reducing white noise, like the thermal noise induced by the signal processing electronics.

Depending on the shape of the analogue signal of a particular trigger tower, the undershoot can occur already within the 5 sample window considered by the filter.

This is typically the case for FCal towers. In these cases, negative coefficients are required for the respective samples in order to implement the Matched Filter scheme. Because of this, the outer two coefficients, c_0 and c_4, are defined as signed integers in the algorithm, while the other three coefficients are variables of unsigned integer type.

Peak Finder

The Peak Finder [52] analyses the output of the filter for each BC. If any of the samples satisfy the condition:

$$f_{i-1} < f_i \geq f_{i+1}, \tag{5.2}$$

BC i is identified by the algorithm and the output of the Peak Finder for that BC is set to '1'. For all other cases, the output is set to '0'.

Energy Calibration Look-Up Table

The energy calibration of the 16-bit filter output occurs in two steps: first, it is reduced to a 10-bit in the DropBits block, then the resulting 10-bit value is mapped to an 8-bit energy value [52]. This implementation was chosen for the PPrASIC as it optimises the resource usage without significantly reducing the resolution of the mapping, and is retained in the CALIPPR FPGA for these same reasons. The functionality of this Drop Bits block is illustrated in Figure 5.6 (b): A configurable 'window' of 10-bit width is selected from the filter output and the remaining bits are dropped. The configuration depends on the choice of filter coefficients, such that the 10-bit range of the LUT is covered in an optimal way.

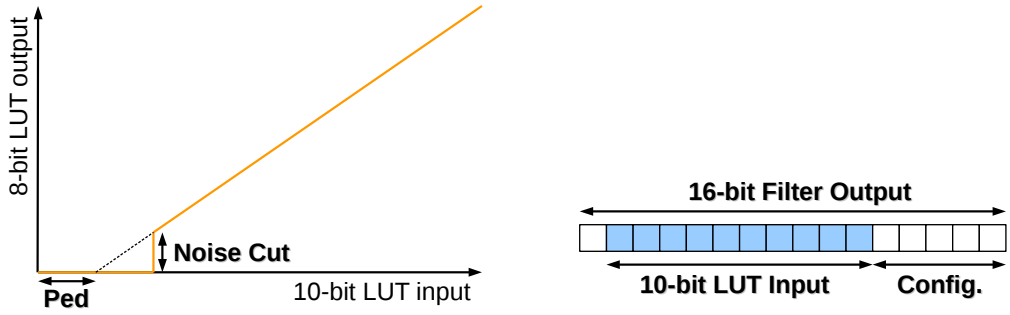

(a) Typical Run-1 LUT Configuration **(b)** Drop Bits Implementation

Figure 5.6: Illustration of LUT implementation. (a) shows the typical configuration of the LUT in Run-1, and (b) shows the functionality of the DropBits logic.

The LUT is implemented as a 1024×8-bit memory, where each entry represents an 8-bit energy value that is associated to a 10-bit LUT input address. By configuring

the memory with the corresponding values, arbitrary functions can be programmed into the LUT.

The typical LUT configuration used in Run-1 is illustrated in Figure 5.6 (a). It follows a linear function, with the LUT output representing transverse energies between 0 and 255 GeV. It also encodes two cuts: a pedestal subtraction and a noise cut. The pedestal subtraction removes the non-zero baseline (see Section 4.3.2) of the input signals, by mapping all filter outputs below this level to zero. The noise cut maps all the LUT inputs below a certain threshold to an output of zero. This suppresses all peaks that are identified by the Peak Finder that originate from low energy signals that are compatible with noise.

Saturated BCID

The SatBCID algorithm operates on the un-filtered 40 MHz ADC samples [52]. It analyses the non-saturated samples on the rising edge of the pulse. The steepness of the rising edge provides an estimation of the position of the peak of the analogue signal. This is illustrated in Figure 5.7 (a). It shows two saturated pulses with different energies. The smaller pulse barely reaches the ADC saturation, such that only the central sample is saturated. For the larger pulse, the sample in front of the maximum is also saturated, leading to a much steeper rise.

The SatBCID algorithm compares all ADC samples to three configurable thresholds: a saturation level, a *high* threshold and a *low* threshold. Whenever an ADC sample exceeds the saturation level, the samples before it are compared to the two other thresholds. The comparison results allow an estimation of the steepness of the edge, so that the peak position can be estimated according to the following logic:

- If ADC$(s-1) >$ *high* and ADC$(s-2) >$ *low*, s is identified.
- If ADC$(s-1) >$ *high* and ADC$(s-2) \leq$ *low*, $s+1$ is identified.
- If ADC$(s-1) \leq$ *high*, $s+1$ is identified.

Here, ADC(n) is the ADC value for BC n, *low* and *high* are the thresholds, s is the first BC for which the ADC output is saturated and the integer increments are used to denote later or earlier BCs relative to s.

Figure 5.7 (b) shows two examples for the algorithm. The bottom pulse is barely saturated. As the samples $s-1$ and $s-2$ are above the *high* and *low* thresholds, respectively, the first saturated sample is correctly identified as the peak. For the larger pulse at the top, the sample before the peak is also saturated. Because $s-2$ is barely above the noise in this case, it is below the *low* threshold, leading to the algorithm correctly identifying $s+1$ as the peak.

External BCID

The External BCID (ExtBCID) algorithm is a further BCID mechanism implemented on the PPM [52]. A voltage comparator on the AnIn board compares the

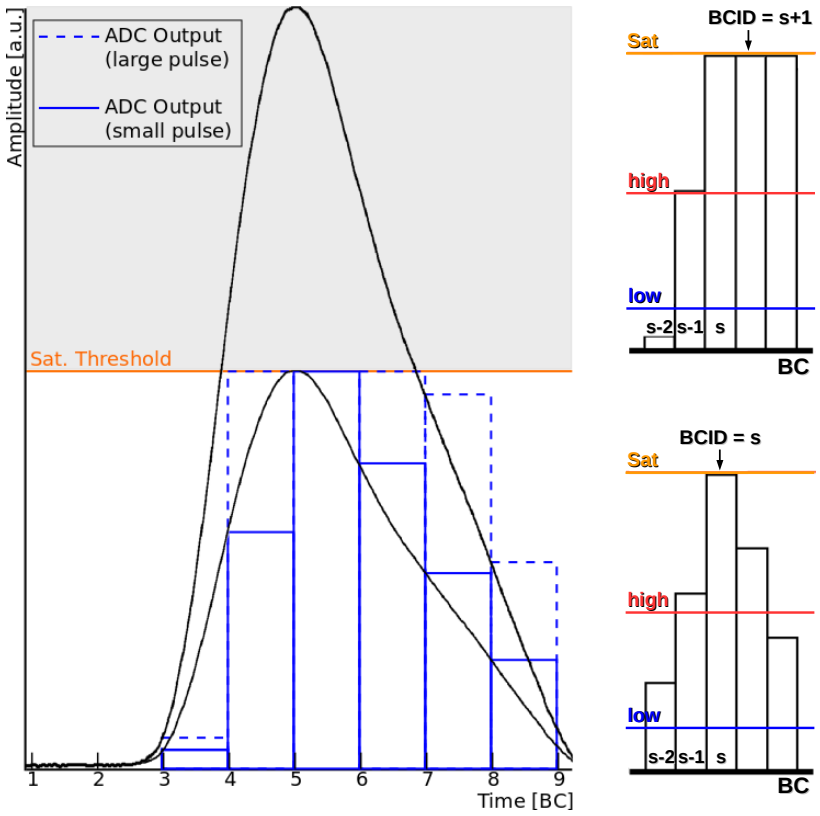

(a) Saturated Pulses **(b)** SatBCID

Figure 5.7: Functionality of the SatBCID algorithm. (a) shows two pulses of different energy, both of which saturate the ADC. (b) shows two cases of the SatBCID decision for pulses like those shown in (a).

analogue signal to a configurable threshold. Depending on the result, the comparator provides either a logic `1` or `0` to the CALIPPR FPGA, where it is received similarly to the ADC result.

The ExtBCID algorithm detects a change from `0` to `1` in the comparator output, indicating a rising edge of a pulse passing the threshold in the respective BC. Depending on a configurable delay, either this BC or a later one is identified by the algorithm.

In practice, the ExtBCID algorithm has never seen use in either Run-1 or Run-2. The implementation has been retained in the CALIPPR FPGA as a complementary backup solution.

BCID Decision Logic

The Peak Finder, SatBCID and ExtBCID algorithms operate in parallel for all BCs. For saturated pulses, all algorithms will identify a BC, though not necessarily the same one. The BCID Decision Logic resolves such situations by combining the

results of the different BCID algorithms to a single decision bit [52].

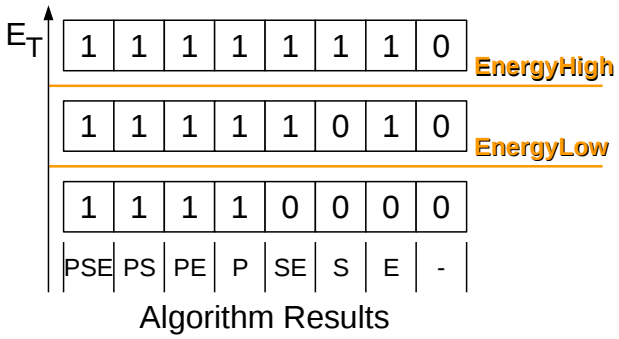

Figure 5.8: Example for configuration of the BCID Decision Logic.

Figure 5.8 illustrates the functionality of the Decision Logic with an example configuration. For each BC, the BCID Decision Logic receives the results of the Peak Finder, SatBCID and ExtBCID algorithms - abbreviated in the figure as P, S and E respectively - as well as both the 10-bit input and the 8-bit output of the LUT and the 10-bit ADC value. It then uses two configurable 10-bit energy levels to divide the energy range up until saturation into three regions - a low, a middle and a high energy region. At each BC the input energy is compared to the two thresholds to find the region the pulse belongs to - where the input energy is either the ADC value or the 10-bit LUT input, chosen again by configuration.

For each of the regions, a configurable register defines which combinations of BCID algorithm results give a decision bit of 0 and which combinations give a decision bit of 1. Eight different combinations of the three algorithms are possible, with e.g. PSE corresponding to P `AND` S `AND` E in the figure. Note that these are treated exclusively, i.e. if P and S are both active for a BC, only the PS configuration is considered by the decision logic, while the individual P and S configurations are ignored. In the example shown in the figure, only the Peak Finder result is used in the low energy region. The middle energy region includes the ExtBCID algorithm result, and the high energy region uses all three algorithms in parallel.

If a BC is identified by at least one of the BCID algorithms, and the respective algorithm is taken into account by the BCID Decision Logic, the BCID output is set to the energy of the pulse. Usually this is the output of the LUT. If the SatBCID algorithm identified the BC, the BCID output can be set to saturation instead, i.e. 255, independent of the LUT output value. If no BCID algorithm identifies a BC or the algorithm is ignored by the decision logic, the BCID output is set to 0.

Whenever a BC with a non-zero BCID output occurs, the BCID Logic sets the output to 0 for the following BC. As none of the BCID algorithms can identify two subsequent BCs, no information is lost in this operation. This resolves any potential disagreements between algorithms - the algorithm that identifies the earlier BC is chosen by the logic.

5.2.4 LVDS Transmission

The BCID results are transmitted to the CP and JEP systems via LVDS links. While this functionality was realised using commercial LVDS chips on the MCM, it is implemented in the CALIPPR FPGA on the nMCM.

Each nMCM provides three LVDS links - two to the CP and one to the JEP - which transmit the data at 480 MHz. The transmission protocol uses a high-state start bit and a low-state stop bit, leaving 10 bits for actual data content for each BC. The way this is utilised differs between the two subsystems, as is described in the following sections.

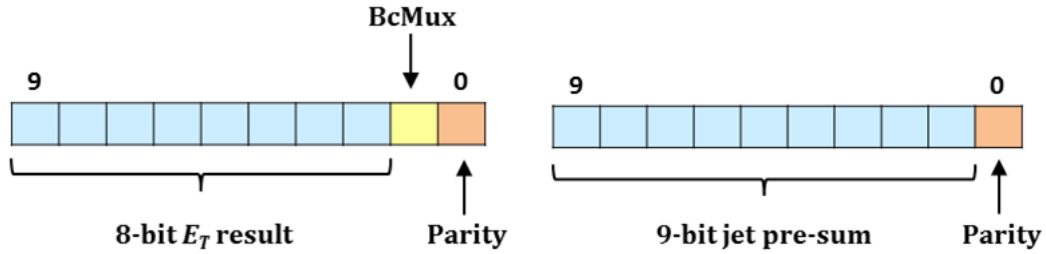

(a) Data Format to CP **(b)** Data Format to JEP

Figure 5.9: Format of data sent to (a) the CP and (b) the JEP systems via LVDS links. [53]

To CP: Bunch-Crossing Multiplexing

The CP requires the 8-bit E_T result for all trigger towers. This means that the BCID results of all four channels on the nMCM have to be transferred to the CP for each BC, which would require four physical links from each nMCM to the CP. This number of required links is halved by multiplexing the data of two channels into two subsequent BCs on the same link. This is made possible by the BCID Decision Logic: As described in the previous section, the BCID output is set to zero after each positive BCID. The resulting empty BC is used to send the data of the other channel.

This *bunch-crossing multiplexing* (BcMux) [52] is illustrated in Figure 5.10. In order to label the contents of which channel are sent in a specific BC, a BcMux flag bit is included in the data sent to the CP (see Figure 5.9 (a) for the data format). In the first BC of the multiplexing, the BcMux flag indicates which channel is being transferred in this BC. If the first channel is being transferred, the flag takes a value of zero, if the second channel is being transferred, the flag is set to one. If both channels contain energy in the same BC, the first channel takes priority and is sent first. In the second BC, the BcMux flag indicates to which of the two BCs the data transmitted in this BC belongs. If the data belongs to the previous BC, the flag is set to 0, otherwise it is set to 1.

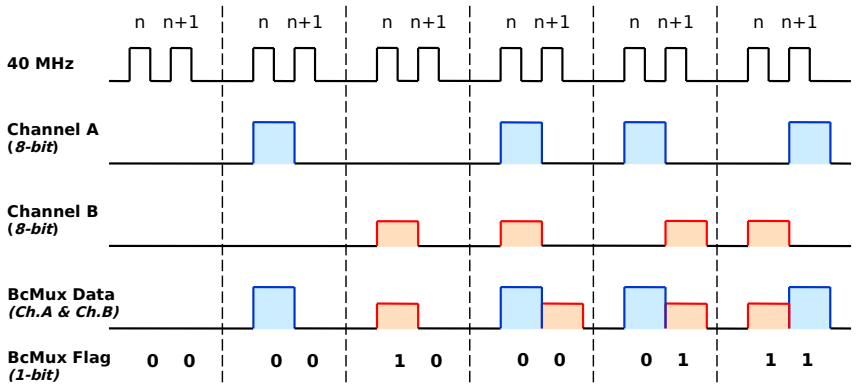

Figure 5.10: Data content and BcMux flag in the LVDS data format to CP for the different possible scenarios of Bunch-Crossing Multiplexing. Taken from [53].

The tenth bit in the LVDS data format is an 'odd' parity bit. This bit is set to 1 if the number of bits set to 1 in the remaining 9 bits is odd. This allows data integrity checks in the receiving system, which calculates the parity of the received data and checks it to be even.

For test purposes, a special configuration allows to discard the BcMux and parity bits and to send the 10-bit LUT input over the LVDS interface instead. This is referred to as the *BypassBcMux* mode.

To JEP: Jet Sums

The JEP algorithm does not use (0.1×0.1) trigger towers, but (0.2×0.2) jet elements. These are formed in the nMCM by summing the BCID results of the four trigger towers [52]. This is done in a two step procedure. In the first step, the 8-bit BCID results of two trigger towers are summed to produce a 9-bit JetPreSum. If one of the trigger towers is saturated, the JetPreSum is also set to its maximum possible value of 511, thus preserving the saturation. In the second step, the two JetPreSums are combined into a 10-bit JetSum, again setting the JetSum to the maximum of 1023 if one of the JetPreSums is saturated.

The 10-bit JetSum is further formatted before sending it to the JEP by dropping either the most or the least significant bit, again preserving the saturation. This reduces the JetSum to 9-bit, and allows the inclusion of a parity bit like in the CP case, as shown in Figure 5.9 (b).

5.2.5 Event Data Readout

Upon receiving an L1A signal, the nMCM provides data for the triggered BC to the ATLAS DAQ system. This data is crucial in order to monitor the system

performance, as it allows offline checks of the BCID decision on an event-by-event basis. Furthermore, the data is used to calibrate the PreProcessor system. For these purposes, the ADC input to the CALIPPR FPGA and the result of the BCID logic for each BC are continously collected in dedicated readout memory pipelines.

Each channel has two such memories, each of which stores 128 11-bit data words, where each word contains data from one BC [52]. One of the memories stores the 10-bit 40 MHz ADC inputs and the 1-bit ExtBCID result, as they are received at the input of the FPGA. The other memory stores the 8-bit BCID result and the three BCID result bits: the Peak Finder, SatBCID and ExtBCID. The memories are constantly cycled, such that after 128 BCs the memories are again filled from the beginning, overwriting the previous entries. This way, the information for each BC remains in the memory for 3.2 µs. This covers the 2.5 µs interval that is available to the L1 trigger to form its decision for each BC (see Section 3.4).

If an L1A is received, a configurable number of words is copied from each of the readout pipelines to another set of memories of equal size. These *derandomisers* serve as a buffering stage to protect the transmission of the event data from fluctuations in the L1A rate. Up to 127 ADC words and up to 7 BCID words can be selected for each event by the nMCM configuration. The information is then picked up from the derandomisers and transmitted to the ReM FPGA over the Serial Interface.

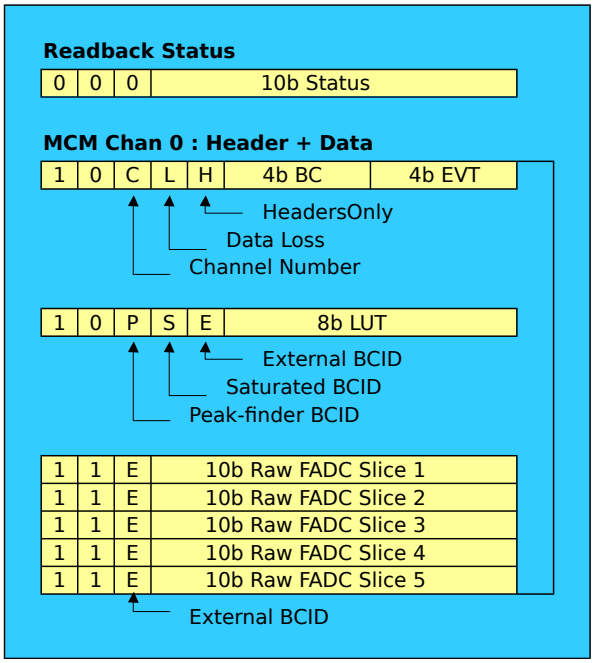

Figure 5.11: Readout Data Format. Adapted from [54].

The typical format of the readout data sent to the ReM FPGA is shown in Figure 5.11. The transmission of each event begins with a Header word. The lower eight bits of the header contain 4-bit BC and event numbers associated to the triggered

event. The nMCM keeps track of both of these numbers in internal counters. Upon receiving an L1A, the lower four bits of each counter are sent to the ReM FPGA, where they are checked against similar counters on the latter device. The remaining three bits of the header encode the channel number and readout status information. These status bits are active if the readout is missing information: If the derandomiser memories are almost filled, only the headers are sent to the ReM, while if the derandomisers are full, data is lost.

The remaining words contain the event data described before. First, the BCID words are picked up from the respective derandomiser and sent to the ReM FPGA. They are followed by the data stored in the ADC derandomiser.

While the nMCM allows for arbitrary numbers of readout words in the bounds mentioned above, the transmission from the ReM FPGA to the DAQ system is only possible in specific, predefined modes. The standard data taking mode, which is shown in Figure 5.11, uses five ADC words and one BCID word. According to the number of words, it is referred to as '5+1' mode. The recorded data in this mode allows to recalculate the filter output for the central of the five samples, and the LUT response to it, which is checked against the BCID output. This is a basic consistency check and is performed for each recorded event by the L1Calo monitoring software.

Other readout modes include the '15+1' mode, which allows the recording of a larger part of the digitised pulse, and the '7+1' mode, that is the minimal mode required in order to properly recalculate the Peak Finder algorithm. These modes are mostly intended for calibration purposes,

5.2.6 Monitoring

Besides the event data readout that provides the possibility to check the proper processing of triggered events, the nMCM also offers two further, more general monitoring tools. These are the Rate-Metering and the Histogramming.

Rate-Metering

The Rate-Metering [52] measures the rate at which energies above a configurable threshold are detected in a trigger tower channel. It consists of two counting registers - a 16-bit time counter and a 20-bit data counter - which start incrementing once the Rate-Metering operation is enabled. The time counter is used to record the time that has passed since the start of the Rate-Metering operation, in units of TimeBlocks, where one TimeBlock equals 1024 BCs. The data counter is used to record the number of BCs for which the energy input of the Rate-Metering exceeds a configurable 10-bit threshold. The energy input can be either the 10-bit ADC output or the 8-bit output of the BCID logic. If the BCID output is selected as energy input, two zero bits are appended to obtain the required 10-bit resolution.

The Rate-Metering operation concludes once one of two conditions is met - either the data counter saturates, or the time counter reaches a configurable threshold. In both cases, the values of the time and data counters are copied to two registers, where

they are available for readout by the ReM FPGA. The Rate-Metering operation then restarts, and both counting registers are reset to zero.

Histogramming

The Histogramming algorithm [52] produces histograms based on either the 10-bit ADC input or the 8-bit BCID result. It uses a memory that can store 256 11-bit items, where each item represents one bin in the histogram, allowing up to 2047 counts per bin. The definition of the bins depends on the chosen input: while the 8-bit BCID output can be mapped 1-to-1 to the memory locations, the 10-bit ADC input has to be truncated by dropping 2 bits. In this case the algorithm allows for three different configurations, dropping either the two least significant bits, the two most significant bits or both the least and the most significant bit.

The resulting 8-bit input is compared to a programmable threshold. If it exceeds the threshold, the respective bin in the histogram is incremented by one count. This operation continues until one of the bins reaches saturation, at which point the Histogramming operation is stopped, and a signal is activated indicating the availability of the finished histogram. The memory contents can then be retrieved by the ReM FPGA.

5.2.7 Testing Facilities

The nMCM provides two internal testing facilities that can be used to check the functionality of various parts of the signal processing. The first of these is a Playback memory in the CALIPPR FPGA, allowing tests of the digital signal processing in the FPGA. The analogue processing on the nMCM can be tested with a Signal Generator circuit.

Playback

The Playback [52] is based on a memory that stores a pattern of 256 11-bit words, where each word corresponds to one 10-bit ADC value and one bit for the ExtBcid result. By operating the nMCM in Playback mode, the 40 MHz ADC result is discarded at the input of the CALIPPR FPGA, and the output of the Playback memory is instead used as input to the processing algorithms. The read address of the memory is incremented each BC, so that the Playback output continuously cycles through the stored pattern. Depending on the configuration, the Playback is operated in either a oneshot mode, where the memory is cycled exactly once, or a continuous mode, where the Playback starts again from the first entry once the last entry has been reached.

The Playback uses the same memory as the Histogramming functionality described in the previous section. This is possible because the two functionalities serve exclusive purposes: The Playback is intended purely for testing, while the Histogramming is only required in data taking runs.

Signal Generator

The Signal Generator circuit provides the possibility to generate analogue inputs to the ADCs. It is a new feature of the nMCM, having been added to the design to enable tests of the analogue processing of the module independently of the input from the calorimeters.

The Signal Generator circuit for one nMCM channel is shown in Figure 5.12. The main component of the circuit is a capacitance, which is connected to five output pins of the FPGA with different resistors in between. Four of these pins (the top four in the figure) can be set to either logic 0 or 1, which result in an effective decrease or increase of charge in the capacitance, or to a high impedance state (TriState), that keeps the charge constant. For two of the pins, the resistance is only 10 Ω, resulting in a relatively big impact on the resulting signal. The other two resistors, set to 220 Ω and 330 Ω, decrease the effect of the respective pins. These are mostly used to fine-tune the produced signal.

The logic states applied to each of these pins are controlled by a configurable memory. It contains 4096 8-bit words, where each word represents the states of the four pins. The state of each pin state is represented by two bits, where one indicates the logic state ('1' or '0') of the pin, and the other indicates the tristate configuration ('on' or 'off'). The memory is operated at 120 MHz so that each entry corresponds to 1 third of a BC, about 8.3 ns. Similar to the Playback, this can be operated in either a oneshot mode or a continuous mode.

The fifth pin from the top in Figure 5.12 is used to equalise the leakage currents in the circuit. These occur when all the other pins are set to TriState mode and lead to a drift of the signal baseline. The logic value of the pin is set to a configurable constant value, while the TriState is controlled according to a configurable 24-bit pattern, that is cycled in regular intervals. Whenever the pattern has been cycled once, a configurable delay has to pass before the pattern is applied again.

A further pin allows the generation of an additional DC type signal that is added to the signal shape produced by the circuit described above. This allows to control the pedestal of the analogue signal, similar to the DAC on the AnIn boards. The logic value of the pin is controlled by a 12-bit register, that defines for which fraction of time the output is set to a value of zero. An internal 12-bit counter, incremented at 80 MHz, is compared to this register at each clock tick. If the counter is smaller than the register value, the output of the pin is set to zero, otherwise it is set to one. Increasing the register setting thus increases the amount of current that is drawn from the pedestal circuit, thus lowering the signal. As the output of this circuit is applied to the negative input of the amplifier when adding it to Signal circuit output, a higher register setting corresponds to higher pedestal value.

5.2.8 Clock Management

The clock signals driving the processes in the CALIPPR FPGA are generated by internal Phase-Locked Loops (PLLs). These components consist of variable frequency

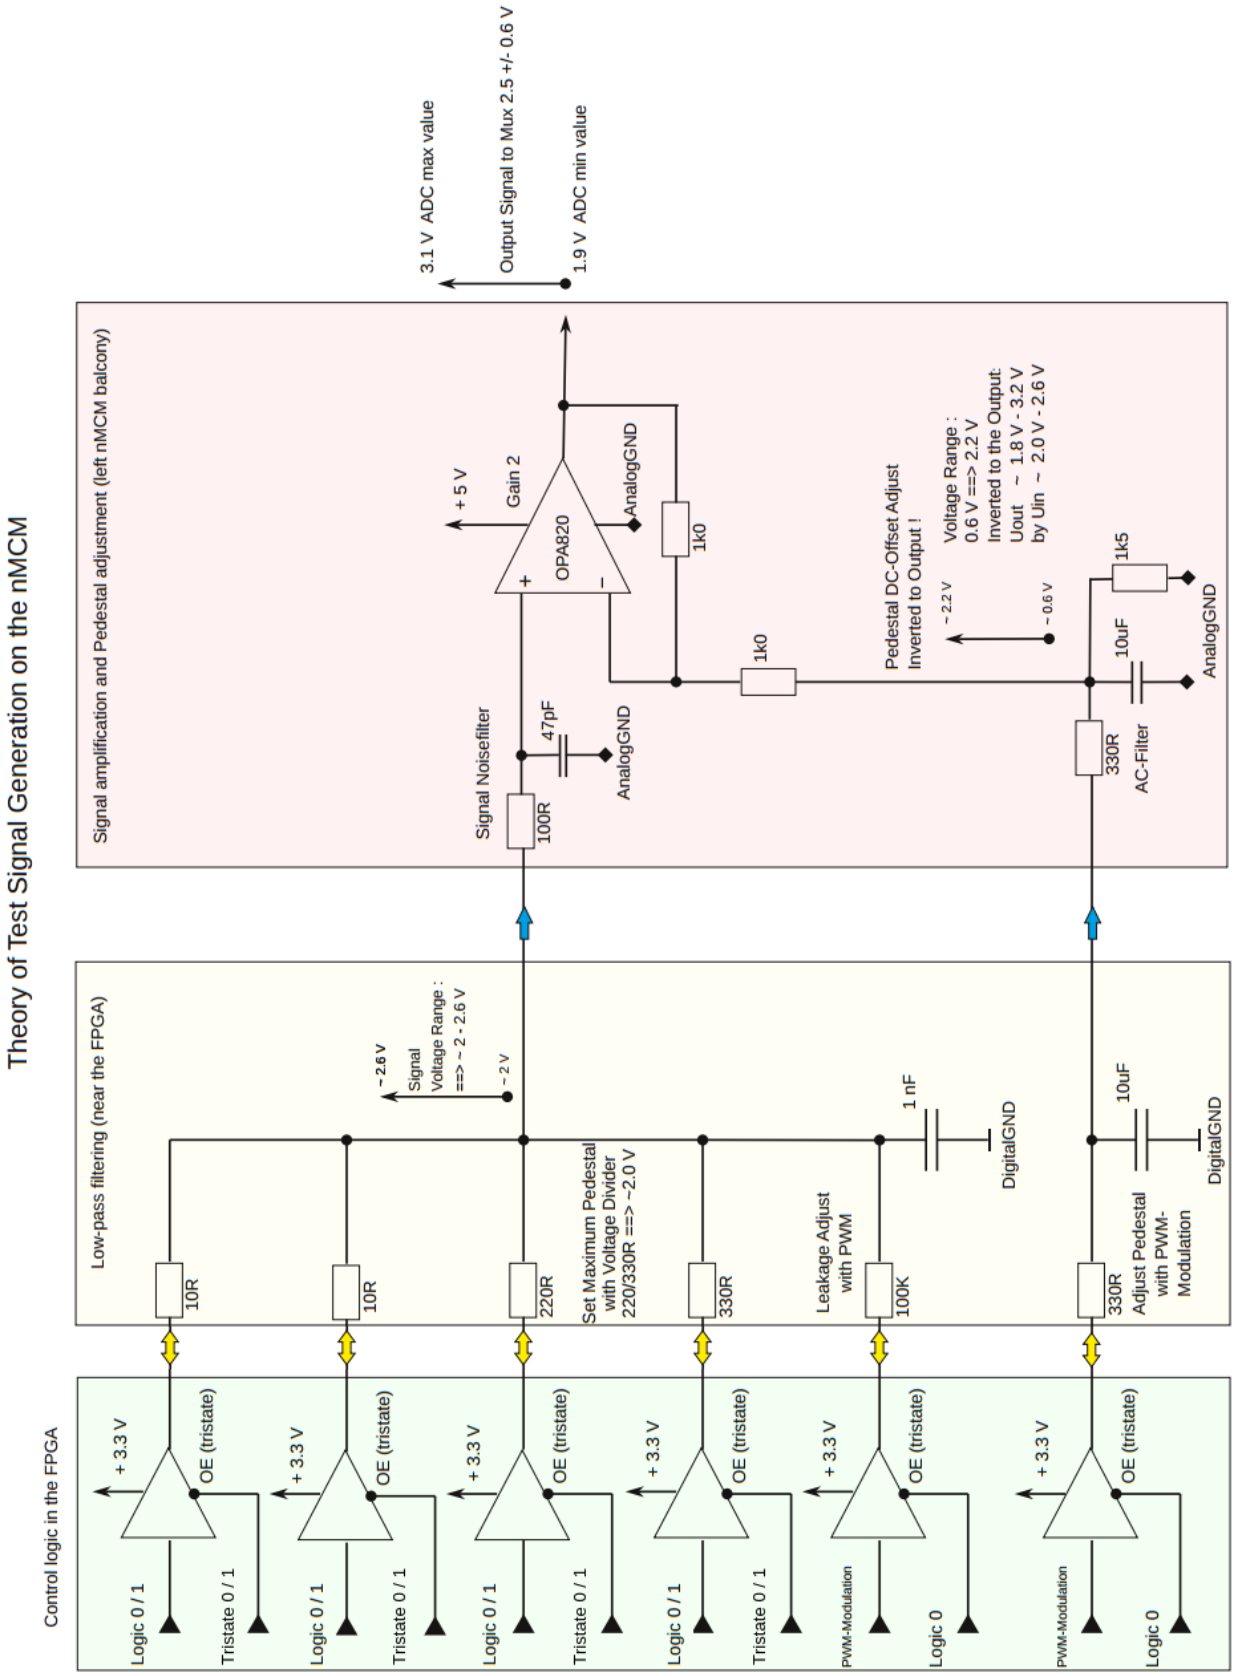

Figure 5.12: Signal Generator circuit on the nMCM. Adapted from [55].

oscillators coupled to phase detectors, such that the signals generated by the oscillators are in phase with the external signal. They also allow for frequency division, generating faster or slower clock signals from the input.

The internal clock generation is a new feature for the nMCM - while the MCM received several different clock signals from the ReM FPGA, only a single 40 MHz signal is routed to the nMCM. This reduces the cross-talk on the PPM, resulting on lower noise on the trigger tower input signals for the nMCM compared to the MCM (see Section 9.1)

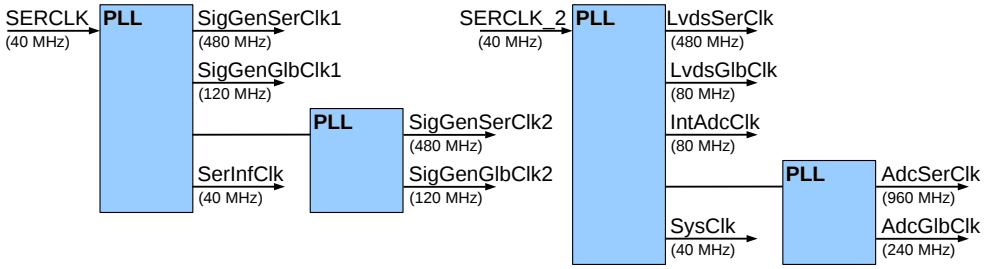

Figure 5.13: Overview of the clock signals generated in the CALIPPR FPGA.

Figure 5.13 shows the clock generation scheme in the CALIPPR FPGA. The 40 MHz LHC bunch crossing clock, that is provided to the PPM by the TTC system is forwarded to the nMCM by the ReM FPGA, and is split into two copies close to the input connector of the CALIPPR FPGA. The two resulting input clocks, SERCLK and SERCLK_2, are used as inputs to separate PLLs.

The SERCLK is used to generate the clocks for the Signal Generator circuit and the Serial Interface. The SERCLK_2 is used to generate the clocks for the LVDS transmission, the clocks required to generate the ADC strobes, the internal 80 MHz clock for the latching of the input ADC data and, most importantly, the internal 40 MHz clock for the majority of the signal processing algorithms.

5.2.9 Configuration

The PreProcessor processes data from various different calorimeter partitions while using equivalent pieces of hardware. To accomodate the differences in signal production, processing and transmission between the calorimeter partitions, all the PreProcessor components are highly configurable in regards to their signal processing parameters.

The nMCM configuration is realised from VME via the ReM FPGA and two serial interfaces. These give access to a set of channel-wise configurable registers and memory locations in the CALIPR FPGA. Additionally, an Inter-Integrated Circuit (I2C) interface is implemented to configure the fine timing functionality. This separate way of configuration has been inherited from the MCM design, where this task was performed by the PHOS4 chip, which required the I2C implementation in order to be accessed.

Serial Interface

The Serial Interface transfers data between the nMCM and the ReM FPGA [52]. Its main uses are the configuration of the nMCM and the transfer of event data from the CALIPPR FPGA to the DAQ system.

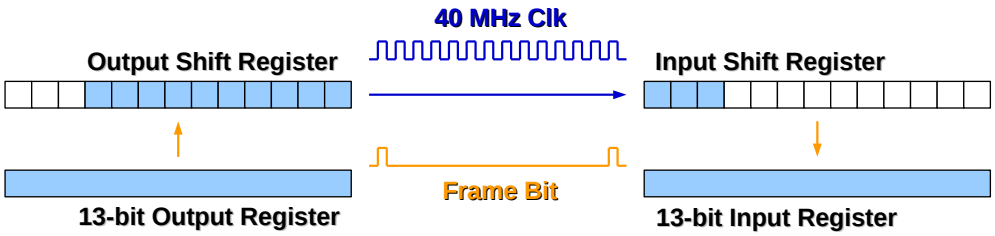

Figure 5.14: Schematic representation of the Serial Interface.

Figure 5.14 illustrates the functionality of the Serial Interface. It is operated at a frequency of 40 MHz. Data is transmitted in frames with a length of 13 BCs. The ReM masters the transfer by providing both the 40 MHz clock and the frame bit signal. In each frame, signalled by the frame bit, the transmitter copies a 13-bit data word from the processing logic to an output shift register. The content of this output register is transferred to an input shift register on the receiving end serially at 40 MHz. With the next frame bit, the content of this register is copied to a 13-bit input register. Each nMCM is equipped with two such interfaces to the ReM FPGA, each of which services a group of two channels.

On the input side, the Serial Interface receives configuration data and commands from the ReM FPGA. The commands initiate the different monitoring processes or activate the Playback mode.

On the output side, the Serial Interface is divided into two channels. The Readout channel is used to transfer the event data to the DAQ when the CALIPPR FPGA receives an L1A. The Readback channel is used to transfer the contents of the configuration registers and memories as well as the results of the monitoring algorithms if an appropriate command from the ReM FPGA is received. If both channels are active at the same time, the Readout channel takes priority. At very high L1A rates, only a minimum of one Readback word is sent between events, which is required to separate the Readout data of the different events on the interface.

If no data has to be transmitted over the Readback channel, it is instead used to send status data to the ReM FPGA. It indicates the state of various internal processes in the CALIPPR FPGA, e.g. the Rate Metering and Histogramming algorithms (see Section 5.2.6).

The format of the status word has been reworked on the nMCM, to the state that is shown in Figure 5.11. In order to reduce cross-talk between the Serial Interface and the trigger tower signals, the first three bits, that function as an identifier, were changed from `001` to `000`. In the idle state, the Serial Interface output is thus

constantly set to `0`. This increases the noise performance of the nMCM compared to the MCM (see Section 9.1).

6 Functional Tests of the new Multichip Modules

The nMCMs were produced in February 2013. The PreProcessor system consists of 124 PPMs, which require a total of 1984 nMCMs. To provide a sufficient amount of nMCMs as potential replacement for malfunctioning modules, a spare policy of 25% was adopted[1]. A conservative approach was chosen to guarantee that this number of working spares is achieved, so that 3000 nMCMs were produced in total.

To sort out any faulty modules before installation, all modules were thoroughly tested for proper functionality at the KIP in Heidelberg, as described in this chapter. Section 6.1 provides an overview of the general test strategy, while the individual functionality tests are described in section 6.2. The results are summarised in section 6.3, and the installation of the modules in the PreProcessor system at CERN is described in section 6.4.

6.1 Test Overview

The strategy of the nMCM functionality tests is based on the one applied for the tests of the PPMs performed in 2007, described in [44]. The tests are ordered into four main stages, which focus on different aspects of the module functionality.

1. The first stage serves to assure that the nMCMs can be safely operated in the PreProcessor crates. The first step is an optical inspection to find any component soldering errors. This is followed by a DC power-up test, in which the nMCMs are mounted on a PPM that is then powered up. If an nMCM is not operating properly, it is detected because the total current drawn by the PPM is lowered. Afterwards, the voltages on the nMCM are checked for their correct magnitude by using a multimeter. If an nMCM passes all these steps, it is equipped with a heat sink and continues to the next stage.

2. The second stage serves to bring the module into an operating condition. Once a set of 16 nMCMs has passed the first test stage, they are mounted on a PPM that is the platform for all following tests. By testing mltiple nMCMs in parallel, the total time required for the test procedure is reduced significantly. The PPM is then inserted into a single-board VME crate setup (see Section 6.2.1). Once the setup is complete, the CALIPPR FPGAs are then programmed over the JTAG interface for the first time. This enables

[1]The same spare policy was already adopted for the MCMs and the PPMs before Run-1

the serial interface the the ReM FPGA, which is used in the next step to transfer the programming file to the EEPROM. Provided no errors occur in this procedure, the ID number corresponding the the mounted heat sink is loaded to the EEPROM as well, enabling the identification of the module.

3. The third stage tests all functionalities that the nMCM has to provide in the PreProcessor system. It consists of a series of tests in the single-board setup, each of which is focussed on one aspect of the nMCM functionality. Detailed descriptions of these tests are given in Section 6.2.

4. The final stage checks the long-term stability of the nMCMs by placing them in a standard VME crate. The crate is fully equipped with PPMs, all of which are loaded with nMCMs. The nMCMs repeat the same tests as in the previous stage, but remain in the crate for a longer period of time, typically one week. After this amount of time is passed, the module is again subjected to all the tests from the previous stage.

Depending on the outcome of the tests, the modules are ordered into several categories:

- nMCMs that pass all tests make up the *Good* category. These modules are used to equip the PreProcessor system and the test setups, in that order of priority. Remaining modules are used as spares.

- nMCMs that fail at least one of the tests are ordered into the *Faulty* category. These modules can not be used in the PreProcessor system. If possible, attempts to repair the modules to recover the functionality will be made, at which point the modules will be re-evaluated.

- For some tests, the definition of what constitutes a failure is not a clear cut. For example, the exact level of noise at which a module is considered 'too noisy' is hard to define. nMCMs that fail such a kind of test are instead ordered into an *Usable* category of modules. The modules from this category can be used as spares, to be decided on a case-by-case basis.

6.2 Functionality Tests

This section provides descriptions of the functionality tests that are performed in the the third and fourth stage described above.

Each of the tests checks one of the functionalities of the nMCM. The tests build upon each other, so that subsequent tests can make use of the functions that have been established in previous tests. The first tests focus on the configuration of the module and the Monitoring and Playback functions. Using the latter, the algorithmic processing in the CALIPPR FPGA and the transmission of the results, both the E_T data to the CP and JEP systems and the event readout information to the

ReM, are checked. Once this digital functionality is established, the analogue input is tested by analysing the performance of the ADC and the proper operation of the input stage of the CALIPPR FPGA.

6.2.1 Test Setups

The tests use two different setups. The single-board setup is used to establish the functionality of all parts of the nMCM. It is used for both the second and the third stage of the tests. The full-crate setup is used to test the long-term stability of the nMCM when put into operational conditions similar to those in the PreProcessor system at CERN. The following two sections provide descriptions for both of these setups.

Single-Board Test Setup

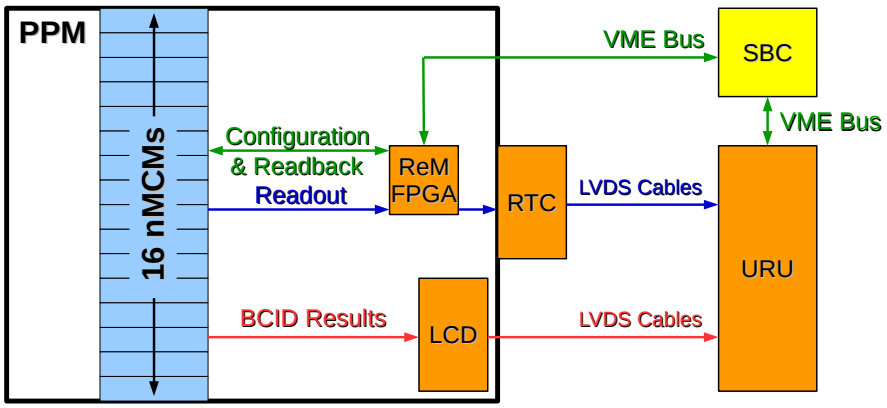

Figure 6.1: Setup of the Single-Board Tests.

The single-board tests are performed in a custom-built VME crate [44] that is equipped as sketched in Figure 6.1. It contains a single PPM, which serves as the test platform for 16 nMCMs at a time. A custom-built Single-Board Computer (SBC) functions as crate controller and VME master.

The crate also houses one *Universal Receiver Unit* (URU), a custom-built VME board that is connected to the SBC as well. It mimics the input stage of the ROD and CP/JEP systems, neither of which is available in the test lab at KIP.

The LVDS signals intended for the CP/JEP are transmitted from the PPM to the URU over 15 m long LVDS cables. This is longer than the cable length of 11 m used in the L1Calo system, and was chosen to make a conservative test of the signal stability. The data sent from the PPM to the ROD is transmitted via an additional *Readout Transfer Card* (RTC), which transforms the data stream to LVDS type signals and sends it to the URU. Both data types are buffered on the URU and made available to the SBC via the VME.

Full-Crate Test Setup

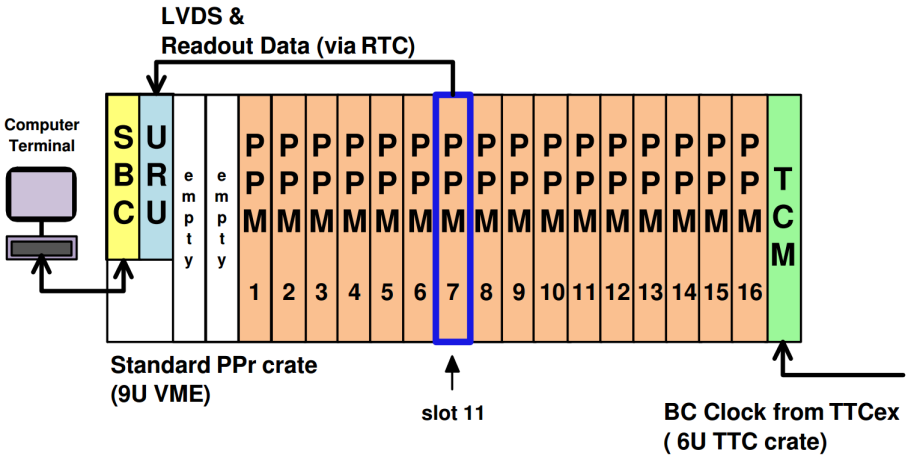

Figure 6.2: Setup of the Full-Crate Tests. Taken from [44].

The Full-Crate tests are performed in a standard PreProcessor VME crate. As shown in Figure 6.2, the crate houses one SBC, 16 PPMs as well as one set of CMCs and an RTC. Additionaly, a TCM connects the PreProcessor crate to a TTC crate, which is used as the source of the 40.08 MHz clock.

The full-crate test is aimed at testing the long-term stability of the nMCM performance. For this reason, only a single PPM slot (PPM 7 / VME slot 11 in the figure) is under active testing at a time. This is the only slot equipped with an RTC and connected to the URU. The remaining nMCMs are configured in the same way and perform the same operations as the ones in the testing slot, but no data is recorded from them.

The PreProcessor crate is a dense environment, and proper heat dissipation is a key challenge for the nMCMs. Figure 6.3 shows the typical temperature distribution of the nMCMs in the crate. The observed structure is a result of the placement of the cooling fans relative to the module position. The air flows from the bottom to the top, resulting in higher temperatures for modules further up in the crate. On the horizontal axis, PPMs in slot 5 and 6 (6 and 7 in Figure 6.2) are the hottest, because they are placed in between two of the cooling fans. For this reason, PPM slot 6 (7) was chosen as the active testing slot, as it puts the largest strain on the nMCMs.

To ensure that all nMCMs remain in the crate for about the same amount of time, a rotation system was adopted. Each time the test procedure finishes, the PPM is removed from the active slot. The PPM that has been in the crate for the longest amount of time is then moved into the active slot, and the resulting empty slot is filled with another PPM that is coming out of the single-board test.

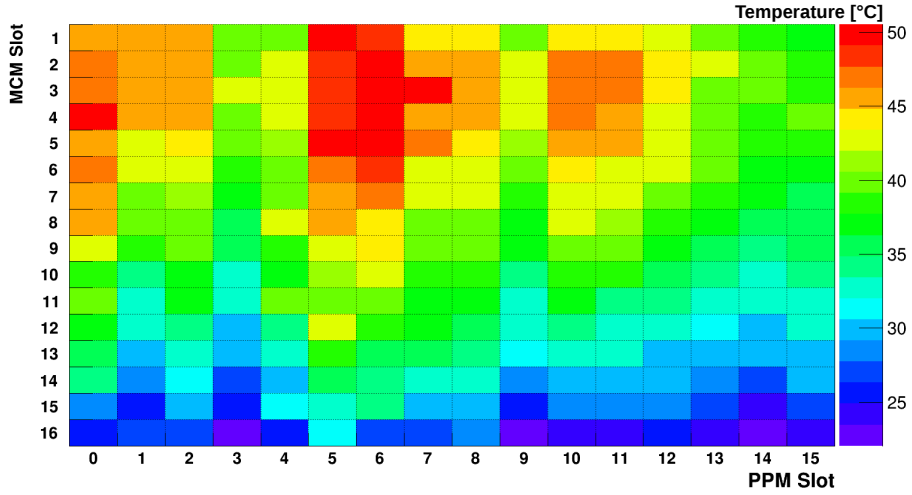

Figure 6.3: Typical temperature distribution of nMCMs in a fully equipped PreProcessor crate.

Simulation Software

Several tests involve the algorithmic processing of known input data in the CALIPPR FPGA. In order to calculate the expected outcome of each test, a suite of simulation software has been included in the test tools. This simulation consists of parts of the L1Calo test software used at CERN. Only the minimum amount of software required to simulate the PreProcessor BCID algorithms has been adapted, dropping the simulation of the other L1Calo processors. This provides flexibility to the testing tools, as the input data can be changed to any arbitrary pattern without requiring lengthy calculations to obtain the expected output.

Logging of Test Results

The results of all individual tests described in this section are stored in automatically generated test logs. For many tests, further information is stored in the form of plots in ROOT files. Once the full suite of tests is completed for the nMCMs on one PPM, a software tool analyses all the log-files to generate summary logs, each of which contains the information of how one nMCM performed in the test. If the nMCM does not pass a specific test, more detailed information regarding the detected failure is also recorded. The log files are produced independently for both the Single-Board and the Full-Crate tests.

All logs and ROOT files are uploaded to a central database, storing the module data of all L1Calo components [56]. Software tools running on the server analyse each uploaded log file for the test result, automatically categorising the nMCM according to their test result. In case of nMCMs not passing a test, the software allows registered users to enter more detailed comments regarding the specific failure of the module. The database remains in use also after the end of the functionality

tests and the nMCM installation, as the locations of all modules and any changes in status are tracked.

6.2.2 Configuration Test

This first test checks whether all configuration registers, as well as the various configurable memories, can be accessed by both read- and write operations. The memories checked by this test are the Playback memory, the two energy calibration LUTs, the EEPROM buffer and the Signal Generator buffer.

The test consists of loading predefined patterns into the registers or memories and subsequently reading back the values stored in them. Any mismatch between the value that is sent to the module and the value that is received upon readback points to an error in either the storage of the values in the CALIPPR FPGA, or in the transmission over the serial interface between the CALIPR and ReM FPGAs.[2]

Name	Bit Pattern	Configuration Test Type
Default	(default configuration)	Registers
Ramp	(increasing ramp)	Playback Memory
all 1's	11111...	Registers & All Memories
all 0's	00000...	Registers & All Memories
Alternating 10's	10101...	Registers & All Memories
Alternating 01's	01010...	Registers & All Memories

Table 6.1: Patterns used in CALIPPR configuration tests.

Table 6.1 lists the patterns used in the configuration tests and which memories are tested using each of them. The first pattern in the list is the 'Default' configuration of the registers. In this test scenario, the registers are left in the default state they aquire after a reset of the CALIPPR FPGA. The second pattern, only used to test the Playback memory, is an increasing ramp, i.e. each entry in the memory is incremented by one compared to the previous entry in the memory, starting from zero. This pattern is used in several of the subsequent tests that make use of the Playback functionality. The remaining patterns use repeating sets of bits to fill up the respective registers or memory entries.

All nMCMs passed this test without a single error being observed. The configuration procedure for the nMCMs is thus ruled out as a source of error for all the following tests.

[2]The transmission over VME to the ReM FPGA has been validated extensively, both by the production tests of the PPMs and by years of operation at CERN, and is thus assumed to function properly for the purposes of this test.

6.2.3 Rate-Metering Test

The second test checks the Rate-Metering functionality of the nMCM (see Section 5.2.6).

To test the Rate-Metering, the CALIPPR FPGA is operated in Playback mode. In this setup, the rate of samples above the Rate-Metering threshold is proportional to the number of samples in the Playback pattern that are above the threshold. The test consists of measuring the rates resulting from a given pattern when applying different thresholds and checking the result against the expectation.

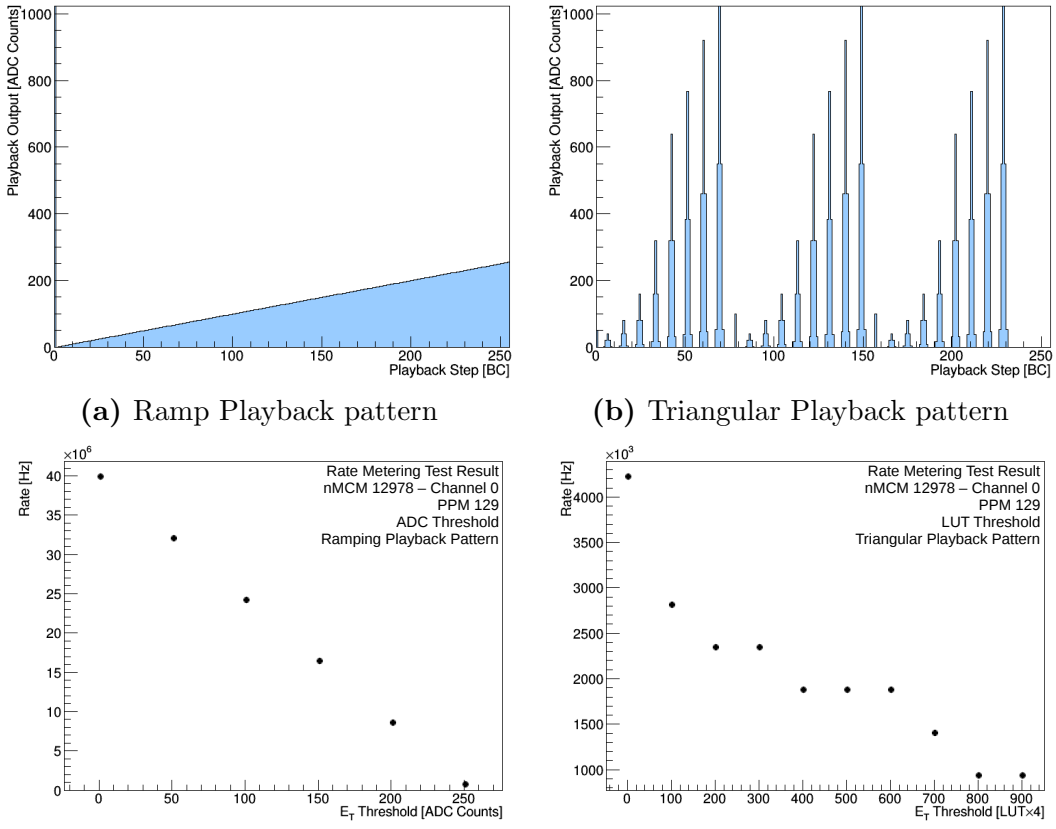

(a) Ramp Playback pattern

(b) Triangular Playback pattern

(c) ADC rates using ramp pattern

(d) LUT rates using triangular pattern

Figure 6.4: Playback patterns used in the Rate Metering test and the corresponding results. The top two plots show the ramping pattern (a) and the triangular pulses pattern (b). The lower two plots show rate metering results obtained for ADC-type input on the ramping pattern (c) and BCID-type input on the triangular pulses pattern (d).

The Rate-Metering is tested for both ADC and BCID-type energy input. To test the ADC input functionality, a ramping pattern covering all values from 0 to 255 ADC Counts is loaded to the Playback. It is shown in Figure 6.4 (a). In this case, the number of samples above the threshold decreases linearly with the threshold, resulting in a linearly falling behavior of the rate as a function of the threshold.

This is confirmed by the test results, shown in Figure 6.4 (c). The six data points, taken at thresholds from 1 to 251 ADC counts in equidistant steps, are equally far apart in resulting rate.

When testing the Rate-Metering using the BCID-type energy input, the ramping pattern is less useful, as only samples that give a positive BCID result are considered in this case. The ramping pattern only contains one peak, the largest sample at the end of the ramp, making a BCID-based Rate-Metering test on this pattern trivial. A different pattern is used instead, that consists of a series of triangular pulses of different amplitudes, as shown in Figure 6.4 (b). With this pattern, the rate as a function of the threshold is not proportional to the number of samples above the threshold, but to the number of peaks above the threshold, resulting in a less linear behaviour of the rate versus the threshold. The test result, shown in Figure 6.4 (d), confirms this expectation.

The measured rates are compared to the expected results to determine whether the nMCM passes the test. The Rate Metering provides two results, the number of TimeBlocks, $N_{\text{TimeBlocks}}$, and the number of counts above threshold, N_{Counts}. In Playback mode, these numbers depend on the number of samples above the threshold in the pattern, N_{Playback}:

$$N_{\text{Counts}} = 4 \cdot N_{\text{TimeBlocks}} \cdot N_{\text{Playback}}. \tag{6.1}$$

The factor of four appears because one TimeBlock is 1024 BCs long, thus covering the Playback pattern exactly four times.

For the ADC-type input, N_{Playback} is equal to the number of samples above the threshold in the Playback pattern. For the BCID-type input, the pattern is first analysed using the simulation to obtain the expected output of the BCID logic. N_{Playback} is then equal to the number of BCID outputs above the threshold.

All nMCMs passed this test.

6.2.4 LVDS Transmission Tests

This test serves to check the transmission of data to the CP and JEP systems over the LVDS links. On the MCM the data was serialised and transmitted by dedicated LVDS Tx chips, while on the nMCM the CALIPPR FPGA implements these tasks.

The test consists of loading specific Playback patterns to each channel of the module and reading back the data that is received by the URU after the digital processing in the CALIPPR FPGA logic and the LVDS transmission. The data is then analysed on the SBC.

Several different tests are performed for both the links to CP and to JEP, which use different playback patterns tailored to the different purposes. Table 6.2 provides an overview of the different tests and which patterns they use.

The Ramp and Triangular Playback patterns are the same ones used in the Rate-Metering test, shown in Figure 6.4 (a) and (b), respectively. In addition to the patterns used in previous tests, this test also uses a saturated and a stress pattern.

LVDS Test Type	Playback Patterns
CP-10-bit	Ramp, Stress, Triangular, Saturated
CP-BcMux	Triangular, Saturated
CP-Parity	Triangular, Saturated
JEP-Sums	Triangular, Saturated
JEP-Parity	Triangular, Saturated

Table 6.2: Playback patterns used in LVDS transmission tests.

The saturated pattern is similar to the triangular pattern, but all test pulses include at least one saturated sample. This way, the saturated BCID logic is included in the test. The stress pattern is generated similarly to the patterns used in the configuration tests, by filling the playback memory with regular patterns of 1's and 0's.

- The **CP-10-bit Test** is performed in BypassBcMux mode, i.e. the BCID and BcMux algorithms are ignored and the input data, in this case the output of the Playback, is sent directly via the LVDS links (see Section 5.2.4).
- The **CP-BcMux Test** uses the normal operation mode, i.e. the BCID and BCMux algorithms are active. This requires Playback patterns that produce non-zero BCID results. The simulation is then used to compute the expected result.
- The **CP-Parity Test** is performed directly on the URU. It computes an odd-parity bit for each of the 10-bit data words it receives from the PPM and checks the result against the parity bit that is included in the data. Mismatches are recorded by incrementing dedicated counters that can be read out via VME.
- The **JEP-Sums Test** verifies the different truncation modes that are available in the construction of the JetSums. It also makes use of different delays in the Synchronisation FIFO, to increase the number of possible different JetSums that can be generated, by delaying data on individual channels.
- The **JEP-Parity Test** is equivalent to the CP-Parity Test - the CMC card checks the parity of the received data to be correct.

Using this test, one nMCM was found to be faulty. It showed no errors in the single-board test, but parity errors were observed for the JEP link in the full-crate setup. The long operation time at high temperature worsened the performance significantly in this case. The board was subsequently marked as 'Faulty'.

6.2.5 Event Readout Test

This test checks the readout of event data from the nMCMs. It uses the Playback to generate a known input and compares the measured output to the expected result,

which is obtained from the simulation.

The Readout is initiated by the ReM FPGA. Upon receiving the appropriate command from the SBC via VME, it issues a PPM internal L1A signal, that is distributed to the nMCMs. The nMCMs subsequently transmit event data to the ReM FPGA where it is collected and formatted.

The data can be accessed in two ways. The first is the standard data transmission to the DAQ system, which in this case means that the data is transmitted to the URU and retrieved from there by the SBC. Secondly, the ReM FPGA provides a set of *Spy Buffers*, where it makes the data it retrieved from the nMCMs available for transmission the the SBC via VME.

Both types of transmission are used for the test procedure. Once the data is collected on the SBC, the ADC data words are compared to the loaded Playback pattern, and the BCID words are checked against the results of the simulation.

The test is repeated for several different readout modes, as listed in table 6.3. For the DAQ readout path, all possible configurations are tested. For the VME path, additional modes are included to test up to the maximum possible 127+7 mode.

Readout Mode	Readout Test Type
3+1	VME & DAQ
5+1	VME & DAQ
7+1	VME & DAQ
9+3	VME & DAQ
11+5	VME & DAQ
15+1	VME & DAQ
15+5	VME
127+7	VME

Table 6.3: Readout modes tested in the Event Readout test.

Once the data has been read out to the SBC, it is compared to the input data. As the command to issue the L1A is sent by the SBC, the readout occurs at an essentially random time. In the first step of the data analysis, the ADC samples in the readout are thus checked against the Playback pattern to locate the position at which the L1A occured. For this to be successful, it is crucial that all pulse shapes used in the pattern are unique, to avoid ambiguities. If the pattern of ADC samples cannot be found in the Playback pattern, the readout is marked as corrupt and the nMCM fails the test. If this syncronisation is successful, the response of the BCID logic for the relevant samples is extracted from the simulation results and compared to the BCID results included in the readout. Any disagreement in this comparison also results in the nMCM failing the test.

No nMCM was found faulty in this test.

6.2.6 Analogue Input Tests

This set of tests ensures the correct operation of the analogue signal processing components on the nMCM. This is done applying DC type input signals to the nMCM, which pass through the OpAmps and are digitised by the ADCs. By collecting the digitisation results that are received by the CALIPPR FPGA, the linearity and noise performance of the analogue components can be analysed.

In order to perform these tests, an analogue input signal to the nMCMs is required. For the purpose of these tests, the 8-bit DAC on the AnIn board is used to provide DC inputs of different values (see Section 4.3.2). Its step size of about 1.64 mV corresponds to roughly 2 ADC Counts per step. It is able to roughly cover the lower quarter of the dynamic range of the ADCs.

An additional method to obtain a DC type input is provided by the pedestal generation circuit of the Signal Generator (see Section 5.2.7). In total, 4096 different pedestal settings are possible, with an increase of three counts in the pedestal setting resulting in a 1 ADC Count increase in the measurement. While this method covers the whole dynamic range of the ADCs, the DC level is less stable and the output less linear than that of the DAC.

The following sections discuss tests using these two methods.

DAC Scan

The DAC Scan tests the response of the ADC to different settings of the DAC installed on the AnIn board. This is done by recording a programmable number of ADC values for each possible DAC setting, using the Spy Buffer method that was already used in the Event Readout test. To avoid any bias coming from operating the DAC, the different DAC settings are investigated in a random order.

The result of a DAC Scan for one nMCM channel is shown in Figure 6.5. It shows the mean ADC output as a function of the DAC setting. For DAC settings below 124 the ADC output is stable at 0, as the resulting DC level is outside the dynamic range of the ADC. The scan reaches values up to just above 250 ADC Counts, covering a bit more than one quarter of the ADC range.

A linear fit is applied to this measurement, the result of which is shown as a red line in the figure. Two parameters of the linear function are extracted - the slope and the x-intercept, referred to as offset - and are checked against limits, listed in table 6.4. The quality of the fit is checked by demanding the χ^2 per degree of freedom to be lower than 1. These comparisons are used to ensure the linearity of the ADC response, and a sufficient coverage of the ADC range to use the DAC for the intended purpose of setting a common baseline between all channels in the PreProcessor system.

The linearity of the ADC is further analysed by checking deviations of the data from the fit for all DAC settings. For this purpose, the difference between the mean value of the measured ADC output and the value of the fitted function is computed individually for each DAC setting. The difference is then compared to a limit, that

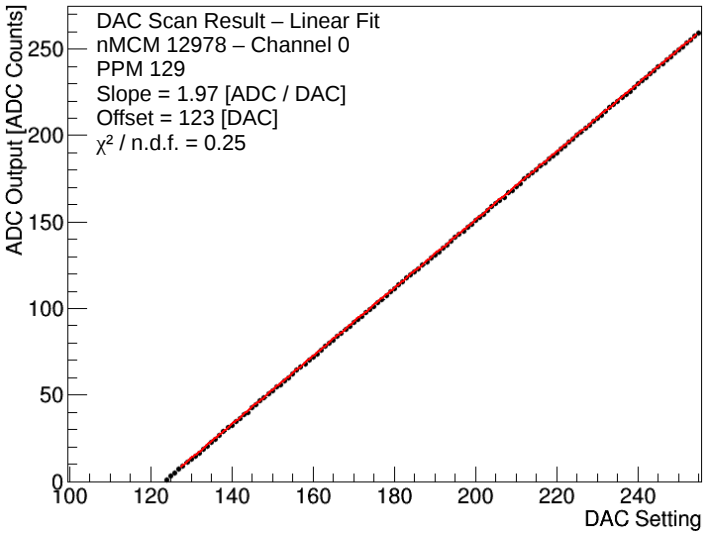

Figure 6.5: Result of a DAC Scan for one nMCM channel. The average ADC Output is shown as a function of the DAC setting. A linear fit is applied and the result shown as the red line.

is defined by the value of the measured slope in ADC Counts.

The DAC Scan data is also used to investigate the noise, by analysing the distribution of ADC Counts. Two types of noise variables are computed: the RMS Noise is equal to the RMS of the distribution, while the Peak-to-Peak Noise is equal to the largest difference between any two ADC values. The measurement is iterated for a programmable number of times, in order to accumulate sufficient statistics. All individual noise measurements are checked against predefined thresholds, which are listed in table 6.4.

Figure 6.6 shows the average of the two noise variables described above over five iterations of the measurement. It shows the average RMS Noise (Figure 6.6 (a)) and the average Peak-to-Peak Noise (Figure 6.6 (b)) as a function of the applied DAC setting. The average noise over all DAC settings is marked by the red line in both plots. No dependence of the noise on the DAC setting is observed.

Variable	Limit
Slope	$\in (2.0 \pm 0.1) \frac{\text{ADC}}{\text{DAC}}$
Offset	$\in (125 \pm 25)$DAC
χ^2/n.d.f.	≤ 1.0
RMS Noise	≤ 0.7 ADC Counts
Peak-to-Peak Noise	≤ 3 ADC Counts

Table 6.4: Limits imposed by the DAC Scan analysis to check the validity of the nMCM.

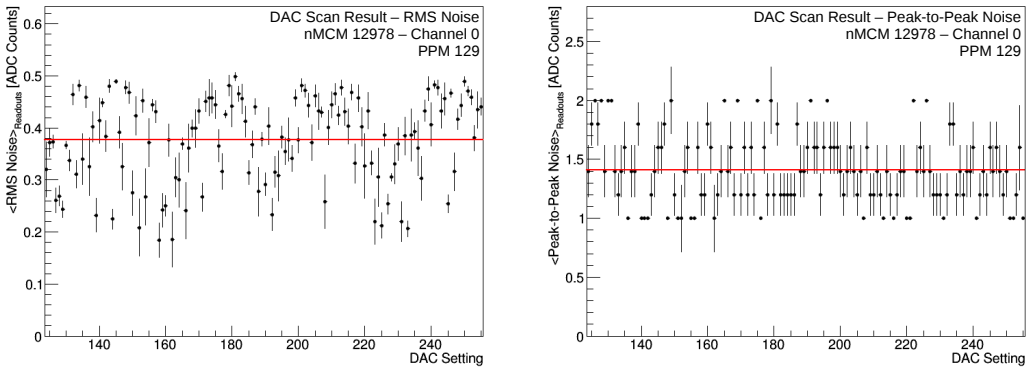

(a) RMS Noise **(b)** Peak-to-Peak Noise

Figure 6.6: Noise in one PPM channel, as measured for different DAC settings. For each DAC setting, the average RMS Noise (a) and Peak-to-Peak Noise (b) for five readouts of 127 samples each are presented. The red line marks the average noise over all DAC settings.

The final check done on the DAC Scan data is the ADC Code analysis. This analysis checks that all possible ADC outcomes - i.e. all ADC values between zero and the maximum recorded value - are present in at least one readout. This is done to ensure that all bits of the ADC operate properly.

As the analogue performance is harder to define using simple predefined thresholds, the test results for all modules that are ordered into the Faulty category by this test were cross-checked to make the final categorisation between 'Faulty' and 'Usable'.

Signal Generator Pedestal Scan

To test the higher bits of the ADC, which are not covered by the range of the DAC, the Signal Generator on the nMCM is used. It allows to scan the entire dynamic range of the ADC, using much finer steps than the DAC Scan.

The Signal Generator Pedestal Scan operates similarly to the DAC Scan by measuring the ADC response to different Signal Generator configurations. In contrast to the DAC Scan, not all 4096 possible configurations are taken into account, since the difference between neighbouring settings is much smaller: while the increase of the DAC setting by one step results in an increase of the ADC output by about two counts, the Signal Generator Pedestal has to be increased by three steps in order to increase the average ADC output by one count. Because of this, a step size of 3 is chosen for the Signal Generator Pedestal Scan, resulting in 1365 steps starting from 0 and ending at the maximum value of 4095.

The result of the scan is shown in Figure 6.7 (a). The whole dynamic range of the ADC is scanned, and the signal generator follows a linear distribution. However, the Signal Generator does not provide the same linearity as the DAC. This is seen

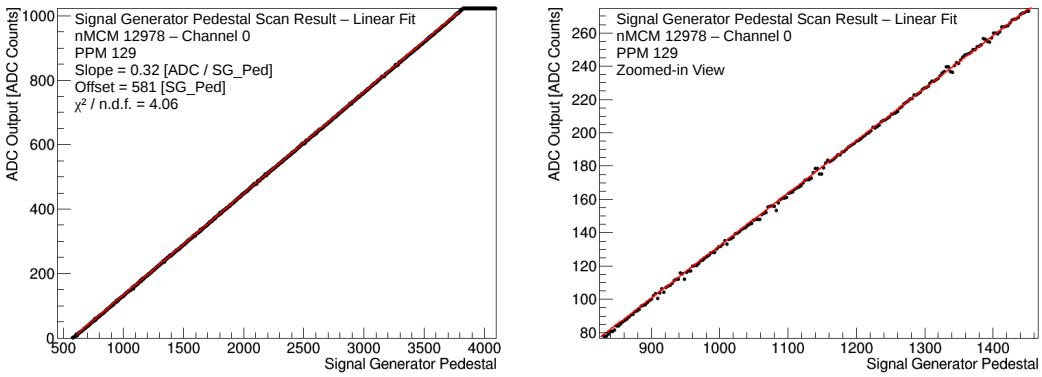

(a) Overall **(b)** Zoomed-In View

Figure 6.7: Result of a Signal Generator Pedestal Scan for one nMCM channel. The average ADC Output is shown as a function of the Signal Generator Pedestal. A linear fit is applied and the result shown as the red line. (a) shows the full result, while (b) shows a zoomed-in view to illustrate the bad linearity of the Signal Generator Pedestal circuit.

in Figure 6.7 (b), which shows a zoom of the previous plot. The ADC output is observed to fluctuate around the linear fit more strongly than in the DAC scan (Figure 6.5). Noise measurements also show that the noise level is higher - the average RMS Noise is of the order of 0.5 ADC Counts and Peak-to-Peak Noise reaches up to 4 ADC Counts, as compared to 0.4 ADC Counts and 2 ADC Counts, respectively, in the DAC Scan.

Because of these effects, the fit and noise results in the Signal Generator Pedestal Scan are not used to evaluate the ADC functionality. The test results are still useful as cross-checks for faulty modules identified by the DAC Scan.

ADC Test Results

The DAC Scan is the most selective test in the nMCM production test procedure. In total, 245 nMCMs failed to pass at least one of the selection criteria outlined above. Because of this, a more in-depth discussion of the results of this test is provided.

Table 6.5 provides an overview of the number of nMCMs failing each of the selection criteria, for both the Single-Board and Full-Crate setups. The total number of nMCMs failing each test are also provided. Note that the different categories are not exclusive, i.e. there are cases of nMCMs failing multiple tests which are included in all relevant counts. The 'Both' category is also not exclusive - i.e. all nMCMs counted in this category are also counted in the 'Single-Board' and 'Full-Crate' categories.

The most selective criteria are the noise requirements and the χ^2 limit. For these criteria, more faulty nMCMs are identified in the full-crate test than in the single-board test. This is due to the higher temperature in the full-crate environment,

nMCM Error in DAC Scan	# nMCMs failing			
	Single-Board	Full-Crate	Both	Total
RMS Noise	68	98	40	126
Peak-to-Peak Noise	42	46	37	51
DAC Fit Slope	9	9	9	9
DAC Fit Offset	25	16	15	26
DAC Fit χ^2	43	56	9	90
Mean ADC Error	14	14	13	15
ADC Code Error	1	1	1	1
Total	140	174	69	245

Table 6.5: Number of nMCMs failing the DAC Scan test due to passing the thresholds defined in Table 6.4.

which impacts the performance of the analogue signal processing on the nMCM.

An overview of results for all nMCMs which were tested in the DAC Scan is shown in Figure 6.8. It presents the distributions of the various quantities that are measured in the DAC Scan for all channels of all tested nMCMs.

Figure 6.8 (a) shows the distribution of the RMS noise. While the noise measurement is iterated several times for each DAC setting as described before, only the highest measured noise value is included in the distribution. The distribution starts around 0.5 ADC Counts and falls steeply, with outliers ranging from 1 ADC Count up to values as high as 3.5 ADC Counts for very noisy nMCMs. The limit of 0.7 ADC Counts manages to select all these outliers, but also part of the tail of the core distribution.

The Peak-to-Peak noise measured for each channel is presented in Figure 6.8 (b). Like for the RMS noise, only the largest noise value for each channel is included. The majority of channels have a noise of 3 ADC Counts, just below the limit demanded by the test. The long tail of channels with noise up to 10 ADC counts is caught by the limit, leaving the core distribution untouched in this case.

For both the slope and the offset parameters of the fit to the DAC Scan data, shown in Figure 6.8 (c) and (d) respectively, no tails of outliers are observed. The distributions peak within the limits defined in the test procedure, and only parts of the tails are selected to not pass the test. Note that in these cases the limits that were derived from the prototype nMCMs do not apply too well to the full series of production nMCM - e.g. the limit on the slope is centered around a value of 2 ADC/DAC, while the distribution peaks slightly before that. Due to this, the fit parameter criteria are only taken into account for clear outliers.

For the χ^2 of the fit, shown in Figure 6.8 (e), the distribution is broader, with a slowly falling edge to higher values. The applied limit cuts into the higher end of the distribution, labelling more than just outliers as bad modules.

The distributions show that many modules fail the limits of the tests only slightly. Due to this, and because many nMCMs only fail one of the various selection criteria,

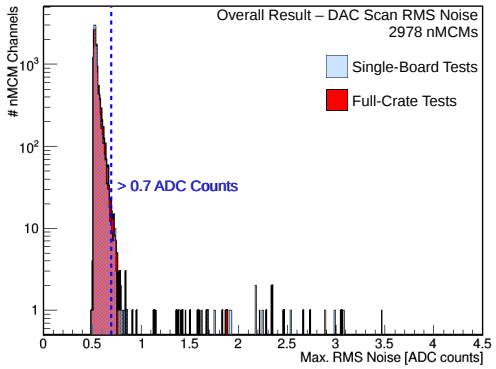

(a) RMS Noise Distribution

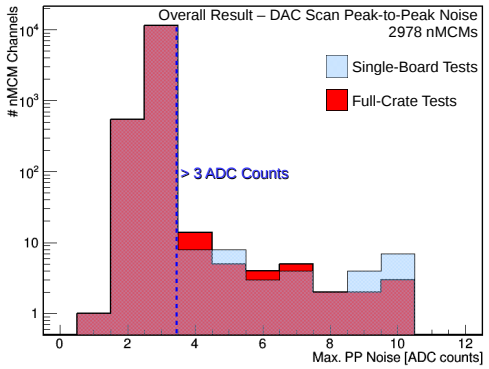

(b) Peak-to-Peak Noise Distribution

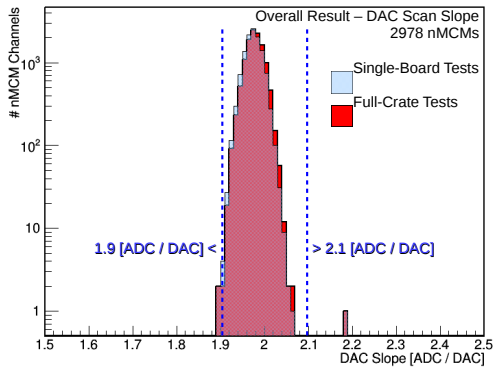

(c) DAC Slope Distribution

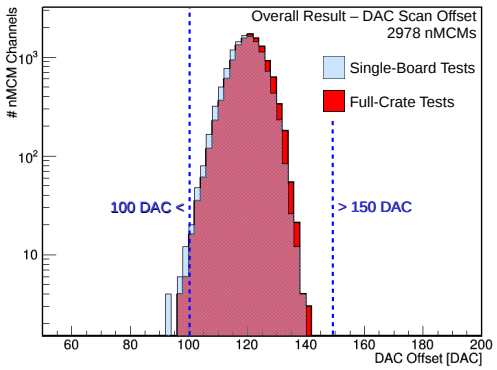

(d) DAC Offset Distribution

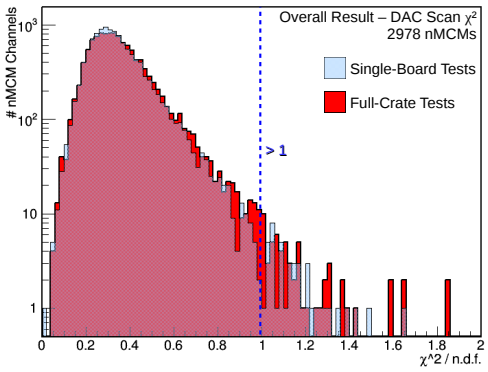

(e) χ^2 Distribution

Figure 6.8: DAC Scan results for all tested nMCMs. All plots show the distribution of one of the quantities derived from the DAC Scan for each channel - the largest measured RMS noise in (a) and Peak-to-Peak noise in (b), the slope and offset of the fit in (c) and (d), and the χ^2 of the fit in (e). Distributions are shown for both the Single-Board and the Full-Crate test. The cuts that are applied to find bad modules are also shown.

most of the nMCMs that do not pass the DAC Scan are assigned to the 'Usable' category. Out of the 245 nMCMs identified by the DAC Scan, only 12 are labelled as 'Faulty' due to severe errors. Figure 6.9 shows two examples of 'Faulty' nMCMs. In one case, Figure 6.9 (a), the ADC response is strongly non-linear, hinting at problems in that component. Figure 6.9 (b) showcases an ADC with a much too high response, resulting in a large ADC output even in the case of a DAC Setting of zero, and to a slope far outside the limits required by the test. Other examples for 'Faulty' nMCMs are modules showing a constant response to all DAC Settings, or extremely large noise up to 20 ADC Counts Peak-to-Peak.

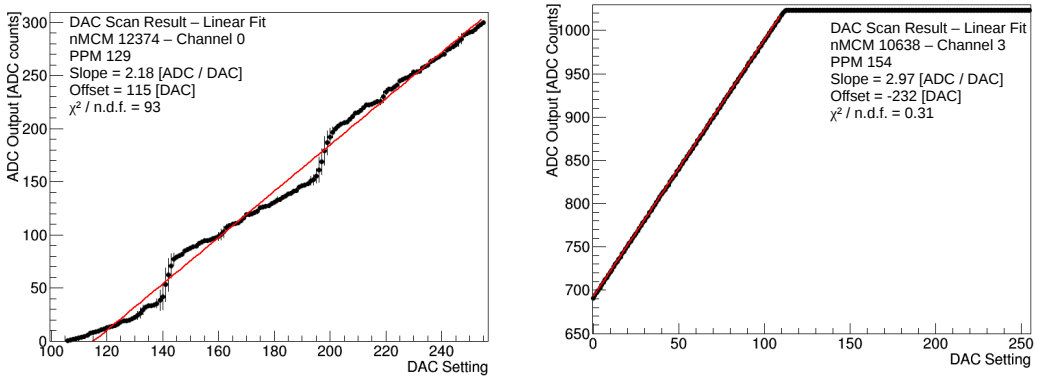

(a) Non-Linear ADC Response **(b)** Negative Offset

Figure 6.9: Examples for DAC Scan results from 'Faulty' nMCMs. (a) shows the case of a strongly non-linear response of the ADC. In (b), the ADC output is strongly offset, such that a DAC Setting of zero already results in a large ADC response.

The Signal Generator Scan does not add anything beyond these results for the majority of nMCMs, though it does confirm the DAC Scan results in all the 'Faulty' cases. One single nMCM is additionally identified as 'Faulty', because the Signal Generator circuit of one channel does not generate any non-zero input to the ADC in this case.

6.2.7 Input Timing Tests

This test checks the functionality of the input stage of the CALIPPR FPGA and the timing configuration of the ADCs.

In total there are four possible configurations for the ADC input, given by the combinations of the two 80 MHz (rising or falling edge) and 40 MHz (first or second sample) latching configurations. The first goal of the input timing test is to check that for each of the 24 possible fine timing delays either the rising or the falling edge of the 80 MHz gives a stable result. This is checked in the 80 MHz Latching test, described below.

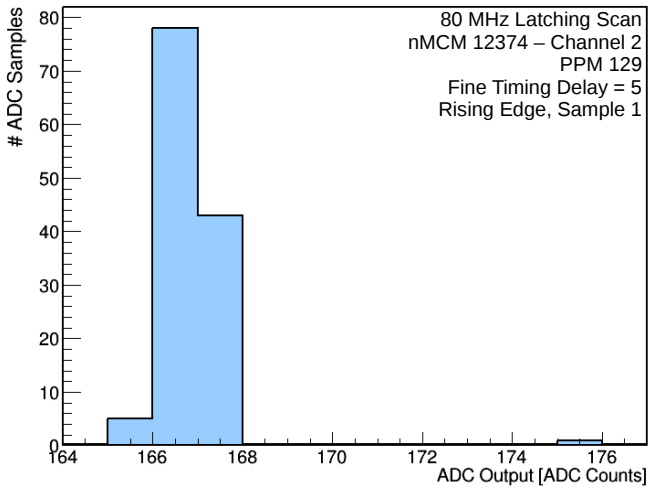

Figure 6.10: Example result of 80 MHz Latching test. The data obtained by one readout operation is shown, with a clear outlier at 175 ADC counts.

Furthermore, switching the latching configuration for the 40 MHz latch, while not changing the time at which a pulse is sampled, results in a *different* ADC sample taken at a *different* time being taken into the internal processing. This is equivalent to introducing an additional shift in the timing in addition to the fine timing delay. To make sure that signals can be properly digitised at any point in time despite this effect, the 40 MHz Latching test is performed.

The final test of the input timing consists of the reconstruction of a signal pulse using the fine timing capability. This forms a cross-check of the results achieved in the previous tests.

80 MHz Latching Test

If a specific combination of input configuration and fine timing delay leads to latching of an unstable signal in the CALIPPR input stage, this is seen as outliers in a measured ADC distribution. The 80 MHz Latching test uses the DAC to generate a stable ADC output, which is read out and recorded like in the DAC test.

Five readouts of 127 ADC samples each are taken for each of the 96 possible combinations of fine timing delay and input configuration. The average over all samples is calculated to serve as a reference value. The difference between each measured sample and the reference is then calculated and compared to a limit of 2 ADC Counts. If the difference is larger than this limit, the input and fine timing configuration used to record the corresponding data is flagged as unstable.

Figure 6.10 shows an example for a channel exhibiting a latching error. It shows a distribution of ADC samples recorded using the rising edge of the 80 MHz clock to latch the signal and the first sample in the 40 MHz reduction step, with a fine timing delay value of 5 (corresponding to a delay of about 5.2 ns). The distribution

peaks around 166 ADC Counts, with some entries in the neighbouring bins, but an outlier is also recorded at 175 ADC Counts. This corresponds to, at least, a one bit latching error: 175 in binary is `1010 1111`, while 167 is `10100111`. In this case, the rising edge configuration is marked as unstable for a fine timing configuration of 5.

An nMCM only fails this test if instabilities like this are found for both the rising and the falling for the same fine timing configuration. As discussed above, these instabilities are expected to be observed to a certain extent - the important thing is that a stable input configuration exists for each fine timing delay.

40 MHz Latching Test

In order to check whether a given signal can be digitised at any point in time, a fine timing scan is performed on a known pulse shape. In this type of measurement, the same pulse is generated many times using the Signal Generator - or any other external source capable of generating input signals to the ADCs. Each time the pulse is generated, the nMCM readout is triggered and 127 ADC samples are recorded. By iterating the fine timing delay between the pulses, the pulse shape is recorded for all delay settings. This allows the reconstruction of the pulse shape with the fine timing resolution of 1.04 ns.

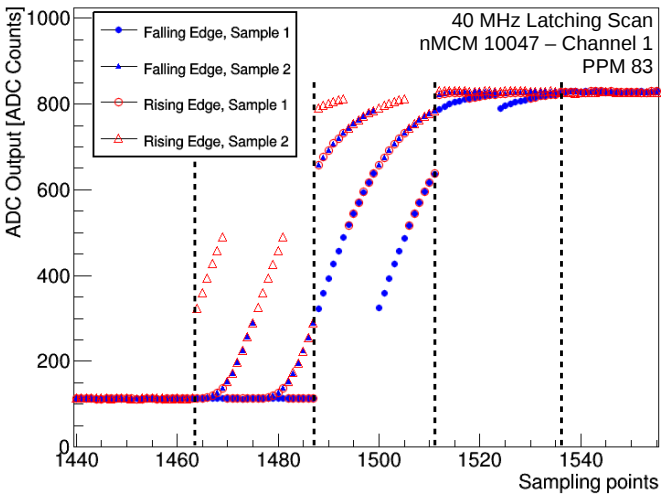

Figure 6.11: Reconstructed signal in the 40 MHz Latching test. The reconstructed signals for all four possible configurations are shown.

Figure 6.11 shows the reconstruction results for a rising edge input signal. The ADC output is plotted as a function of the time of digitisation, where the distance between two sampling points is equal to the fine timing resolution. Each set of 24 subsequent sampling points corresponds to the same sample in the readout, but recorded with different fine timing delays. The borders between two such sets are marked by the dashed vertical lines in the figure.

The Figure shows four different curves, each being the result for one of the four possible latching configurations of the CALIPPR input stage. The most striking feature of all the curves are the incontinuities: each set of 24 subsequent points cover the same 12 values twice. This is due to the nature of the fine-timing delay, as it covers the whole time between two BCs, while the ADC digitises the signal twice in this time. A fine timing delay of 12 thus corresponds to shifting timing of the ADC output by half the distance between two BCs, i.e. by 12.5 ns, which is exactly the clock frequency used to latch the output in the FPGA. This means that all pairs of delay settings that have a difference of 12 between each other produce the same result.

A further complication arises due to the non-zero timescales involved in the digitisaton process: the propagation of the ADC strobe signal from CALIPPR to the ADC, the actual digitisation and updating of the output signal in the ADC and finally the propagation of the digital result back to the FPGA. The processing time in the ADC is the largest of these contributions, with up to 6 ns passing between the reception of the strobe signal and the output acquiring its new value [51]. An additional offset equivalent to 6 delay settings is introduced between the rising and falling edge latching modes, as this is the timing difference between the clock edges. As a result of this, the incontinuities occur at delays 0 and 12 for the falling edge, while it is seen at delays 0, 6 and 18 for the rising edge.

Figure 6.12 shows a timing diagram of the ADC output and the CALIPPR input stage to illustrate these effects. The top two rows show the ADC strobe signal and the ADC output signal, while the lower three rows show the signals in the CALIPPR input stage: the 80 MHz clock used to latch the signal, the resulting latched ADC data, and the selected sample, all for the case of the rising edge and first sample latching setting. Figure 6.12 (a) shows the case for a delay setting of 0, i.e. the input clock and the ADC strobe are in phase. The ADC output, displaced from the ADC strobe by the propagation delay, is latched by the input clock and the first sample is selected for the 40 MHz processing. In Figure 6.12 (b), the fine timing delay is increased to 6, moving the ADC clock by half a clock cycle relative to the input clock. This results in the analogue input being sampled at different points, indicated by the labels of the ADC signals. It also shifts the timing of the ADC output, such that the CALIPPR input clock now registers an earlier sample in the latching process. Effectively, the ADC samples get pushed into later clock cycles in the CALIPPR by the fine timing delay, resulting in the jumps observed in Figure 6.11.

The test consists of finding those settings that taken together result in the reconstruction of a single, continuous curve. This is done by iteratively testing for incontinuities between two samples. All settings that do not give a continuous continuation of the curve are excluded for the respective fine timing setting. If no setting fulfills this requirement for a given sample, the construction is restarted from that sample. Once 24 subsequent samples have been found that form a continuous curve, the test is stopped and the resulting good latching settings for each fine timing delay are stored.

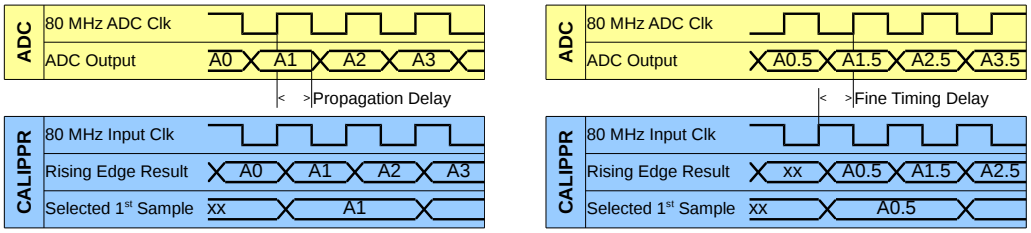

(a) Fine Timing Delay 0 **(b)** Fine Timing Delay 6

Figure 6.12: Timing diagram for the ADC and the CALIPPR input stage for the case of the rising edge and first sample used as the latching configuration. (a) shows the situation for a fine timing delay of 0 and indicates the propagation delay, while (b) shows it for the highlighted delay of 6.

Pulse Reconstruction Test

The Pulse Reconstruction test forms a cross-check of the results in the previous two tests and of the consistency of the latching behaviour across different modules. It consists of a fine timing scan similar to the one described above, with the difference that only one set of latching settings is used for each fine timing delay. These settings have been determined from tests of the preproduction series of nMCMs, and are given in table 6.6.

Fine Timing Delay	0-10	11-16	17-22	23
Sampling Edge	Falling	Rising	Falling	Rising
80 MHz Sample	First	First	Second	Second

Table 6.6: Nominal latching configuration used in the Pulse Reconstruction test.

Figure 6.13 shows the result of the pulse reconstruction for one of the tested nMCMs. Two parameters are extracted from the measured pulse, as indicated in the figure: the height and the width. Here, the height is defined as the difference between the highest measured point in the pulse and the pedestal level, which is determined as the average of 50 points outside the pulse. The width is defined as the time between the 5% and 95% height points on the rising edge of the pulse. Both parameters are compared to predefined limits to decide whether the nMCM passes the test - the height has to be at least 700 ADC Counts, and the width has to be less than 100 ns.

Input Timing Test Results

In total, 22 nMCMs do not pass the Input Timing test. 11 of these are also identified in the ADC tests, showing e.g. large noise or nonlinear ADC response, resulting in instabilities being observed in the timing test and being (wrongly) interpreted as

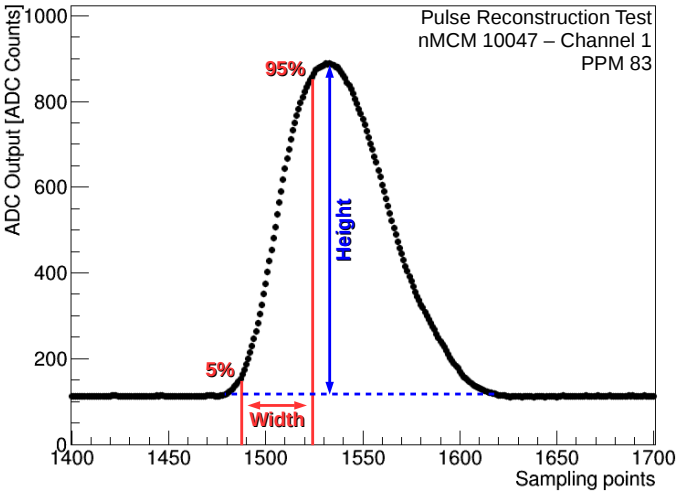

Figure 6.13: Example of the Pulse Reconstruction test. The height and width parameters determined in the test are also indicated.

timing issues. An example is shown in Figure 6.14, which shows the result of the 40 MHz test for an nMCM with strongly non-linear ADC response (see also Figure 6.9 (a)).

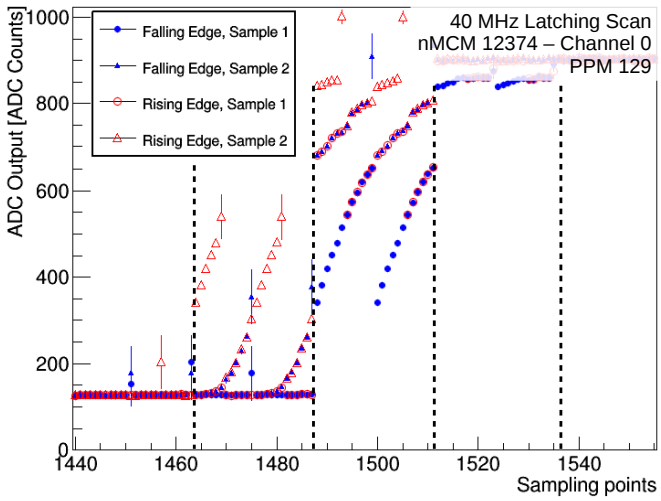

Figure 6.14: Example for an nMCM not passing the 40 MHz Latching test.

Of the remaining eleven nMCMs, eight do not pass the 80 MHz test. For most of these this is an effect of slightly wider than usual ADC sample distributions, leading to samples outside the applied cuts. Furthermore, the majority of them is found in either only the Single-Board or the Full-Crate test. Because of this, all but one of the nMCMs, which shows instabilities for almost all fine timing settings, are ordered to the 'Usable' category.

Two more nMCMs do not pass the 40 MHz test, because the signal shape is distorted at higher ADC values. This is due to bad performance of the ADCs at higher bits, which are not covered by the DAC Scan.

The remaining 'Faulty' nMCM exclusively fails the Pulse Reconstruction test, due to not giving stable results for the default latching configuration.

6.2.8 Histogramming Test

This test checks the proper functionality of the histogramming feature of the nMCM.

It is not possible to use the Playback in order to test the Histogramming operation, because both functionalities use the same memory resources in the initial CALIPPR FPGA implementation. Instead, the test uses the Signal Generator to produce an input to the ADC, allowing to test both the ADC and the BCID input modes. The same pulse shape as in the Pulse Reconstruction test is used.

By operating the Signal Generator in loop mode, the pulse is produced with a frequency of about 2.34×10^5 Hz. The Histogramming threshold is set to 100 counts, equivalent to 400 ADC Counts if dropping the two least significant bits from the ADC input. This way, the pedestal is cut out and only the samples in the pulse contribute to the histogramming. The test then checks whether the Histogramming operation concludes, which is the case as soon as one of the histogramming bins saturates. If this is not the case within a programmable time limit, the nMCM is considered to have failed the test.

Four nMCMs do not pass the histogramming test. In three of these cases this is due to a malfunctioning ADC, and the nMCMs are also detected by the previous tests. For the remaining nMCM, the Signal Generator does not produce a non-zero output, leading to no inputs in the Histogramming. This nMCM is also detected in the other tests using the Signal Generator.

6.3 Summary of Test Results

Out of the 3000 produced nMCMs, 22 do not pass the first electrical test. 16 of these were sorted out from the beginning due to problems occuring in their production, five were found to behave incorrectly in the tests and one was found to be equipped with the wrong type of connector.

The 2978 remaining nMCMs were checked by an automated series of tests as described above. In total, 257 modules did not pass all the tests flawlessly, as listed in Table 6.7. In most cases, the failures are of a minor nature, and the nMCMs are thus ordered into the 'Usable' category. Only 15 modules are taken as 'Faulty', mostly due to badly performing or malfunctioning ADCs. The exceptions from this rule are one module for which transmission errors are detected in the LVDS test, one module for which the Signal Generator does not produce any output to the ADC and one module that is detected as unstable in all possible configurations in the 80 MHz latching test.

nMCM Test	# nMCMs failing			
	Single-Board	Full-Crate	Both	Total
Configuration	0	0	0	0
Rate Metering	0	0	0	0
LVDS Transmission	0	1	0	1
Event Readout	0	0	0	0
DAC Scan	140	174	69	245
80 MHz Latching	6	7	1	12
40 MHz Latching	6	8	6	8
Pulse Reconstruciton	8	5	5	8
Histogramming	4	4	4	4
Total	145	184	72	257

Table 6.7: Number of nMCMs failing each type of nMCM test.

With 2717 modules passing all tests, the overall efficiency is 90.05%. The requirements described in Section 6.1 are easily met - the number of 'Good' nMCMs is sufficient to equip the PreProcessor system and to fulfil the 25% spare policy.

The main goal of the tests, that is to ensure that only fully functional nMCMs are installed in the PreProcessor system at CERN, is also fulfilled: since the installation in 2014, only 5 nMCMs have shown problems in two years of operation. Compared to dozens of MCMs which were replaced by spares during Run-1, the nMCM has also proven more reliable than its predecessor.

6.4 Installation of nMCMs at CERN

The nMCMs were installed in the PreProcessor system in three batches over the course of 2014. In a first step, one crate was equipped with the new modules in May of that year, with three more crates following in July, covering one half of the system. The final batch was installed at the end of August, completing the process.

The installation process consists mainly of replacing the MCMs by nMCMs on all PPMs, but also includes a minor amount of soldering work, removing a few resistors, on the PPM in order to include the nMCMs in the JTAG chain. The typical process to upgrade one PPM thus consists of a few steps. First, the input cables are removed and the module is taken out of the crate and moved to the soldering station. After the soldering is done, the MCMs are taken off the module, to be replaced by nMCMs. Finally, the module is put back into the crate and the input cables are re-inserted.

For each crate, basic configuration and JTAG tests were done immediately after completing the nMCM installation, to make sure that the relevant interfaces operate properly. Further commissioning was done in the remainder of LS1, consisting mostly of updating the software tools that operate the modules and integrating additional trigger algorithms in the CALIPPR FPGA implementation (see Chapters 7 and 8).

7 Improved Bunch-Crossing Identification on the nMCM

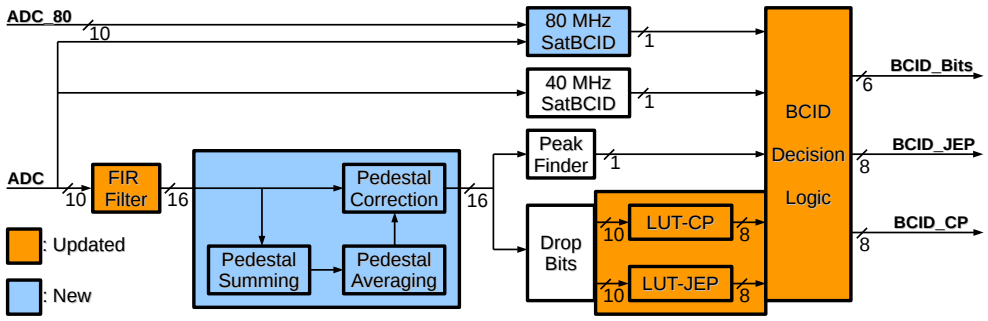

Figure 7.1: Diagram of the PreProcessor BCID logic and the updates in the CALIPPR FPGA implementation. Updated algorithms are marked in orange, while new functionalities are highlighted in blue.

As described in Chapter 5, the LHC upgrade during LS1 imposes several challenges on the L1Calo trigger. The increased luminosity leads to more pile-up, resulting in increased noise in the trigger tower signals and large fake rates for several L1Calo trigger items. The higher centre-of-mass energy leads to larger trigger tower signals, increasing both the probability of signals saturating the ADCs and the degree of saturation when they do.

All of these effects impact the trigger tower signals. They can thus be overcome by improving the processing of the individual signals in the channels of the CALIPPR FPGA. For this reason, improving the BCID logic is one of the main motivations of the PreProcessor upgrade.

The improved BCID logic builds upon the initial version described in Chapter 5.2.3. Figure 7.1 shows a block diagram of the improved BCID logic, highlighting those parts that have been either updated or newly introduced compared to the algorithms operating in Run-1 (see Figure 5.3). The improvements are the following:

- The FIR filter has been updated to allow a wider range of filter coefficient values. This makes it possible to apply different filtering schemes that are optimised for Run-2 noise conditions.

- The output of the filter is further processed in a newly introduced Pedestal Correction logic. This logic block subtracts pile-up effects on the pedestal independently for each BC. The correction is dynamically computed from the filter output in two processing blocks parallel to the main signal processing.

- The energy calibration LUT has been split up into two identical blocks. One of these provides the E_T calibration for the CP system while the other one is used for the output to the JEP system. This allows to separately optimise the E_T calibration for each of the two processors.

- The processing of saturated signals is improved by introducing an additional, enhanced version of the SatBCID algorithm. This algorithm makes use of the 80 MHz sampling frequency of the ADCs on the nMCM. By analysing the rising edge of saturated pulses at this higher resolution, the correct BC can be identified more precisely.

The following sections describe these changes in more detail.

7.1 Digital Filter Extension

The central component in achieving a high efficiency for the Peak Finder BCID algorithm is the correct choice of the FIR filter coefficients c_n. Depending on which filter coefficients are chosen, different types of noise contributions to the trigger tower signal are reduced more effectively. The possible filter configurations in Run-1, as described in Section 5.2.3, are not sufficient to apply filtering schemes optimised for Run-2 conditions. The CALIPPR FPGA implementation of the FIR filter has thus been updated to allow for a wider definition of coefficients.

The following sections briefly describe the noise conditions expected in Run-2, the different filtering strategies applied in Run-1 and Run-2 and the update of the FIR filter implementation.

Trigger Tower Noise

The noise in the trigger tower signals can be split into two categories. The first is the thermal noise. It describes the fluctuation of the electrical signal due to thermal movements of the charge carriers. All signal processing components and cables along the signal transmission path from the calorimeters to the PreProcessor inputs contribute to it. The second type of noise is the pile-up noise. It is caused by statistical fluctuations of the pile-up, which leads to variations in the signal induced in the detector cells.

Figure 7.2 shows the noise conditions expected in the PreProcessor for Run-2. It shows the thermal noise and the pile-up noise on the trigger tower signal, as a function of the position of the trigger tower in $|\eta|$. The thermal noise is shown in Figure 7.2 (a). It is of the order of 200 to 400 MeV, depending on the partition of the calorimeter.

The pile-up noise is shown in Figure 7.2 (b). For the hadronic calorimeter, the pile-up noise is lower than the thermal noise for most of the trigger towers, only exceeding it in the forward region. For the EM calorimeter on the other hand, the pile-up noise is larger than the thermal contribution for all trigger towers. This

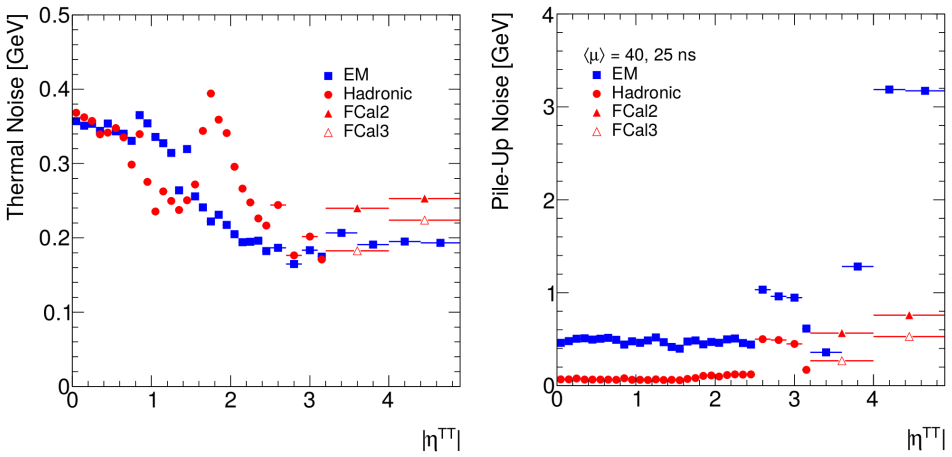

(a) Thermal Noise **(b)** Pile-Up Noise

Figure 7.2: Noise in the PreProcessor system as determined for trigger towers at different $|\eta|$. (a) shows the thermal noise while (b) shows the pile-up noise. Note the different scales. [57]

behaviour is expected, because the pile-up contribution scales with the position and size of the trigger towers. The larger a trigger tower is and the closer it is to the interaction point, the higher the pile-up noise contribution becomes. It is thus maximal in the larger FCal towers in the inner, electromagnetic layer.

Updated Filter Coefficients

The standard filter coefficient configuration during Run-1 was the Matched Filter scheme. In this configuration, the coefficients for each trigger tower are chosen such that they resemble the shape of the digitised pulse. This is the most effective filter scheme to reduce thermal noise [58].

For Run-2, the increased luminosity leads to larger pile-up contributions compared to Run-1, while the thermal noise is not expected to change. In this case, filter coefficients that are optimised for these specific noise conditions perform better. The filter scheme chosen for Run-2 are the Autocorrelation Filters, which are computed based on both the correlation of noise samples to each other and the expected pulse shape [57] [59]. The coefficients are derived from minimum bias data [60], using the methods developed in [57]. In the absence of pile-up noise, where only the thermal noise is present, the noise samples are uncorrelated. In this case, the Autocorrelation Filters revert to the Matched Filter scheme described before. If pile-up noise dominates, the noise samples are strongly correlated, leading to filter coefficients that divert strongly from the Matched scheme. Typically, the presence of pile-up noise results in large negative filter coefficients.

The Run-1 filter implementation can not realise the Autocorrelation filter scheme, because the central three coefficients can only acquire positive values (see Section 5.2.3). This limitation is removed in the CALIPPR FPGA implementation. The FIR

filter is updated such that all five coefficients can take both positive and negative values, by including an additional sign bit for the central three coefficients. For the outer two coefficients, the amplitude is increased by one bit, bringing all five samples to the same definition as 5-bit signed integers. The possible values for the five coefficients in both the PPrASIC and the CALIPPR FPGA are listed in Table 7.1.

	PPrASIC		CALIPPR FPGA	
Coefficient	Bits	Values	Bits	Values
c_0	3+sg	[-8,7]	4+sg	[-16,15]
c_1	4	[0,15]	4+sg	[-16,15]
c_2	4	[0,15]	4+sg	[-16,15]
c_3	4	[0,15]	4+sg	[-16,15]
c_4	3+sg	[-8,7]	4+sg	[-16,15]

Table 7.1: Bits used and possible values for filter coefficients in the PPrASIC and the CALIPPR FPGA. '+sg' denotes an additional sign bit.

7.2 Dynamic Pedestal Correction

The Dynamic Pedestal Correction algorithm subtracts the average pile-up contribution from the input signal. The correction is computed from the input signal as the deviation of the average from the expected baseline. This computation is performed independently for each BC in the LHC orbit.

Section 7.2.1 explains how pile-up affects the input signals and the performance of the L1Calo system. The algorithm design is motivated in Section 7.2.2 and explained in detail in Section 7.2.3. Sections 7.2.5 and 7.2.4 cover the readback functionality the synchronisation mechanism, which are important in the commissioning and operation of the algorithm.

7.2.1 Pile-Up Effects on L1Calo Input Signals

As described in Chapter 4, the trigger tower signals are formed by summing the signals of individual calorimeter cells in a defined region of the detector. Pile-up affects the signal formation in the earliest stage: pile-up particles deposit additional energy in the cells, which is added to the trigger tower signal. Because the calorimeter signals extend over several BCs, each such pile-up hit affects the following BCs as well. The total pile-up effect for one BC is thus the sum of pile-up signals originating in earlier BCs and those originating in the BC itself.

Because the positive peak and negative undershoot of the signal are of equal area, the total pile-up contribution is zero on average for many subsequent pile-up signals. As the LHC beams are structured into bunch trains (see Section 3.1.2),

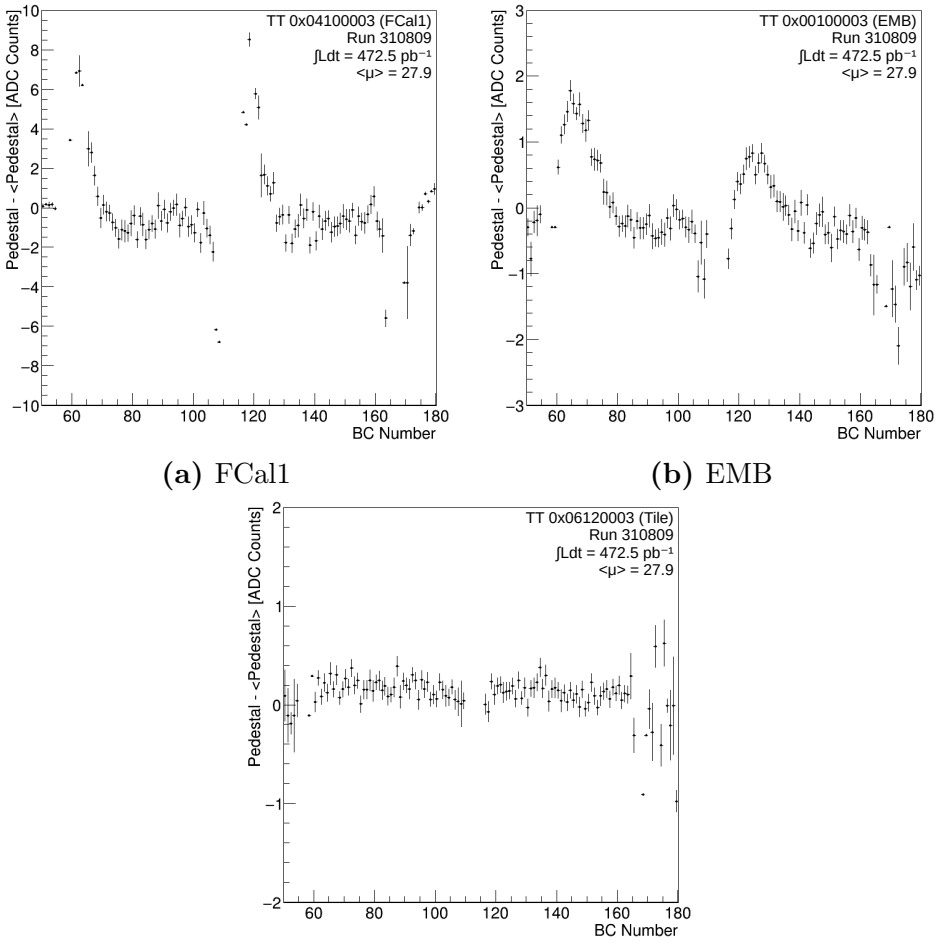

(a) FCal1 **(b)** EMB

(c) Tile

Figure 7.3: Deviation of the pedestal from the average in different calorimeter partitions as function of the BC number. (a), (b) and (c) show a trigger tower from the FCal1, EMB and Tile calorimeter partition, respectively. Note the different scales. BC numbers 110-115 correspond to the gap between bunch trains, so no data is recorded for these. Produced using [61].

this cancellation is not given for all BCs in the LHC orbit. BCs at the beginning of bunch trains are preceeded by empty BCs, resulting in an non-zero average pile-up effect.

This effect is shown in Figure 7.3. It shows the difference between the pedestal measured for each BC and the average of the pedestal over all BCs for trigger towers in different calorimeter partitions (see Section 3.2.3): the FCal1, the EMB and the Tile calorimeter. It is plotted for BCs in a part of the LHC orbit that covers two subsequent bunch trains.

The effect is most pronounced in the FCal1 trigger tower (Figure 7.3 (a)). About ten BCs at the beginning each bunch train show an increased pedestal, with a

maximum increase of 8 ADC counts. This corresponds to 2 GeV of additional energy that is deposited in the trigger tower.

In the EMB trigger tower (Figure 7.3 (b)) the pedestal increase reaches a maximum of around 2 ADC counts. The effect is reduced compared to FCal1 because the trigger towers in the EMB are smaller: (0.1×0.1) compared to (0.4×0.4) in $(\Delta\eta \times \Delta\phi)$. Furthermore, the effect is reduced in the second bunch train compared to the first in the EMB, while the effect is of the same size in both bunch trains in FCal1. This is due to the differences of the pulse shapes between the calorimeter partitions: the undershoot in the FCal1 is much shorter than in the EMB (see Figure 4.4).

Because pile-up particles typically have low momentum, they are mostly absorbed in the electromagnetic calorimeter. The hadronic calorimeter behind it is thus expected to be unaffected by pile-up. This is confirmed by Figure (Figure 7.3 (c)). It shows the pedestal variation for a trigger tower in the Tile calorimeter, which does not deviate significantly from the average.

While this pile-up effect distorts the energy measurement in all trigger towers, it affects empty trigger towers the most. These are usually mapped to an E_T of zero by the BCID algorithms and the LUT noise cut. In the BCs at the start of the bunch trains, the increased pedestal can be larger than the noise cut, resulting in a non-zero E_T result.

Because this increase occurs synchronously in all trigger towers, L1Calo trigger items that sum many trigger towers are most affected by it. The effect is demonstrated in Figure 7.4. It shows the rates of several different triggers as a function of instantaneous luminosity, as measured during Run-1. While the rates of most single-object triggers are linear as expected, the rates of E_T^{miss}-, jet- and multi-jet-triggers show strongly non-linear rises. These point to large fake rates, which are caused by the mechanism described above.

The only way to reduce this effect with the Run-1 BCID logic is to increase the noise cut in the LUT. But optimising the noise cut to the BCs with the increased pedestal affects all the remaining BCs as well, for which the increased cut is too high. To properly correct for the pile-up effect, a BC-wise treatment is thus required.

7.2.2 Pedestal Correction Algorithm

The pedestal correction algorithm corrects for the pile-up induced pedestal distortion described in the previous section. The nature of these distortions imposes several requirements on the algorithm design:

- Because the magnitude of the pedestal increase depends on the position within the LHC orbit, the correction is applied **BC-wise**.

- The magnitude of the distortion also depends on the position of the trigger tower. The correction thus is applied **channel-wise** in the PreProcessor logic.

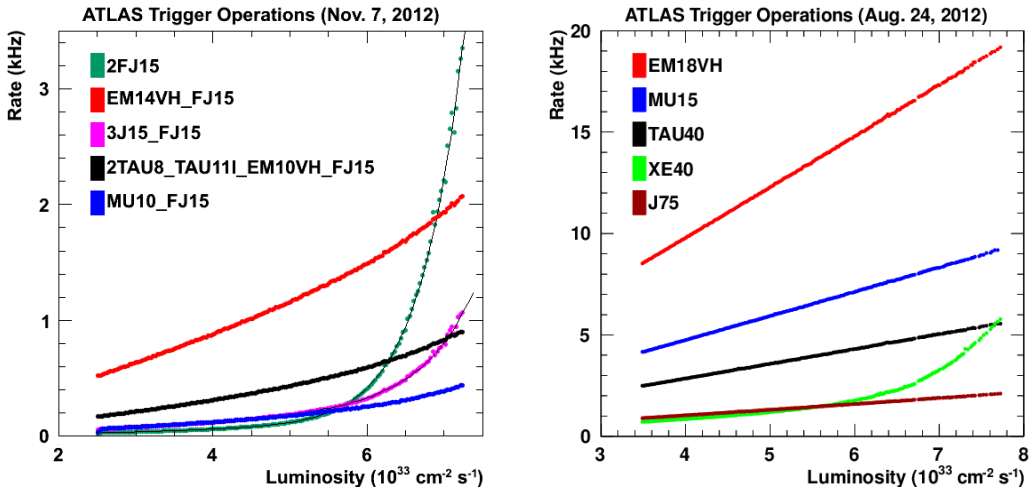

Figure 7.4: Rates of several L1 trigger items as a function of instantaneous luminosity. `EM`, `MU` and `TAU` are triggers for electrons/photons, muons and taus respectively, jet triggers are denoted by `J` and E_T^{miss} triggers by `XE`. The numbers after the objects define the threshold in GeV, while the numbers before the objects define the required number. [62].

- The pile-up depends on the instantaneous luminosity. It thus decreases over the course of an LHC fill. Because of this, the correction is **time dependent**.

- From a technical point of view, the algorithm has to be fast, so that the L1Calo latency is not strongly affected by it.

The implementation that has been chosen is a correction that is dynamically calculated from the input signal itself. The calculation is performed individually for each BC in the LHC orbit and for each channel, thus satisfying the first two requirements. The algorithm follows the luminosity trend by regularly re-computing the correction. To keep the latency minimal, the computation of the correction and its application are executed in parallel. The dynamic approach is also intrinsicaly flexible, as the correction follows all unforeseen trends of the pedestal.

The pedestal correction logic operates on the output of the FIR filter. This is due to the improved resolution and reduced noise contribution compared to operating on the ADC output.

The correction is based on the average filter output $\bar{f}$. The algorithm computes the average individually for each BC and compares it against a configurable constant f^{ped} that describes the expected pedestal value. At each BC, the difference between these two numbers is subtracted from the filter output f, giving the corrected filter output f^{corr} as:

$$f_i^{corr} = f_i - (\bar{f}_i - f^{ped}), \tag{7.1}$$

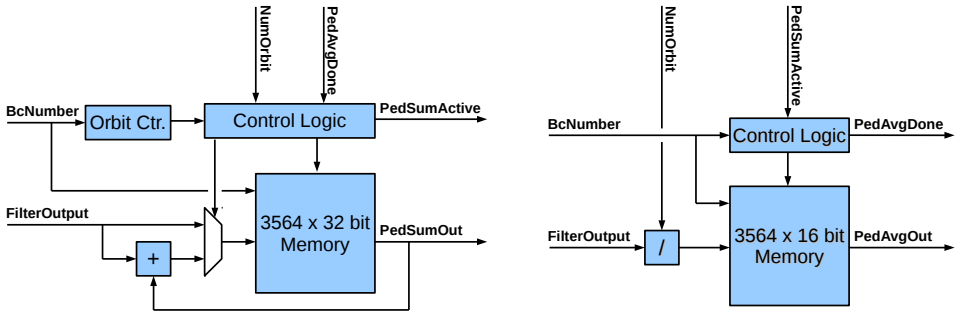

(a) Summing Logic **(b)** Averaging Logic

Figure 7.5: Block diagrams of the two stages of the computation of filter averages: the summing (a), and the averaging (b).

where $i = 0, \ldots, 3563$ is the BC number. The algorithm thus ensures that, on average, the corrected filter output is equal to f^{ped}.

The value for f^{ped} is usually set to the filter output that is expected for a constant, flat pedestal from the ADC. It is derived from the filter coefficients c_n and the ADC pedestal ADC^{ped}:

$$f^{ped} = \sum_{n=0}^{4} c_n ADC^{ped} \tag{7.2}$$

The next section describes the algorithm implementation in detail.

7.2.3 Pedestal Correction Algorithm Implementation

As can be seen in Figure 7.1, the implementation of the pedestal correction algorithm is split into two parallel paths. Along the main signal path, the Pedestal Correction block subtracts the correction from the filter output. The average filter output that the correction is based on is computed in the second path, in the Pedestal Summing and Pedestal Averaging blocks.

The two paths are discussed individually in the following subsections.

Computation of Average Filter Output

The central part of the pedestal correction algorithm is the computation of the BC-wise average filter output. The average $\bar{f}$ is defined as

$$\bar{f} = \frac{1}{N} \sum_{n=0}^{N} f_n, \tag{7.3}$$

where the filter output f is summed over a configurable number of orbits N. The process is divided into two steps: the Pedestal Summing, which computes the filter output sum, and the Pedestal Averaging, which performs the division.

The summing logic is illustrated in Figure 7.5 (a). A 3564×32-bit memory stores the filter output sums for all BCs in the LHC orbit. At each BC, the filter output is added to the memory entry for that BC and the result is used to overwrite the entry in the memory.

The number of orbits that have passed is tracked by a counter. It is incremented by one whenever the BC number reaches the last entry in the orbit, 3563. Once the orbit counter is equal to N, the summing operation is halted, and the averaging logic is activated.

Figure 7.5 (b) shows a block diagram of the averaging logic. It is built around a 3564×16-bit memory block that holds the average filter output for each BC. Once the summing process is complete, the averages are updated within one orbit. At each BC in this orbit, the corresponding sum is read out from the summing memory and divided by the number of orbits N. The result is written to the corresponding entry in the average memory, relacing the value from the previous iteration.

In order for the division operation to be fast, it is realised by a bitshift operation. Thus, to make the division exact, the number of orbits is limited to a power of two.

Once all entries have been updated, the summing process is restarted and the orbit counter reset to zero. In the first orbit of the new summation the filter output is written directly to the memory, instead of performing the summing step. This way, the sums are reset for the next iteration.

Correction of the Filter Output

Using the averages calculated from the filter output as described in the previous section, the correction can be computed. It performs the calculation described by Equation 7.1.

The subtraction is performed in two steps. First, the correction is computed as the difference between the configurable pedestal value f^{ped} and the average filter output $\bar{f}$. In the second step, the correction is subtracted from the filter output. The algorithm detects both underflow and overflow of the resulting corrected filter output. In the first case, the corrected filter output is set to zero, while it is set to the saturated 16-bit value `0xffff` in the latter case.

7.2.4 Synchronisation of the Pedestal Correction

At the start of a standard ATLAS data-taking run, all PPMs are reset and configured sequentially. As the pedestal correction starts its operation as soon as the module is configured, the averages are not updated synchronously for trigger towers that are processed on different PPMs. A synchronous update of the averages over the whole system at a known point in time is preferred, as it simplifies the monitoring of the algorithm.

A synchronisation mechanism has thus been implemented in the pedestal correction logic. It aligns all channels in the PreProcessor system such that the update of the pedestal averages happens at the same time in all of them.

The synchronisation mechanism requires an external signal, that is synchronous in all PreProcessor channels. It uses the event counter reset (ECR), which is distributed by the TTC system roughly every five seconds (see Section 3.4.1).

Because a typical ATLAS run starts before the beams of the LHC reach their optimal stable collision configuration, ECRs are typically not sent out for some time at the start of a run. For this time, between the configuration of the modules and the issueing of the first ECR, the pedestal correction is running asynchronous between the different channels.

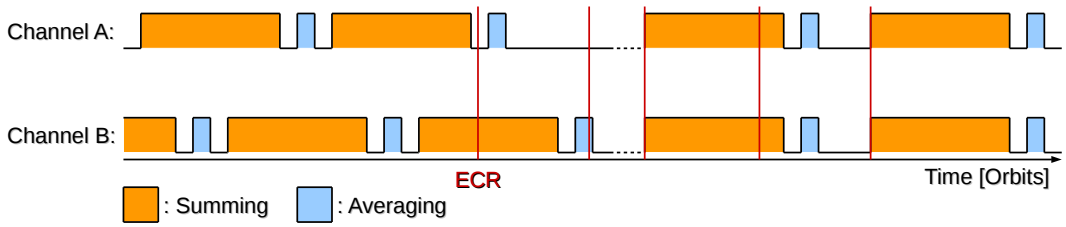

Figure 7.6: Pedestal Correction synchronisation mechanism. The two asynchronous channels are synchronised by the ECR signals.

As shown in Figure 7.6, the detection of the first ECR that is detected after configuration of the nMCM initiates the synchronisation process. The summing and averaging processes that are already running conclude, but no new process is started. Only after a configurable number of ECRs has been detected in the nMCM, the process starts again.

The number of ECRs before restarting the calculation is chosen such that the time between the first ECR and the restart is longer than the time of one iteration of the average calculation.

In addition to the synchronisation, the ECRs are also used to initiate the transition from one set of averages to the next. This is achieved by only starting the averaging process upon reception of an ECR. The updates of the averages are thus tied to controlled signals, making them easier to trace. This mode of operation is activated with the reception of the first ECR. Before that the averages update automatically as described in the previous section.

7.2.5 Readback of Pedestal Average Memory

In order to be able to test the functionality of the pedestal correction algorithm, it is crucial to have access to the values involved in the calculation. The applied correction, $\bar{f}_i - f^{ped}$, is included in the event readout for the central slice in the readout window (see Section 8.2.1). While this is sufficient to monitor the behavior of the pedestal correction over a run, it does not allow to reconstruct the full memory of averages at a specific point in time. To be able to obtain such a snapshot of the memory of averages, a special readback mode has been implemented in the nMCM.

This readback operation sends the contents of the average memory to the ReM FPGA over the serial interface. This has to be accomplished without interfering

with the standard pedestal correction operation. In particular, the readback is not allowed to control the read-pointer of the memory. The latter has to be in a fixed relation to the BC number at all times to make sure that the correct information is used in the pedestal correction. As the averages constitute a large amount of data, the readback operation requires many cycles of the serial interface, which is also required for the readout of event information in case of an L1A.

Due to these limitations, a very conservative approach is taken in the design of the readout algorithm, by only sending a single value over the serial interface each orbit. This minimises the impact on the event readout, and allows to pick up each value from the memory as the read-pointer reaches the BC number in question.

As the average filter output has a width of 16-bit, it can not be packed into a single frame of the serial interface. Instead, it is split into two 8-bit words, that are sent in subsequent orbits. This results in a total of 0.63 s[1] to read out all averages of one channel.

In the first implementation of this readout mechanism, all channels on the PPM were read out in one go. As this meant that no other readback operation could be performed for about 40 s, it was deemed unusable for standard data taking, where regular access to e.g. rate metering and histogramming results is crucial in monitoring the PreProcessor performance. Since implementing the channel-wise readout, the commissioning of this feature has continued, but was stalled again due to other bug-fixes and operational duties having higher priorities. As of the time of writing this document during the 2016/17 end-of-year shutdown, the final commissioning of the controlling software is supposed to be completed in the near future, so that the readout can be used to directly monitor the pedestal correction in the 2017 data taking period. Nevertheless, it has been succesfully used in test setups relying on separate software tools, in order to verify the functionality of the pedestal correction and its various controlling mechanisms.

7.3 Separate LUTs for CP and JEP Systems

In Run-1, the same LUT was used for both the CP and JEP systems. The two systems search for different signatures, which are treated quite differently in the offline calibration and data analysis. Using different calibrations for the two systems allows to treat the signatures independently already in the L1Calo trigger, which can improve the selectivity of the trigger.

The CALIPPR FPGA thus implements two identical LUTs operating in parallel, as shown in Figure 7.1. The output of one of them is used for the CP, while the other is transferred to the JEP. Both receive the same filter output and use the same configuration for the DropBits block, but they are configured independently from each other.

Both LUTs operate exactly as the original implementation (Section 5.2.3) and

[1] 3564×2 orbits $\times (3564 * 25) \frac{10^{-9}\,\mathrm{s}}{\mathrm{orbit}}$

their outputs are treated equivalently in the BCID Decision Logic. The BCID output for the CP is then transferred the BcMux logic, while the output for the JEP is used to form the JetSum (see Section 5.2.4).

7.4 Improved BCID of Saturated Signals

As is described in Section 5.2.3, the SatBCID algorithm analyses the rising edge of saturated signals to identify the BC that contains the peak of the analogue signal. It operates in parallel to the Peak Finder and the results of the two algorithms are combined in the BCID Decision Logic.

In Run-1, no configuration of the SatBCID algorithm was found that reliably identifies the correct BC over the whole saturated regime, i.e. trigger tower energies from 250 GeV to 4 TeV. Because of this, the algorithms were operated in a special setup that still guarantees full BCID efficiency for saturated signals of even the highest possible energy: The Peak Finder and SatBCID algorithms are operated simultaneously for all saturated signals.

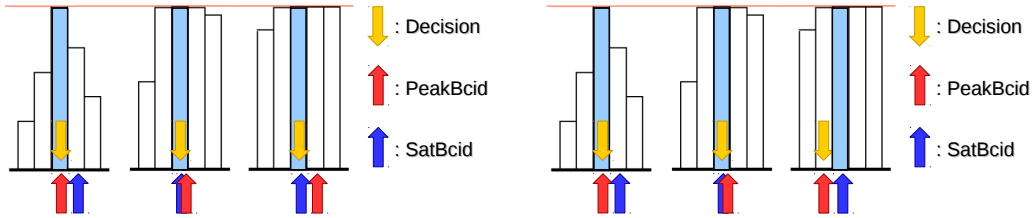

(a) Matched Filter **(b)** Autocorrelation Filter

Figure 7.7: Behaviour of the special Run-1 setup for the BCID for saturated pulses. (a) shows the result using Matched Filters, (b) shows it using Autocorrelation Filters.

The setup relies on the behaviour of the Matched Filter scheme in the saturated regime. For pulses that are just slightly above the saturation energy, only a single sample saturates and the Matched Filters still produce a peak at the correct position. As the energy increases and more samples saturate, the Matched Filter tends to produce peaks for later samples. This is due to the steepness of the two edges of the pulse: the falling edge falls slower than the rising edge rises, so that more samples saturate on the falling than on the rising edge. This effect is enhanced by the analogue saturation, that distorts the falling edge. As the Matched Filter coefficients are mostly positive, the larger samples later in the pulse are stronger weights that move the peak of the filter output to the back of the pulse. At very high energies, the Matched Filter thus produces peaks further back in the pulse, resulting in the Peak Finder identifying a wrong, late BC.

As the BCID Decision Logic always takes the first identified BC in the case where the two BCID algorithms disagree, it has to be ensured that the SatBCID algorithm

decides on the correct sample once the Matched Filter peaks late. At the same time, the SatBCID algorithm must not identify an earlier sample at lower energies, where the Matched Filters peak in the correct BC. This can be achieved with a trivial setup of the SatBCID algorithm. By setting the high threshold to saturation level, 1023 ADC Counts, and the low threshold below the pedestal level, e.g. to 5 ADC Counts, the algorithm will always identify the second saturated sample, $s + 1$, as the peak (see the algorithm definition in Section 5.2.3 for the sample definition).

This leads to the behaviour shown in Figure 7.7 (a). For slightly saturated pulses, the Matched Filters correctly identifies the lone saturated sample, while the SatBCID decicison, which points to the later sample, is ignored. As the energy rises, the sample before the peak saturates, making the second saturated sample the correct choice. At these energies, both the Matched Filter and the SatBCID algorithm identify the correct BCID. At very high energies, the Matched Filter peak moves, such that the Peak Finder identifies a late sample. In these cases it is the SatBCID that takes the correct, earlier decision.

This setup has been studied extensively during Run-1, and the correct functionality and full BCID efficiency have been demonstrated up to centre-of-mass energies of 8 TeV [63], the highest value reached in Run-1 operations. This corresponds to a maximum of 4 TeV in one trigger tower, given a situation where the entire energy of the pp collision is divided into two final state particles, one of which deposits its entire energy in a single trigger tower.

For Autocorrelation filters, which are the default for Run-2, this setup is no longer valid. This is illustrated in Figure 7.7 (b). The large negative coefficients in this filter scheme reverse the behaviour in the saturated regime - in contrast to the Matched Filters, the Autocorrelation Filters identify earlier samples as the energy increases further. With the Decision Logic always taking the earliest BCID decision, this leads to wrongly identified BCs in the saturated regime.

In order to solve this situation, the Peak Finder algorithm decision has to be ignored by the Decision Logic for highly saturated pulses. The obvious solution of ignoring the Peak Finder result as soon as one sample saturates is not feasible, due to the calibration problems for the SatBCID algorithm mentioned before. As any higher energies are not accessible with the threshold method implemented in the Run-1 BCID Decision Logic, it has been upgraded to extend its reach into the saturated regime. This upgrade is described in Section 7.4.1.

With this upgrade, only the SatBCID algorithm operates on the highest saturated pulses. It is thus critical that this algorithm identifies the correct sample up to the highest energies that can be achieved in Run-2. With the centre-of-mass energy increasing up to 13 TeV, this means that energy depositions up to 6.5 TeV in one trigger tower have to be covered.

While the SatBCID has been validated to work correctly up to a trigger tower energy of 4 TeV in Run-1 [63], the performance beyond this point is not guaranteed. The latest BC the SatBCID algorithm can identify is the one corresponding to the second saturated sample, $s + 1$. In the case that two samples before the peak saturate, the SatBCID algorithm can no longer identify the correct sample.

This situation can be solved by using the higher digitisation frequency of the ADCs on the nMCM to improve the SatBCID algorithm. This enhanced version of the algorithm is described in section 7.4.2.

7.4.1 Updated BCID Decision Logic

The BCID Decision Logic, as described in Section 5.2.3, applies energy thresholds to decide which BCID algorithms are considered for each sample. As the energy measurement of the PreProcessor is limited to below 256 GeV, the BCID decision logic can not distinguish pulse energies beyond this point. To achieve a distinction of saturated pulses, more information than just the height of one sample is required.

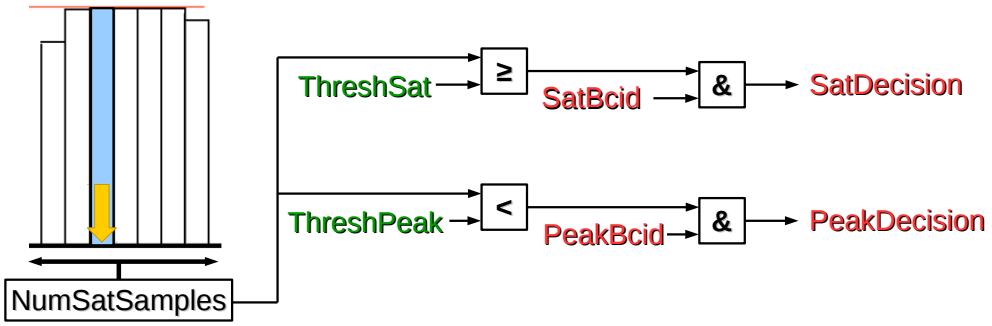

Figure 7.8: Functionality of the updated BCID Decision Logic.

The functionality of the updated Decision Logic is shown in Figure 7.8. It counts the number of saturated ADC samples in a fixed window around the investigated BC. It considers seven samples, starting two BCs before the central one and ending four BCs after it[2]. The number of saturated samples is compared to two thresholds, one for the Peak Finder and the other for the SatBCID algorithm. Depending on the outcome of the comparison, the algorithm results are either taken into account or ignored by the Decision Logic.

For the Peak Finder, the threshold functions as an upper limit: The Peak Finder result is taken into account only if the number of saturated samples contained in the pulse is lower than the threshold. The SatBCID threshold is implemented in the inverse way, that is the number of saturated samples has to be at least as high as the threshold for the SatBCID result to be considered in the BCID decision.

Using the updated Decision Logic, the inefficiency of the Autocorrelation Filters for saturated pulses explained in the previous section can be circumvented. This is illustraded in Figure 7.9, which shows one possible configuration for the two thresholds. In this example, the Peak Finder result is ignored once more than three samples are saturated. The wrong peak produced by the Autocorrelation filter is thus ignored, and the correct SatBCID decision is used instead. The example also

[2]This is the maximum number of samples that can be considered without altering the latency of the BCID logic.

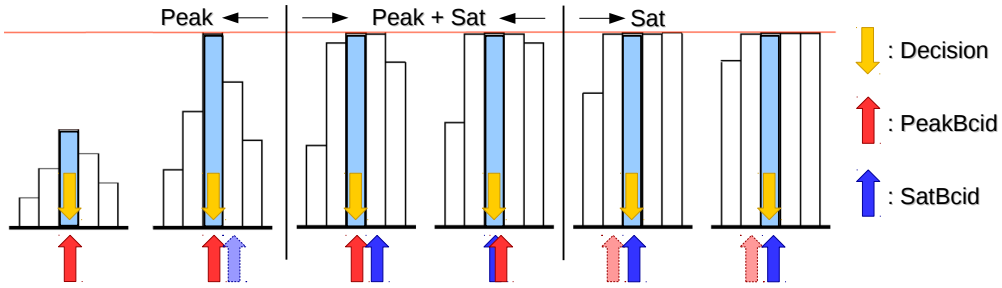

Figure 7.9: Example for the updated BCID Decision Logic.

shows the SatBCID only being active once at least two samples are saturated. By using the algorithm only in a part of the saturated phase space, its calibration can be optimised purely for the higher energies.

7.4.2 80 MHz Saturated BCID Algorithm

In order to improve the SatBCID algorithm for the highest possible energies, an enhanced version of the algorithm is included in the BCID logic. This version of the algorithm functions similar to the SatBCID algorithm described in Section 5.2.3, as it compares ADC samples on the rising edge of saturated pulses to identify the correct BC. In contrast to the SatBCID, the new Sat80BCID algorithm uses the full 80 MHz resolution of the ADCs.

As described in Section 5.2.1, the input stage of the CALIPPR FPGA reduces the 80 MHz stream of ADC data to a 40 MHz stream by only taking every other sample. The remaining samples form a second 40 MHz stream, that is used in the Sat80BCID algorithm. Following the notation introduced in the description of the SatBCID algorithm, the additional samples are described by half-integer numbers. The Sat80BCID algorithm then reads:

- If ADC$(s-1) > high$ and ADC$(s-1.5) > low$, s is identified.
- If ADC$(s-1) > high$ and ADC$(s-1.5) \leq low$, $s+1$ is identified.
- If ADC$(s-1) \leq high$ and ADC$(s-0.5) > high$, $s+1$ is identified.
- If ADC$(s-1) \leq high$ and ADC$(s-0.5) \leq high$, $s+2$ is identified.

Compared to the SatBCID algorithm, this logic has two major changes. Firstly, the BC $s-2$ is no longer considered by the algorithm logic. Instead, the additional sample at $s-1.5$ is compared to the *low* threshold. This change is especially important for the first of the four cases listed above. In this case, the sample $s-2$ is 2 BCs before the peak. This corresponds to the rise time of the trigger tower signals of about 50-60 ns. The sample $s-2$ is thus mostly given by noise contributions, and less sensitive to the height of the pulse. The sample $s-1.5$ is situated closer to the peak, making it the better choice.

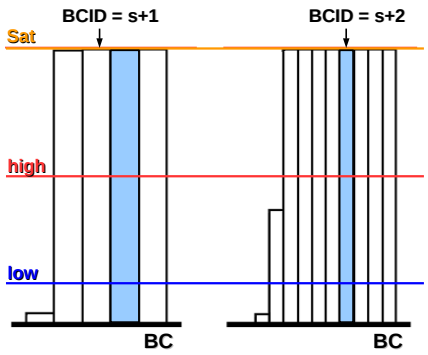

Figure 7.10: Comparison of the SatBCID (left) and Sat80BCID (right) algorithms.

The second change is an additional comparison in the case where the sample $s-1$ is below the *high* threshold. This is used to improve the algorithm performance in for strongly saturated pulses by making it possible to identify the BC $s+2$. This is illustrated in Figure 7.10. It shows the response of the SatBCID and the Sat80BCID algorithms to the same highly saturated pulse. The SatBCID algorithm response is shown on the left. Because two samples in front of the peak saturate, the algorithm can by its definition not identify the correct BC. Instead it identifies the BC before it, because the non-saturated sample in front of the peak is very low. In case of the Sat80BCID algorithm, shown on the right in the figure, the additional 80 MHz sample on the rising edge is taken into account. This allows the algorithm to identify the correct BC $s+2$.

While the two saturated BCID algorithms are implemented independently, they are not used for the BCID at the same time. The configuration of the BCID Decision Logic allows to choose between using either the SatBCID or the Sat80BCID. While the other algorithm still operates, and its decision is included in the event readout (see Section 8.2.1), it is not taken into account for the actual BCID decision. This implementation allows to operate one of the algorithms in a *monitoring mode*, and was introduced to calibrate the Sat80BCID without impacting the L1Calo performance.

8 Additional Monitoring and Testing Facilities on the nMCM

The upgrades to the BCID algorithms described in the previous chapter significantly alter the way the L1Calo trigger decision is formed. Similar enhancements to the monitoring and testing facilities on the nMCM are thus required to ensure the same level of data quality checks as applied in Run-1.

Section 8.1 describes changes in the Rate Metering and Histogramming facilities, while Section 8.2 provides a description of updates of the event data readout mechanism. Updates of the Playback are presented in Section 8.3.

8.1 Monitoring Algorithms

The Rate Metering and Histogramming algorithms monitor the PreProcessor performance unbiased by the L1 trigger decision (see Section 5.2.6). Both can be configured to monitor either the ADC data that is latched in the input stage of the CALIPPR FPGA or the calibrated energy output of the BCID logic. Due to the separation of the energy calibration LUTs for the CP and JEP systems, both monitoring algorithms have to be updated accordingly.

Both the Rate Metering and the Histogramming have been similarly split into two separate copies. The copies have identical functionality, but different inputs, with one copy receiving the BCID output to the CP and the other copy operating on the BCID output to the JEP. The copies are configured independently, so that the respective thresholds and energy inputs can be chosen as required.

8.2 Event Data Readout

The nMCM provides the DAQ system with event data upon receiving an L1A signal from the CTP (see Section 5.2.5). The information is used for two main purposes: the monitoring and the calibration of the system. To ensure a monitoring quality on the same level as in Run-1, additional data words are included in the readout for each event, as described in Section 8.2.1. In order to calibrate the Sat80BCID algorithm, pulse shape information at 80 MHz resolution is required. For this purpose, an additional readout type, that includes the full 80 MHz data stream from the ADCs in the readout, has been added. This is described in Section 8.2.2.

8.2.1 Additional Readout on the nMCM

During Run-1, the readout content in the standard readout mode consists of a header word, five ADC data words and one BCID data word for each trigger tower (see Figure 5.11). This is sufficient information to recalculate the energy output of the BCID logic from the ADC data and to compare it against the value contained in the BCID word in the monitoring software. To perform the same comparison with the upgraded BCID logic described in Chapter 7, additional information in the event data readout is required.

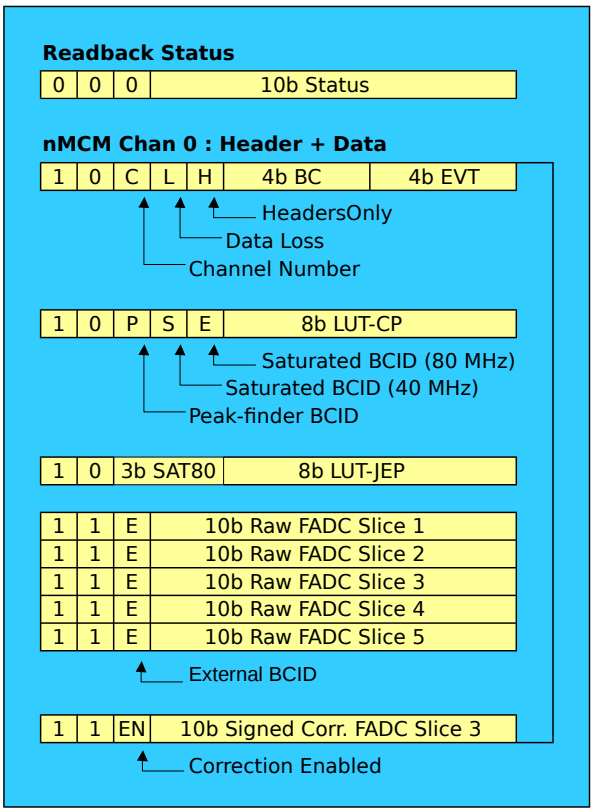

Figure 8.1: Updated Readout format on the nMCM. Adapted from [54].

Figure 8.1 shows the updated readout format for the PreProcessor in Run-2. Compared to the Run-1 format, it includes two additional data words.

First, the split of the energy calibration LUT into two separate tables for the CP and JEP systems means that an additional BCID-type word is required. It contains the 8-bit energy output to the JEP, while the original BCID word is used for the output to the CP.

The three remaining bits in each of the BCID words are used to encode the decision of the BCID algorithms. Similar to the Run-1 implementation, the CP word includes three bits that indicate the decisions of the BCID algorithms. To

include the decision of the Sat80BCID algorithm, the ExtBCID bit is dropped from the readout format.

Additional information on the Sat80BCID algorithm is included with the JEP word. The three remaining bits in this words encode the result of the threshold comparison of the two 80 MHz samples used in the algorithm, s-0.5 and s-1.5 (see Section 7.4.2). For a saturated pulse, the two samples can be arranged in six ways as shown in Figure 8.2. These six situations are mapped to 3-bit patterns ranging from `001` to `110`. The two remaining patterns are used to signify non-saturated pulses (`000`) and invalid pulse shapes (`111`)[1].

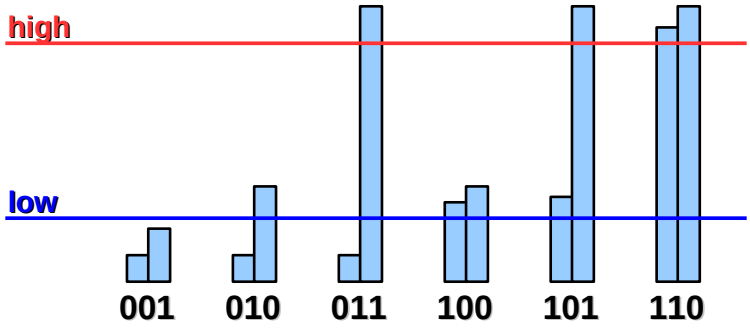

Figure 8.2: Readout bits on Sat80BCID included in the JEP word. The two samples s-0.5 (left) and s-1.5 (right) can be arranged in six different ways for saturated pulses. These situations are mapped to the 3-bit patterns shown below them.

The second additional readout word contains information on the pedestal correction (see Section 7.2).. The applied correction value, $\bar{f}_i$-f^{ped}, is truncated to a 10-bit signed value, and packed with an *Enable Bit* that is set to `1` if the pedestal correction algorithm was active and `0` otherwise. The truncation operation respects saturation, i.e. the value is set to 511 (-512) if the pedestal correction is larger than (less than) that value.

The two additional readout words make it possible to keep the same event monitoring quality in Run-2 that was achieved in Run-1, by recomputing the output of the BCID logic from the ADC input and the applied pedestal correction and comparing it to the energy responses included in the readout.

8.2.2 Readout of 80 MHz ADC Data

While the ADCs on the nMCM sample the analogue input at 80 MHz, the input stage of the CALIPPR FPGA reduces the input data stream to 40 MHz as described in Section 5.2.1. An additional readout mode has been introduced to have access to

[1] A saturated pulse shape is considered invalid if the the sample s-1.5 is larger than s-0.5.

the samples that are otherwise dropped. This 80 MHz readout mode is required to calibrate and test the new Sat80BCID algorithm.

The 80 MHz readout mode can be activated by setting a configuration bit from VME via the ReM FPGA. The samples that are dropped in the input logic are collected in an additional readout pipeline (see Section 5.2.5). If the 80 MHz readout mode is active, upon reception of an L1A, samples from both the standard and the additional ADC sample pipeline are copied to the ADC derandomiser. The number of samples that is taken from each pipeline depends on the readout configuration and is chosen such that the central ADC sample is taken from the standard ADC pipeline. Figure 8.3 shows an example of this for the standard readout configuration, which includes five ADC samples. In this case, three samples are taken from the standard pipeline, while two samples are taken from the additional pipeline.

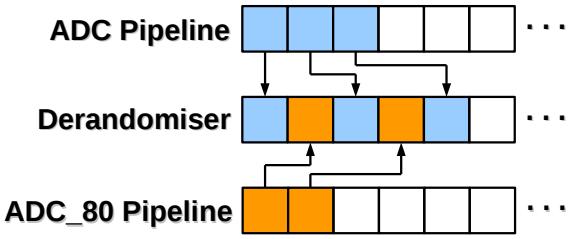

Figure 8.3: Handling of ADC data if the 80 MHz readout mode is activated.

8.3 Updates to Playback

Like the monitoring, the testing facilities on the nMCM are also adapted to the changes in the BCID algorithms. Specifically, the Playback functionality (see Section 5.2.7) is upgraded to add an 80 MHz mode to test the Sat80BCID algorithm. In addition, the Playback and Histogramming algorithms, that use the same memory resources in the initial implementation (see Section), are split. This is done to achieve an equivalent treatment of the Histogramming for CP and JEP, as described above.

8.3.1 Playback of ADC Data at 80 MHz

To be able to test the functionality of the Sat80BCID algorithm, a test data stream operated at 80 MHz is required. The Playback functionality is thus updated to include an additional 80 MHz mode.

Similar to the special Readout mode described above, the 80 MHz Playback mode is activated from the VME via the ReM FPGA. If the Playback is activated while configured in this mode, it's output value is updated at 80 MHz. If the configuration is not set, the update frequency is 40 MHz.

8.3.2 Separation of Histogramming and Playback

While the Histogramming and Playback functionalities share the same memory resource in the initial implementation, the CALIPPR FPGA provides sufficient resources to separate the two algorithms. This change is mostly introduced to make the duplication of the Histogramming into two separate copies more reliable, but it also allows to use the Histogramming in digital tests with the Playback.

9 Performance of the nMCM

The nMCMs have been installed in the PreProcessor system in 2014 and have since been operated for two years of data taking in the ATLAS experiment. During this time, the changes to the processing logic described in the previous chapters have improved the performance of the PreProcessor significantly. This chapter provides an overview of the improvements that have been gained with the nMCM compared to the MCM.

9.1 Input Noise

The noise performance of the nMCM was optimised by removing or reducing several sources of cross-talk between the trigger tower signal lines and other signals on the PPM:

- The number of clock signals that are sent to the nMCM from the ReM FPGA is reduced. The signals are instead derived in PLLs in the CALIPPR FPGA (see Section 5.2.8).
- The format of the Serial Interface is reworked such that all transferred bits are zero in the idle state (see Section 5.2.9).

The impact of these changes on the noise performance of the nMCM is demonstrated in Figure 9.1. It shows the result of a noise measurement that consists of performing the DAC scan procedure discussed in Section 6.2.6 for all possible fine timing configurations (see Section 5.2.2). Figure 9.1 (a) shows the average of the noise over all channels on the PPM versus the fine timing delay. For the MCM result, the noise varies with the position of the sampling point. This is a consequence of cross-talk between the digital and the analogue plane of the MCM. For the nMCM, this effect is not observed, and the noise is constant with the fine timing delay. This is achieved by removing the majority of the clock signals that are transferred across the PPM, thus minimising the source of the cross-talk.

Figure 9.1 (b) shows the noise for each of the 64 channels on the PPM, averaged over all fine timing delays. For the MCM, two populations of channels are observed with different levels of noise. This separation corresponds to the placement of the ADCs on the MCM: The signals that are digitised by the two ADCs that are placed closer to the PPrASIC (see Figure 4.9) experience more noise. This is caused by the longer analogue signal path over the MCM, so that the signal is more likely to be affected by cross-talk. This effect is removed on the nMCM, as the routing of all four trigger tower signals to the dual-channel ADCs is equivalent.

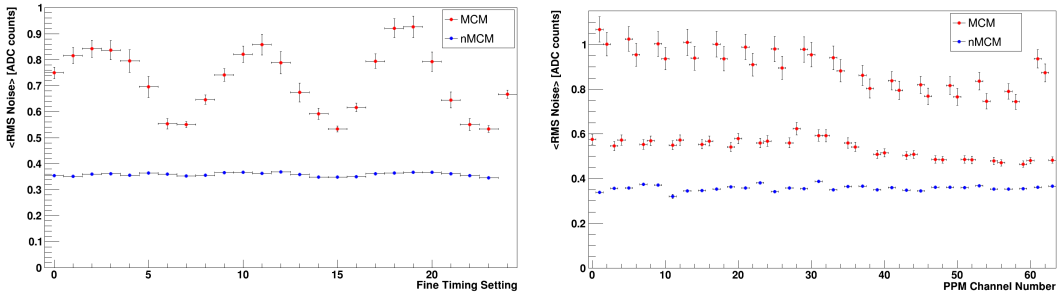

(a) **(b)**

Figure 9.1: Noise for all channels on one PPM. (a) shows the average noise per channel as a function of the fine timing delay. (b) shows the average noise over all fine timing delays per channel.

In addition to removing the described systematic effects, the nMCM implementation also lowers the overall noise level. This is achieved by a combination of all the described improvements.

Similar noise measurements for the whole PreProcessor system were performed after the installation of the nMCMs at CERN. For this purpose, the distribution of ADC values recorded in each trigger tower channel is measured in time periods with no beam in the LHC and no calorimeter calibration systems being active. The gaussian width of this distribution is then taken as a measure of the noise.

In contrast to the standalone tests of the PPM, in these measurements all sources on the input signal path contribute to the noise. In addition to the PreProcessor noise, this includes the noise on the high voltage of the calorimeter cells, in the cables and in the Receivers.

Figure 9.2 (a) shows the results of such a measurement. It shows an η-ϕ map of all trigger towers in the electromagnetic layer, taking into account the lower granularity in the forward region. Each trigger tower is colour-coded according to the measured noise value. Figure 9.2 (b) shows the result for a measurement of the same type that was performed in Run-1, when the PreProcessor system was still equipped with MCMs.

In both measurements, the noise is maximal in the central region and becomes smaller for trigger towers at larger $|\eta|$. This is due to the reweighting of the cell signals to E_T - the factor $\sin\theta$ that is applied to the analogue signal reduces the absolute noise for the trigger tower.

For the MCM, the two populations of channels are again observed as the structure of alternating trigger towers that is visible in the central and endcap regions ($1.5 < |\eta| < 2.4$). No such systematic effect is observed for the nMCM. The better noise performance of the nMCM that was observed for a single PPM is also confirmed in the full system, with a reduction of the noise by

Similar studies are performed on the correlation of the noise between different trigger tower channels that are processed on the same PPM. To cancel contributions from upstream processing systems, a gain of zero is applied in the Receiver system

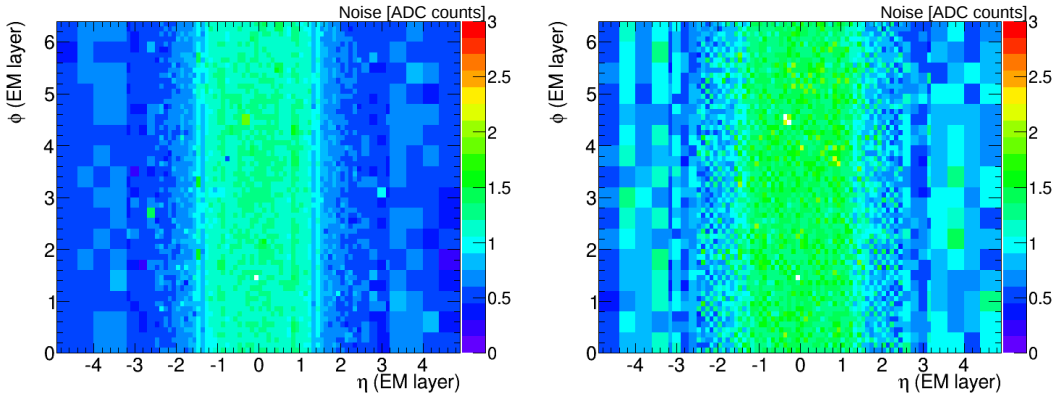

(a) **(b)**

Figure 9.2: Noise measured in all trigger towers of the electromagnetic calorimeter. (a) shows the result for the nMCM, (b) shows it for the MCM.

for this type of measurement. As a typical example, Figure 9.3 shows the linear correlation coefficients[1] for all pairs of the 64 channels processed on one of the PPMs. The results are presented for both a PPM equipped with nMCMs (a) and MCMs (b).

While correlation coefficients of up to 0.8 are observed in the MCM case, they are strongly reduced with the nMCM. This is due to the optimisations described before, which remove effects that affect many channels on the board in the same way.

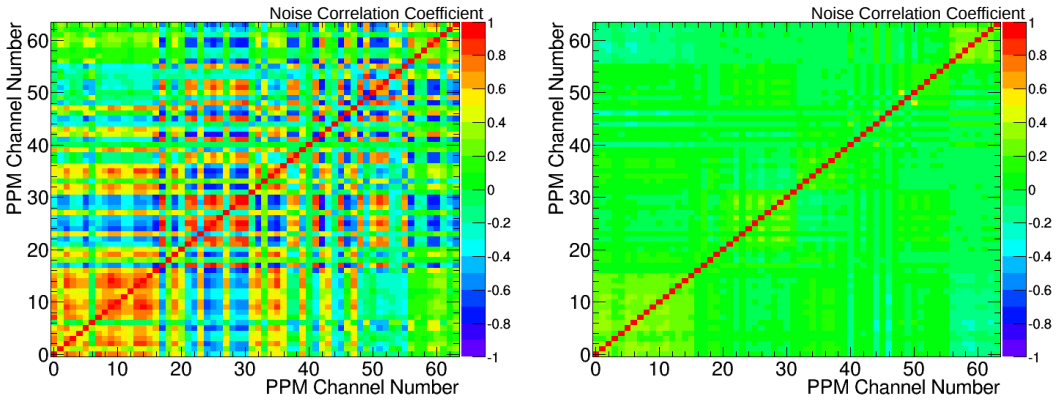

(a) **(b)**

Figure 9.3: Correlation coefficient for noise samples between channels on one PPM. (a) shows the result for a PPM equipped with MCMs, (b) shows it for a PPM equipped with nMCMs.

[1]The linear correlation coefficient is defined as: $\frac{n\sum x_i y_i - \sum x_i \sum y_i}{\sqrt{n\sum x_i^2 - (\sum x_i)^2}\sqrt{n\sum y_i^2 - (\sum y_i)^2}}$

9.2 Autocorrelation Filters

Next to the noise performance, the signal detection efficiency of the PreProcessor is also improved on the nMCM. The Autocorrelation filter configuration is optimised on the Run-2 noise conditions, where pile-up noise is the dominant contribution. Their use is made possible by the extension of the filter definition in the CALIPPR FPGA implementation.

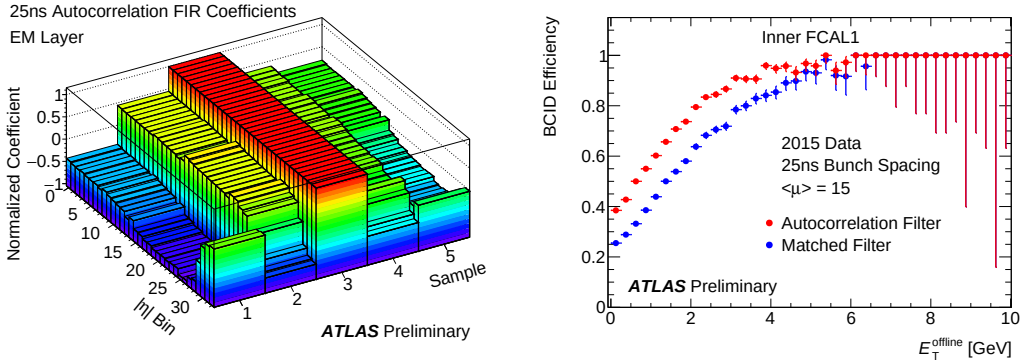

(a) **(b)**

Figure 9.4: Performance of the Autocorrelation filter configuration. (a) shows the filter coefficients as a function of the $|\eta|$-bin. (b) shows the BCID efficiency for the Autocorrelation and Matched filter schemes as a function of E_T for FCal1 trigger towers. [62]

Figure 9.4 (a) shows the Autocorrelation filter setup optimised for the electromagnetic calorimeter in 2015. The filter coefficients are the same for all trigger towers in the same $|\eta|$-bin, where $|\eta|$-bins are defined as the trigger tower row in $|\eta|$ (see e.g. Figure 4.2), starting with 0 in the central EMB and ending at 32 for the most forward FCal1 trigger tower. Note that the coefficients are normalised to one in the figure, while the actual configuration in the nMCM uses integer values.

For most of the electromagnetic calorimeter, the coefficients follow the same behaviour, with their value decreasing with distance from the centre. This reflects the behaviour of the Matched Filter scheme, where the coefficients resemble the pulse shape. At lower $|\eta|$-bins only the first coefficient takes on negative values, with the last one also becoming negative in the EMEC region, around $|\eta|$-bin 20. In the FCal1 region, represented by the last two $|\eta|$-bins, the coefficients behave very differently. Here, the outer coefficients become positive, while the second and fourth coefficientstake on negative values. This is due to the increased pile-up contribution in this calorimeter region, due to larger trigger towers and closer position to the beam pipe. This configuration is only possible on the nMCM, which extends the coefficient definition to allow for negative values for other coefficients but the outer ones.

Figure 9.4 (b) shows the BCID efficiency[2] as a function of the offline E_T[3] for the inner FCal bin, i.e. $|\eta|$-bin 32, which is the are most affected by pile-up. It is shown for both the Matched and the Autocorrelation filter scheme. With the Autocorrelation filter, a higher efficiency is achieved at very low energies. This is crucial in obtaining a high trigger efficiency for low energy objects.

9.3 Pedestal Correction

The Pedestal Correction algorithm has been introduced to subtract the pile-up contribution from the input signals to the PreProcessor. It is one of the central enhancements of the PreProcessor for Run-2.

As described in Section 7.2.2, the algorithm computes the average of the filter output for each BC in the LHC orbit and subtracts the difference between this average and the expected pedestal value from the filter output, as described by Equation 7.1. The difference $\bar{f}_i$-f^{ped} is included in the readout information for each triggered event (see Section 8.2.1) and is used for monitoring during data taking.

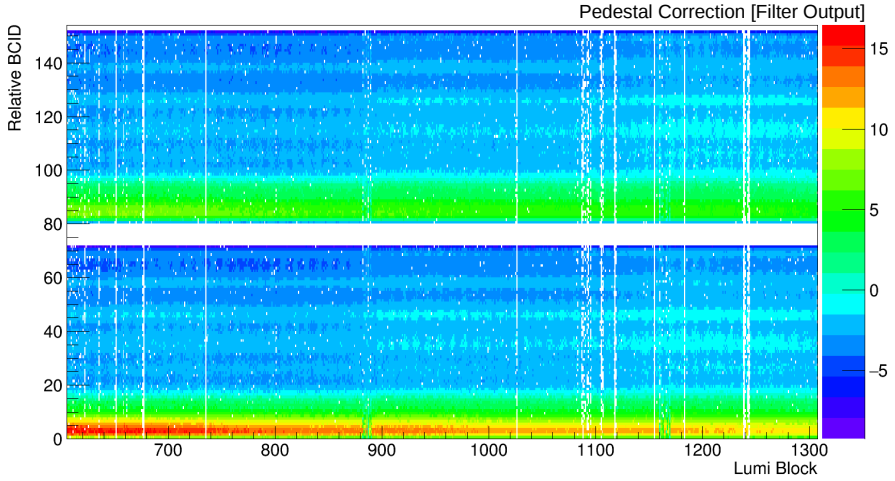

Figure 9.5: The average pedestal correction in the EMB as a function of time in lumi blocks and of BC Number relative to the beginning of a bunch train after a long gap. Produced using [61].

Figure 9.5 shows an example for the behaviour of the Pedestal Correction over the course of a data taking run. It shows the applied correction as a function of

[2]The BCID efficiency is defined as the fraction of events for which the Peak Finder algorithm identifies the correct BC for the respective filter setup.

[3]The *offline* E_T is the energy contained in the trigger tower as measured by adding the energies of the individual cells contained in the trigger tower, in contrast to the 8-bit calibrated *online* E_T determined by the PreProcessor.

the time in lumi blocks[4] along the x-axis and of the position in the orbit along the y-axis. The latter is given as the BC number relative to the start of a bunch train after a long gap, and reaches up to the beginning of the following long gap, covering two bunch trains and the short gap between them. The contributions of all such bunch train pairs in the LHC are averaged to increase the statistics. For the same reason, the correction is averaged over all trigger towers in the EMB calorimeter.

For a given lumi block, the pedestal correction can be examined as a function of the BC number. It quickly rises to its maximum value at the beginning of the first bunch train, then falls and reaches a plateau at lower values that extends up to the end of the bunch train. In the second bunch train the the behaviour is the same, but the maximum correction value is reduced compared to the first bunch train. In total, the pedestal correction follows the same behaviour as the pedestal itself (see Figure 7.3 (b)).

As a function of time, the correction slowly decreases over the course of the run. This is best seen at the maximum of the first bunch train, which falls by about 30% from the start to the end of the run. This reflects the behaviour of the pile-up contribution to the trigger tower signal, which follows the luminosity decay over time.

The effect of the pedestal correction on the trigger performance is demonstrated in Figure 9.6. It shows the rate of an L1 E_T^{miss} trigger as a function of the instantaneous luminosity for two data taking runs in 2015. In one of the runs the pedestal correction is active, while it is de-activated for the other run. Otherwise, the L1Calo configuration is the same in both runs.

For the run without the pedestal correction, the trigger rate rises sharply as a function of the luminosity, as was already seen in Run-1 (see Figure 7.4). The non-linearity of the curve points to large fake rates caused by the pile-up. For the other run, the pedestal correction subtracts the pile-up contribution, resulting in a linear behaviour of the trigger rate with luminosity.

A further comparison of the same two runs is shown in Figure 9.7. It shows the rate of a further L1 E_T^{miss} trigger as a function of the BC number for one bunch train. Figure 9.7 (a) depicts the result for the run with pedestal correction, while (b) shows it for the run without pedestal correction. For the later case, an increase of the trigger rate is observed for the BCs at the beginning of the bunch train. With the pedestal correction active, the effect is removed and the trigger rate is constant as a function of the BC number.

Next to the E_T^{miss} the pedestal correction also reduced fake rates for forward jets and low threshold multi-jet L1 trigger items. Due to this, the Pedestal Correction makes it possible to maintain low trigger thresholds and reduced prescale factors for these signatures even at the high instantaneous luminosities reached in Run-2 of the LHC.

[4]A lumi block is defined as a period of time with approximately constant operating conditions, like luminosity, detector module voltages, gas pressure, etc. . It typically lasts for one minute, but shorter lmi blocks are possible if changes in the detector configuration are performed.

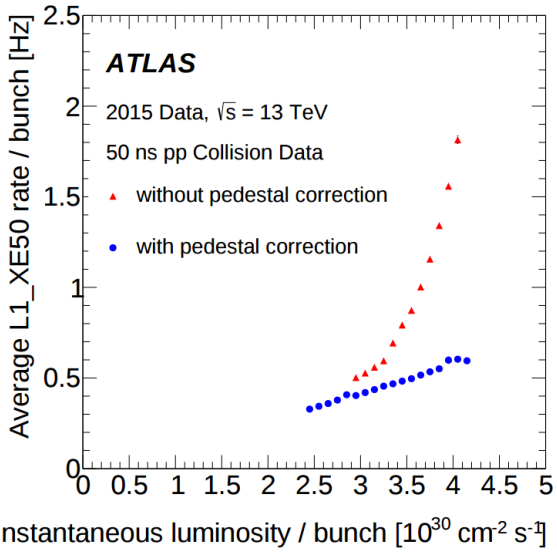

Figure 9.6: Rate of a L1 trigger item that requires a missing transverse energy of at least 50 GeV as a function of the instantaneous luminosity. Both quantities are normalised to the number of colliding bunches per beam. Measurements from two runs are shown, one with Pedestal Correction active (blue dots) and one with Pedestal Correction inactive (red triangles). [41]

9.4 Separate Energy Calibration for CP and JEP

The separation of the energy calibration LUT allows independent calibrations for the CP and JEP systems. For the CP, the energy resolution of the E_T output is improved by changing the scale of the LUT output from 1 to 0.5 GeV for each count [41]. This allows more refined isolation cuts in the CP algorithms, especially considering the change to energy dependent isolation thresholds for Run-2 (see Section 4.4.2). While reducing the LUT scale also lowers the maximum energy that can be detected for each trigger tower, the resulting maximum of 127.5 GeV is sufficient for electromagnetic trigger items.

For the JEP, a non-linear energy calibration that is optimised to improve the jet energy resolution and the E_T^{miss} trigger rate has been studied. Because the resulting calibration is too sensitive to changes in the operating conditions of the LHC, this calibration is not currently in use in the PreProcessor, and the JEP-LUT remains in the standard Run-1 configuration described in Section 5.2.3.

9.5 BCID of Saturated Signals

The BCID algorithms for saturated signals underwent several changes over the course of Run-2, which are described in more detail in Section 7.4. Table 9.1 lists the performance of these setups, in terms of the rate of mis-identification of saturated signals that resulted in early triggers. As saturated signals always result in objects

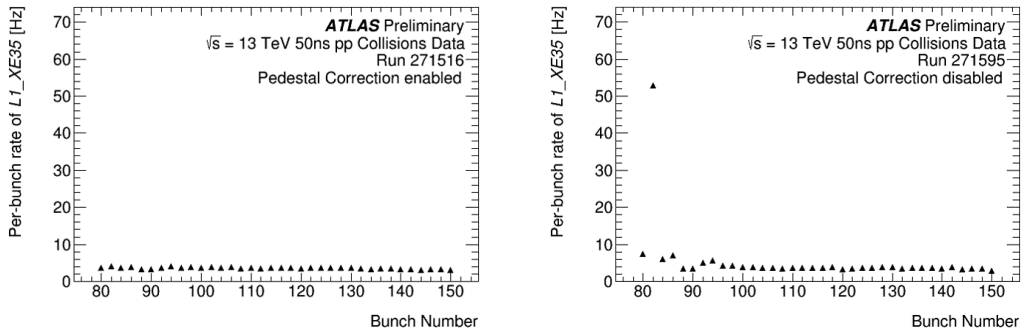

(a) **(b)**

Figure 9.7: Bunch-normalised rate of a L1 trigger item that requires a missing transverse energy of at least 35 GeV as a function of the BC number for one bunch train. Results from two runs are shown, one with Pedestal Correction active (a), and one with Pedestal Correction inactive (b). [62]

passing L1Calo thredholds, this type of mis-identification can lead to early triggers where the information of the actual BC being lost, with the consequences described in Section 4.3.1.

Saturated BCID Setup	$\mathcal{L}$ [fb^{-1}]	Early Trigger Rate [per fb^{-1}]
Legacy (Run-1 Setup)	4.0	30
Improved Decision Logic	31.1	1.3
Sat80BCID Algorithm	5.1	0

Table 9.1: Number of events that were triggered one BC too early because of the Saturated BCID logic in the PreProcessor during Run-2.

At the start of Run-2, the algorithms still operated in the same way as in Run-1. As described earlier, this *legacy setup* results in the BCID logic mis-identifying the saturated trigger tower at an early BC due to the autocorrelation filters. This was fixed by the improvements to the BCID Decision Logic, which ignore the Peak Finder algorithm for highly saturated pulses. While the mis-identification rate is reduced by this change, it remains at a non-zero value of about 1.3 events per fb^{-1}.

The remaining inefficiency is a consequence of the trivial setup of the SatBCID algorithm. Because it always identifies the second saturated sample, the BCID Decision Logic can not be set to a configuration where only the SatBCID algorithm is active without the sufficient amount of saturated signals. Thus, a small chance remains for the Autocorrelation Filter to mis-identify a BC.

The Sat80BCID algorithm was activated for the last weeks of 2016, after the calibration was established by operating the algorithm in monitoring mode (see Section 7.4.2). This allows to adapt the thresholds in the BCID Decision logic, because the Sat80BCID correctly identifies the peak already for fewer saturated

samples than the SatBCID algorithm. Since switching to this configuration, no additional mis-identified events have been found.

10 Summary

The increased centre-of-mass energy and luminosity due to the upgrade of the LHC in LS1 impose additional challenges to the ATLAS TDAQ systems for Run-2. Several TDAQ systems have thus been upgraded during LS1, to ensure low thresholds for unprescaled trigger items and a high trigger efficiency while keeping the total L1 trigger rate below the maximum of 100 kHz.

For the L1Calo trigger, the increased pile-up enhances the noise on the trigger tower signals, reducing the detection efficiency for low E_T signals. Out-of-time pile-up effects result in a large increase of the rate of triggers that rely on many trigger towers, like those requiring missing transverse energy and multi-jet signatures. Furthermore, the increased centre-of-mass energy leads to a higher probability for signals to saturate the ADCs in the PreProcessor hardware.

To address these challenges, the MCM in the L1Calo PreProcessor has been exchanged by the improved nMCM. Its implementation makes use of modern signal processing components to increase the signal sampling rate to 80 MHz and to port the signal processing algorithms from the hard-coded ASIC used in Run-1 to a flexible FPGA. This allows to both enhance the existing Run-1 signal processing and to include additional algorithms that further improve the PreProcessor performance.

All nMCMs were tested for their proper functionality before installation at CERN. A high production efficiency resulted in 90% of the modules passing all tests without showing any errors. Since the nMCMs have been installed at CERN, only four modules showed problems during two years of operation.

The signal processing improvements gained with the nMCM significantly increase the performance of the L1Calo trigger, as demonstrated already in two years of operation. The minimisation of cross-talk between the digital and analogue signal paths on the PPM reduces the noise on the analogue input signals. For low energy signals, the updated digital filter implementation significantly improves the detection efficiency in the pile-up dominated environment provided in Run-2. In addition, independent energy calibrations for the JEP and CP systems provide flexibility for separate optimisations and algorithm improvements in the respective systems. For saturated signals, the association to the correct BC is assured by the increased flexibility of the BCID Decision Logic and the new 80 MHz saturated BCID algorithm. The pile-up subtraction provided by the dynamic pedestal correction removes a large contribution of fake triggers, making it possible to reach the maximum rate requirement of 100 kHz without impacting important physics trigger items.

The nMCM thus significantly improves the event selection of the L1Calo trigger, making it possible to efficiently operate the system for many years to come.

A Lists

A.1 List of Figures

3.1 CERN Accelerator Complex . . . 19
3.2 ATLAS Detector . . . 23
3.3 ATLAS Inner Detector . . . 24
3.4 ATLAS Calorimeter . . . 26
3.5 EMB Module Structure . . . 27
3.6 Tile Module Structure . . . 29
3.7 ATLAS Muon Spectrometer . . . 31
3.8 ATLAS Magnets . . . 32
3.9 ATLAS TDAQ System . . . 34

4.1 ATLAS L1Calo Overview . . . 40
4.2 Trigger Towers in Electromagnetic Calorimeter . . . 41
4.3 L1Calo Input Signal Path . . . 42
4.4 Trigger Tower Pulse Shapes . . . 43
4.5 PreProcessor Functionality . . . 44
4.6 PreProcessor Module . . . 47
4.7 PreProcessor Module Coverage in Hadronic Calorimeter . . . 48
4.8 AnIn Board Block Diagram . . . 48
4.9 Multichip Module . . . 49
4.10 ReM FPGA Blockdiagram . . . 51
4.11 CP Algorithm . . . 53
4.12 JEP Algorithm . . . 55

5.1 New Multichip Module . . . 57
5.2 nMCM Power Conversion . . . 59
5.3 Overview of Initial CALIPPR FPGA Implementation . . . 61
5.4 CALIPPR FPGA Input Timing Diagram . . . 62
5.5 Fine Timing . . . 63
5.6 LUT and Drop Bits Implementation . . . 65
5.7 Saturated BCID Functionality . . . 67
5.8 BCID Decision Logic . . . 68
5.9 LVDS Data Formats . . . 69
5.10 Bunch-Crossing Multiplexing Scheme . . . 70
5.11 Initial Readout Format . . . 71
5.12 Signal Generator Circuit . . . 75
5.13 Clock Generation in CALIPPR FPGA . . . 76

5.14 Serial Interface Illustration 77

6.1 Single-Board Test Setup 81
6.2 Full-Crate Test Setup 82
6.3 nMCM Temperature in PreProcessor Crate 83
6.4 Rate Metering Test Results 85
6.5 DAC Scan Linear Fit 90
6.6 DAC Scan Noise Results 91
6.7 Signal Generator Pedestal Scan Linear Fit 92
6.8 Overview of DAC Scan results for all nMCMs 94
6.9 Faulty nMCMs in DAC Scan 95
6.10 80 MHz Latching Test Example 96
6.11 40 MHz Latching Test Signals 97
6.12 CALIPPR Input Timing Diagrams 99
6.13 Pulse Reconstruction Test 100
6.14 40 MHz Latching Test Failure 100

7.1 Block Diagram of the BCID Logic 103
7.2 Calorimeter Noise 105
7.3 Pedestal Distortion by Pile-up 107
7.4 Trigger Rates in Run-1 109
7.5 Pedestal Correction Averaging Logic 110
7.6 Pedestal Correction Synchronisation 112
7.7 BCID for Saturated Pulses 114
7.8 Updated BCID Decision Logic 116
7.9 BCID Decision in Run-2 117
7.10 Sat80BCID Algorithm 118

8.1 Updated Readout Format 120
8.2 Sat80BCID Readout Bits 121
8.3 80 MHz ADC Readout 122

9.1 Noise MCM vs nMCM 125
9.2 PreProcessor System Noise 126
9.3 Correlated Noise MCM vs nMCM 126
9.4 Autocorrelation Filter Performance 127
9.5 Pedestal Correction Monitoring 128
9.6 Pedestal Correction Rate Reduction 130
9.7 Pedestal Correction Rate Reduction Per BC 131

A.2 List of Tables

2.1 Overview of Particles in Standard Model 16

3.1 Calorimeter Energy Resolution 26

6.1 nMCM Configuration Test Patterns 84
6.2 LVDS Transmission Test Playback Patterns 87
6.3 Event Readout Test Modes 88
6.4 DAC Scan Limits 90
6.5 nMCMs failing DAC Scan Test 93
6.6 Nominal Latching Configuration 99
6.7 nMCM Production Test Results 102

7.1 Filter Coefficient Definition in Run-1 and Run-2 106

9.1 Saturated BCID Performance 131

B Bibliography

[1] J. Dalton. "A New System of Chemical Philosophy". Manchester, 1808.

[2] J. J. Thomson. "Cathode Rays". *Philosophical Magazine*, **44**:293, 1897.

[3] E. Rutherford. "The Scattering of α and β Particles by Matter and the Structure of the Atom". *Philosophical Magazine*, Series 6, **21**:669–688, 1911.

[4] E. Rutherford, T. Royds. "The Nature of the α Particle from Radioactive Substances". *Philosophical Magazine*, **17**:281–6, 1909.

[5] F. Soddy. "The Atomic Weight of "Thorium" Lead". *Nature*, **98** (2468):469, 1917.

[6] F. Soddy. "Intra-atomic Charge". *Nature*, **92** (2301):399–400, 1913.

[7] F. W. Aston. "Isotopes". London: E. Arnold, 1922.

[8] F. W. Aston. "Mass-Spectra and Isotopes". London: E. Arnold, 1933.

[9] J. J. Thomson. "XIX. Further experiments on positive rays". *Philosophical Magazine*, Series 6, **24** (140):209, 1912.

[10] J. J. Thomson. "LXXIII. Rays of positive electricity". *Philosophical Magazine*, Series 6, **20** (118):752, 1910.

[11] O. Hahn, F. Strassmann. "Über den Nachweis und das Verhalten der bei der Bestrahlung des Urans mittels Neutronen entstehenden Erdalkalimetalle". *Naturwissenschaften*, **27** (1):11–15, 1939.

[12] L. Meitner, O. R. Frisch. "Disintegration of Uranium by Neutrons: a New Type of Nuclear Reaction". *Nature*, **143** (3615):239, 1939.

[13] M. L. E. Oliphant, A. R. Kempton, Lord Rutherford. "The Accurate Determination of the Energy Released in Certain Nuclear Transformations". *Proceedings of the Royal Society A*, **149** (867):406–416, 1935.

[14] H. A. Bethe. "Energy Production in Stars". *Physical Review*, **55** (5):434–56, 1939.

[15] S. H. Neddermeyer, C. D. Anderson. "Note on the Nature of Cosmic-Ray Particles". *Physical Review*, **51** (10):884–886, 1937.

[16] J. C. Street, E. C. Stevenson. "New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron". *Physical Review*, **52** (9):1003–1004, 1937.

[17] G. D. Rochester, C. C. Butler. "Evidence for the Existence of New Unstable Elementary Particles". *Nature*, **160** (4077):855–857, 1947.

[18] V. D. Hopper, S. Biswas. "Evidence Concerning the Existence of the New Unstable Elementary Neutral Particle". *Physical Review*, **80** (6):1099, 1950.

[19] M. Gell-Mann. "The Interpretation of the New Particles as Displaced Charge Multiplets". *Il Nuovo Cimento*, **4**:848–866, 1956.

[20] M. Gell-Mann. "A Schematic Model of Baryons and Mesons". *Physics Letters*, **8** (3):214–215, 1964.

[21] G. Zweig. "An SU(3) Model for Strong Interaction Symmetry and its Breaking". *CERN Report*, No.8182/TH.401, 1964.

[22] G. Zweig. "An SU(3) Model for Strong Interaction Symmetry and its Breaking: II". *CERN Report*, No.8419/TH.412, 1964.

[23] S. L. Glashow. "Partial Symmetries of Weak Interactions". *Nuclear Physics*, **22**:579–588, 1961.

[24] A. Salam. "Elementary Particle Physics". *Proc. 8th Nobel Symp.*, Almqvist and Wiksell, Stockholm, 1968.

[25] S. Weinberg. "A Model of Leptons". *Physical Review Letters*, **19** (21):1264–1266, 1967.

[26] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". *Physics Letters B*, **716**:1–29, 2012.

[27] CMS Collaboration. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". *Physics Letters B*, **716**:30, 2012.

[28] C. Patrignani *et al.* (Particle Data Group). "2016 Review of Particle Physics". *Chinese Physics C*, **40**:100001, 2016.

[29] L. Bergstrom. "Non-baryonic dark matter: Observational evidence and detection methods". *Reports on Progress in Physics*, **63** (5):793–841, 2000.

[30] L. Canetti, M. Drewes, M. Shaposhnikov. "Matter and Antimatter in the Universe". *New J.Phys.*, **14**: 095012, 2012.

[31] P. Bryant, L. Evans. LHC Machine. *JINST **3** (2008) S08001.*

[32] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST **3** (2008) S08003*.

[33] CMS Collaboration. The CMS experiment at the CERN LHC. *JINST **3** (2008) S08004*.

[34] The LHCb Collaboration. The LHCb Detector at the LHC. *JINST **3** (2008) S08005*.

[35] ALICE Collaboration. The ALICE experiment at the CERN LHC. *JINST **3** (2008) S08002*.

[36] "CERN's Accelerator Complex", OPEN-PHO-CHART-2013-001. `https://cds.cern.ch/record/1621583`, October 2013.

[37] M. Bajko et. al. "Report of the task force on the incident of 19 September 2008 at the LHC". LHC-PROJECT-Report-1168, 2012.

[38] ATLAS Experiment - Public Results (Run-1 Luminosity). `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults`.

[39] ATLAS Experiment - Public Results (Run-2 Luminosity). `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2`.

[40] M. Capeans et al. "ATLAS Insertable B-Layer Technical Design Report". ATLAS-TDR-19, 2010.

[41] ATLAS Collaboration. "Performance of the ATLAS Trigger System in 2015". CERN-EP-2016-241, arXiv:1611.09661, 2016.

[42] ATLAS Collaboration. "Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System". ATLAS-TDR-023, 2013.

[43] R. Achenbach et al. "The ATLAS Level-1 Calorimeter Trigger". *JINST **3** (2008) P03001*.

[44] V. Andrei. "The Data Path of the ATLAS Level-1 Calorimeter Trigger Pre-Processor". Dissertation, HD-KIP 10-67, 2010.

[45] L1Calo Collaboration. "Cable Mappings and Crate Layouts from Receivers to Processors", ATL-DA-ES-0036. `https://edms.cern.ch/document/399348/2.2`, 2007.

[46] N. J. Buchanan et al. "ATLAS liquid argon calorimeter front end electronics". *JINST **3** (2008) P09003*.

[47] ATLAS Collaboration. "ATLAS tile calorimeter: Technical Design Report". ATLAS-TDR-3, 1996.

[48] B. Asman et al. "The ATLAS Level-1 Calorimeter Trigger: PreProcessor implementation and performance". *JINST* ***7*** *(2012) P12008.*

[49] P. Hanke. "The Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger", Version 1.1. October 2009.

[50] T. Toifl, P.Moreira, A. Marchioro and R. Vari. "4-channel rad hard delay generation ASIC with 1ns timing resolution for LHC". *IEEE Trans. Nucl. Sci.* ***46*** *(1999) 139.*

[51] Analog Devices. "AD9218 10-bit, 40/65/80/105 MSPS 3V Dual Analog-to-Digital Converter Data Sheet (Rev. C)". 2006.

[52] D. Husmann et. al. "Pre-Processor Asic User and Reference Manual". 2011.

[53] V. Andrei. "Tile Rear-Extension Module Specifications". 2016.

[54] L1Calo Collaboration. L1Calo Twiki (internal page). `https://twiki.cern.ch/twiki/bin/viewauth/Atlas/L1CaloUpgrade`, 2015.

[55] K. Schmitt, P. Hanke. "New Multichip Module Specifications". 2013.

[56] L1Calo Hardware Database. `http://www.kip.uni-heidelberg.de/atlas/db/DbPPr/welcome.html`, June 2009.

[57] A. C. Daniells. "Pile-Up Suppression in the ATLAS Level 1 Calorimeter Trigger and Searches for Higgs Boson Pair Production". Dissertation, 2016.

[58] J. Garvey and D. Rees. "Bunch Croessing Identification for the ATLAS Level-1 Calorimeter Trigger". ATL-DAQ-96-051, 1996.

[59] W. E. Cleland and E. G. Stern. "Signal processing considerations for liquid ionization calorimeters in a high rate environment". *Nucl. Inst. Meth.*, **A338**:467–497, 1994.

[60] F. Bartels. "Kalibration des Finite-Impulse-Response-Filters im PreProcessor des ATLAS Level-1 Calorimeter Triggers für den LHC Run-2". Bachelor Thesis, HD-KIP 15-89, 2015.

[61] "L1Calo Tier-0 Monitoring Tools (internal page)". `https://cds.cern.ch/record/1621583`.

[62] ATLAS Experiment - Public Results (Level-1 Calorimeter Trigger). `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults`.

[63] H. Martinez, N. Besson and N. Morange. "Update on the study and validation of the BCID of the L1Calo at very high transverse energies". ATL-COM-LARG-2012-025, 2012.

Acknowledgements

I thank Prof. Dr. Hans-Christian Schultz-Coulon for giving me with the opportunity to work for many years on an interesting topic, and for providing a great environment at our working group at KIP. I never felt like I missed anything I required to do my work, and I realise this is in great part to a lot of organisational work 'behind the scenes'.

I want to thank Prof. Dr. Norbert Herrmann for finding time to referee my thesis while taking on additional responsibilities.

A huge thanks to Victor Andrei. He taught me a lot about what it means to work with hardware and firmware, about both of which I knew only very little when I first started my PhD. He was always ready to help when I got stuck. I also learned to appreciate his (sometimes wicked) sense of humor, and the discussions about 'real life' we had over one beer or another.

Furthermore I want to thank Martin Wessels, who provided a lot of support in getting the nMCM to run and properly calibrated. His experience with the intricacies of ATLAS operations was especially valueable.

Many people from L1Calo supported me during my time in the collaboration. Most of all I want to thank Bruce Barnett and Murrough Landon for many lessons on L1Calo functionality and helpful advice on being a better On-Call.

For proof-reading my thesis I want to thank Victor Andrei, Martin Wessels, Rainer Stamen, Claire Antel, Hanno Meyer zu Theenhausen, Manuel Geisler, Sebastian Weber and Falk Bartels.

Besides these professional acknowledgements I want to take the time to thank quite a few people on a personal level. In all endeavours of my life, I have always valued the time spent with other people in a good atmosphere the most. My most treasured memories from my time at KIP (a total of 5.5 years by now!) are thus of times we spent talking to each other, laughing together, having parties, going for ice cream at the Uni-Shop or just generally enjoying life as a group.

During my time at CERN, I had two great office mates in the form of Claire and Silvia. Claire, having been a summer student, introduced me to many people around CERN, easing my start into the social life there by a lot. Silvia brought a lot of energy into the group and her detailed questions about L1Calo often had me doubt my own knowledge of the system. I also want to thank Ismet for being a companion for many lunches, dinners, parties and boardgame club meetings - and for 'Everything is Awesome'.

Konrad and Huangshan were my office mates for the majority of my time spent at KIP. Thanks to both of you for tolerating my (sometimes random) blabbering and being the great guys you are. Good luck to both of you for your upcoming thesis

writing!

I also want to thank a certain group of BIOQUANT people for many, many coffee breaks. Michi, Pete, Verena, Priyata, Sam, Vineet, and all the others who came and went over the last few years: I am really happy to have found a group of friends like you, and want to thank you for keeping me rooted in the world outside the bubble of particle physics.

I want to thank other KIP people from F8 and F11 in no particular order: Hanno, for your biceps and philosophical discussions; Stanislav, for many hikes around the Heidelberg mountains; Manuel, for being sexy; Daniel, for discussions of One Piece and life as a PhD student; Alessandra, for being a partner in suffering during thesis writing; Julia, for introducing us to water skiing; Patrick, for calling me "the prophet of the group" once; Tobias, for being a good guy behind a facade of (sometimes hilarious) jokes and jabs; Thomas, for fun talks over a few (too many) drinks.

Finally I want to thank my family, for the support you all provided me in your own ways during almost three decades of my life.