

**Simulation Studies of Digital Filters for the
Phase-II Upgrade of the
Liquid-Argon Calorimeters of the
ATLAS Detector at the High-Luminosity LHC**

DISSERTATION

zur Erlangung des akademischen Grades

Doctor rerum naturalium

(Dr. rer. nat.)

vorgelegt

**dem Bereich Mathematik und Naturwissenschaften
der Technischen Universität Dresden**

von

M. Sc. Nico Madysa

geboren am 14. September 1990 in Hoyerswerda.

Eingereicht am 17. Februar 2020.

Die Dissertation wurde in der Zeit vom Januar 2016 bis Februar 2020
im Institut für Kern- und Teilchenphysik angefertigt.



Eingereicht am 17. Februar 2020.
Verteidigt am 30. September 2020.

1. Gutachter: Prof. Dr. Arno Straessner
2. Gutachter: Prof. Peter Krieger, Ph. D.

Kurzreferat

Am Large Hadron Collider und am ATLAS-Detektor werden umfangreiche Aufrüstungsarbeiten vorgenommen. Diese Arbeiten sind in mehrere Phasen gegliedert und umfassen unter Anderem Änderungen an der Ausleseelektronik der Flüssigargonkalorimeter; insbesondere ist es geplant, während der letzten Phase ihren Primärpfad vollständig auszutauschen. Die Elektronik besteht aus einem analogen und einem digitalen Teil: während ersterer die Signale verstärkt und sie zur leichteren Abtastung verformt, führt letzterer einen Algorithmus zur Energierekonstruktion aus. Beide Teile müssen während der Aufrüstung verbessert werden, damit der Detektor interessante Kollisionsereignisse präzise rekonstruieren und uninteressante effizient verwerfen kann.

In dieser Dissertation werden Simulationsstudien präsentiert, die sowohl die analoge als auch die digitale Auslese der Elektromagnetischen Kalorimeter optimieren. Die Korrektheit der Simulation wird mithilfe von Kalibrationsdaten geprüft, die im sog. Run 2 des ATLAS-Detektors aufgenommen worden sind. Der Einfluss verschiedener Parameter der Signalverformung auf die Energieauflösung wird analysiert und die Nützlichkeit einer erhöhten Abtastrate von 80 MHz untersucht. Des Weiteren gibt diese Arbeit eine Übersicht über lineare und nichtlineare Energierekonstruktionsalgorithmen. Schließlich wird eine Auswahl von ihnen hinsichtlich ihrer Leistungsfähigkeit miteinander verglichen.

Es wird gezeigt, dass ein Erhöhen der Ordnung des Optimalfilters, der gegenwärtig verwendete Algorithmus, die Energieauflösung um 2 bis 3 % verbessern kann, und zwar in allen Regionen des Detektors. Der Wiener Filter mit Vorwärtskorrektur, ein nichtlinearer Algorithmus, verbessert sie um bis zu 10 % in einigen Regionen, verschlechtert sie aber in anderen. Ein Zusammenhang dieses Verhaltens mit der Wahrscheinlichkeit fälschlich detektierter Kalorimetertreffer wird aufgezeigt und mögliche Lösungen werden diskutiert.

Abstract

The Large Hadron Collider and the ATLAS detector are undergoing a comprehensive upgrade split into multiple phases. This effort also affects the liquid argon calorimeters, whose main readout electronics will be replaced completely during the final phase. The electronics consist of an analog and a digital portion: the former amplifies the signal and shapes it to facilitate sampling, the latter executes an energy reconstruction algorithm. Both must be improved during the upgrade so that the detector may accurately reconstruct interesting collision events and efficiently suppress uninteresting ones.

In this thesis, simulation studies are presented that optimize both the analog and the digital readout of the Electromagnetic Calorimeters. The simulation is verified using calibration data that has been measured during Run 2 of the ATLAS detector. The influence of several parameters of the analog shaping stage on the energy resolution is analyzed and the utility of an increased signal sampling rate of 80 MHz is investigated. Furthermore, a number of linear and non-linear energy reconstruction algorithms is reviewed and the performance of a selection of them is compared.

It is demonstrated that increasing the order of the Optimal Filter, the algorithm currently in use, improves energy resolution by 2 to 3 % in all detector regions. The Wiener filter with forward correction, a non-linear algorithm, gives an improvement of up to 10 % in some regions, but degrades the resolution in others. A link between this behavior and the probability of falsely detected calorimeter hits is shown and possible solutions are discussed.

Contents

1	Introduction	1
2	An Overview of High-Energy Particle Physics	3
2.1	The Standard Model of Particle Physics	3
2.2	Verification of the Standard Model	7
2.3	Beyond the Standard Model	8
3	LHC, ATLAS, and the Liquid-Argon Calorimeters	11
3.1	The Large Hadron Collider	11
3.2	The ATLAS Detector	16
3.3	The ATLAS Liquid-Argon Calorimeters	24
4	Upgrades to the ATLAS Liquid-Argon Calorimeters	35
4.1	Physics Goals	35
4.2	Phase-I Upgrade	37
4.3	Phase-II Upgrade	41
5	Noise Suppression With Digital Filters	49
5.1	Terminology	49
5.2	Digital Filters	54
5.3	Wiener Filter	56
5.4	Matched Wiener Filter	59
5.5	Matched Wiener Filter Without Bias	61
5.6	Timing Reconstruction, Optimal Filtering, and Selection Criteria	62
5.7	Forward Correction	65
5.8	Sparse Signal Restoration	67
5.9	Artificial Neural Networks	71
6	Simulation of the ATLAS Liquid-Argon Calorimeter Readout Electronics	75
6.1	AREUS	75
6.2	Hit Generation and Sampling	79
6.3	Pulse Shapes	83
6.4	Thermal Noise	86
6.5	Quantization	91

6.6	Digital Filters	93
6.7	Statistical Analysis	96
7	Results of the Readout Electronics Simulation Studies	99
7.1	Statistical Treatment	99
7.2	Simulation Verification Using Run-2 Data	103
7.3	Dependence of the Noise on the Shaping Time	105
7.4	The Analog Readout Electronics and the ADC	108
7.5	The Optimal Filter (OF)	111
7.6	The Wiener Filter	113
7.7	The Wiener Filter with Forward Correction (WFFC)	114
7.8	Final Comparison and Conclusions	118
8	Conclusions and Outlook	123
A	Derivation of the Liquid-Argon Electronics Pulse Shape	127
	List of Figures	131
	List of Tables	133
	Glossary	135
	Bibliography	141
	Acknowledgments	165
	Versicherung	167

Chapter 1

Introduction

Throughout the field of high-energy particle physics, the Standard Model is considered the best-proven theory to explain the fundamental interactions of the universe. Since the discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012, all of its predictions have been verified successfully. Nonetheless, several issues remain: the Higgs boson self-coupling has not yet been observed; there are still open questions regarding the electroweak phase transition; the measured CP violation is too weak to explain the matter–antimatter imbalance observed in the universe; Standard Model extensions like *supersymmetry*, which predict multiple Higgs bosons to exist, cannot be ruled out; and finally, there are multiple phenomena (like dark matter) that are explained only insufficiently by the Standard Model. Thus, the search for physics beyond the Standard Model continues.

The LHC is currently the most powerful particle accelerator in the world. Two general-purpose experiments – ATLAS and CMS – as well as several specialized experiments make use of the proton–proton and heavy-ion collisions at the LHC to realize a rich physics program. The ATLAS experiment has, among others, two important goals: high-precision measurements of Standard Model parameters and processes (like the W mass or the Higgs self-coupling) and the search for physics beyond the Standard Model. To achieve these goals, the detector must identify the final state of each collision with high accuracy. Particularly important is the identification and reconstruction of electrons, photons, and particle jets. For this task, the most vital subsystem of the detector are the liquid argon (LAr) calorimeters.

The LHC and its experiments are planned to undergo an upgrade in two phases: one that has started in 2019 and one scheduled for 2025. The goal of this is to increase the rate at which the detectors record data and to improve the data quality. The LAr calorimeters are also affected by this: the accuracy of their energy measurements must be improved and they must send data to the ATLAS trigger at higher resolution, so that more “uninteresting” events may be suppressed. For the second upgrade phase, it is planned to completely replace both the analog and the digital readout electronics of the calorimeters on the main readout path.

The readout electronics work as follows: Each time a particle passes through a calorimeter cell, a voltage pulse is induced in the analog electronics. The digital electronics sample this pulse and pass the samples through an *energy reconstruction* algorithm, which estimates the deposited energy based on the amplitude of the pulse. This algorithm must be both fast and

accurate, even in the face of strong overlap between pulses that are caused by subsequent events. The upgrade of the electronics opens the opportunity to improve this algorithm. With the new hardware being faster and having more memory available, more complex algorithms may be considered.

Extensive studies are required to determine how both the analog and the digital readout electronics can be optimized in terms of energy resolution. These studies are necessarily based on simulations of the future upgraded electronics, as measurements do not exist yet. The computer program AREUS is used for this purpose. It performs a detailed simulation of the signal digitization and energy reconstruction algorithm. It also allows the user to pass in custom pulse shapes and electronic-noise descriptions, which may be obtained from dedicated simulations. In contrast to other simulations, AREUS is *free-running*, i.e. the electronics are simulated not only for disjunct events, but for contiguous runs of data taking. This makes it possible to study long-term effects and correlations in the energy reconstruction.

The analog electronics may be studied by transforming the pulse shapes and noise descriptions in certain, well understood ways. If the transformation depends on some parameter, the data can be transformed for each possible value and passed to the simulation. The results can then be compared to study the impact of the parameter on the energy resolution or other figures of merit. Energy reconstruction algorithms may be compared similarly; where an algorithm depends on hyperparameters, these may be varied to analyze their influence.

It is important to perform this optimization carefully, to study the quality of the relationship between parameters and energy resolution, and to look for unexpected influences. Otherwise, solutions may turn out to be statistical fluctuations or the result of *over-training*. A sign of the latter would be, for example, if an algorithm performs well in some regions of the detector, but considerably worse in others.

In this thesis, parameters of the analog readout electronics are tuned for an optimal energy resolution. A number of energy reconstruction algorithms is presented, their hyperparameters are optimized, and their performance is compared to the Optimal Filter, the current state-of-the-art algorithm.

In chapter 2, an overview of the current state of high-energy particle physics is given to motivate this research. In chapter 3, the LHC, the ATLAS detector, and the LAr calorimeters as the underlying experimental setup are presented. In chapter 4, the upgrade project in general and the concrete changes of the LAr calorimeter readout in particular are described. In chapter 5, a brief introduction to the theory of linear filters is given. The Optimal Filter and the Wiener filter as the most important examples of linear filters are introduced. Furthermore, a number of non-linear approaches to the problem of energy reconstruction are reviewed. In chapter 6, the setup of the simulation studies is documented. The data that has been used is cited, the most important algorithms of the simulation are described, and the figures of merit are defined. In chapter 7, the studies are presented, their results analyzed, and conclusions are drawn. A summary of this work and an outlook on potential future studies are given in chapter 8.

Chapter 2

An Overview of High-Energy Particle Physics

The field of particle physics at high energies has seen a rapid development since its beginning in the second half of the twentieth century. At its core, it deals with interactions of subatomic particles at energies higher than the particles' rest energies. The investigated phenomena are hence intrinsically subject to both quantum physics and special relativity. It is the equal importance of both theories that distinguishes this field of physics from most others.

This chapter aims to give a brief overview over high-energy particle physics. Section 2.1 begins by introducing the Standard Model, a theory that consolidates our current understanding of particle physics. Section 2.2 details the ways in which the Standard Model has been proven so far. Section 2.3 concludes by showing the limits of the Standard Model and by naming a select few ongoing experimental efforts to find new physics in contradiction to it.

2.1 The Standard Model of Particle Physics

The Standard Model (SM) currently is the best-verified and most comprehensive theory of elementary particle physics. It emerged in the early 1970s [1] from two independently evolving theories: quantum chromodynamics [2–5] and the electroweak theory [6–8] extended by the Brout–Englert–Higgs mechanism [9–14]. Whenever new particles have been discovered since its inception, it has been extended rather than replaced by another theory [15, 16].

The Standard Model is a *quantum field theory*, i.e. a relativistic field theory in which all observable particles are interpreted as quantized excitations of their underlying fields. Additionally, the three¹ fundamental interactions described by the Standard Model are introduced by imposing certain *local gauge symmetries*, i.e. by requiring that the theory be invariant under certain continuous, location-dependent transformations of its fields. The symmetries are commonly denoted by the Lie groups of these transformations.

¹The fourth fundamental interaction, gravitation, is generally neglected, both because it is vastly weaker than the other interactions on the subatomic scale, and because its quantization still is an unsolved problem [17].




Figure 2.1: Particle content of the Standard Model of particle physics [19]. The fermions are arranged by family and generation, the gauge bosons by symmetry group. Each particle is annotated with its mass in GeV (top) and gauge bosons or the charges that they couple to (bottom). The fermions are additionally annotated with their color C (left) and electric charge Q (right), if they carry any. The hand symbol signifies a coupling to the weak isospin L .

Concretely, the symmetries of the Standard Model are:

$$SU(3)_C \times SU(2)_L \times U(1)_Y. \quad (2.1)$$

Following convention, each symmetry is subscripted with the charge it conserves according to Noether's theorem [18]: The color charge C , the weak isospin L , and the weak hypercharge Y . However, due to electroweak symmetry breaking (EWSB), L and Y are *not* conserved, only a combination of them: $Q := \frac{1}{2}Y + L_z$, the conventional electric charge. This is described in more detail in section 2.1.3.

Figure 2.1 shows the particle content of the Standard Model. The particles are usually grouped according to their spin: *fermions* with spin $\frac{1}{2}$, which are commonly understood as matter particles; *vector bosons* with spin 1, which are mediators of the fundamental interactions; and *scalar bosons* with spin 0, of which there is only one after EWSB.

2.1.1 Fermions

The fermions are categorized into two families: *quarks*, which carry color charge, and *leptons*, which do not. Each family contains six *flavors* of particles, which are arranged into three *generations*. They are arranged such that each generation is identical w.r.t. the electric and color charges of its constituents.

Quarks, as carriers of color charge, are subject to the confinement effect and cannot appear as free particles. Instead, they can only be found in color-singlet bound states called *hadrons*.

These are further divided into *mesons* (quark–antiquark pairs) and *baryons* (quark triples). The two lightest baryons, protons and neutrons, are able to form stable atomic nuclei and make up the majority of visible matter in the universe.

Leptons can be further categorized based on their electric charge. The three neutral flavors are called *neutrinos*. Although they are known to have a small, non-zero mass², the Standard Model treats them as massless. They only participate in the weak interaction, which makes it exceedingly difficult to detect them. The three *charged leptons*, by contrast, are additionally subject to electromagnetism and much easier to detect. Only the lightest among them, the electron, is stable. Together with protons and neutrons, it forms atoms, the fundamental building blocks of all ordinary matter.

2.1.2 Vector Bosons

The vector, or gauge bosons are naturally divided by the interactions they mediate. In the Standard Model, each particle carries certain charges and each gauge boson only couples to certain charges. Hence, not all particles experience all interactions.

The *gluons* are the mediators of the strong interaction, modeled by the $SU(3)_C$ group. They couple to the color charge, which is three-dimensional with axes *red*, *blue*, and *green*. The gluons are massless, but carry color charge themselves due to the non-Abelian nature of their gauge symmetry. This causes them to self-interact and gives rise to several characteristic properties of the strong interaction, such as:

- asymptotic freedom: the strong coupling decreases at shorter distances;
- confinement: only color-singlet bound states can be observed;
- chiral symmetry breaking: gluon-bound compound states have a large mass that is primarily generated by the interaction, not by the particles' rest masses.

The strong interaction not only binds quarks into hadrons; its residual effects also stabilize atomic nuclei via an attractive force between nucleons, i.e. protons and neutrons.

The *W* and *Z bosons* mediate the weak interaction, modeled by the spontaneously broken portion of the $SU(2)_L \times U(1)_Y$ symmetry. They are the only massive vector bosons in the Standard Model and their mass causes the weak interaction to be of extremely short range. Additionally, *W* and *Z* are the only bosons that break *parity symmetry* by coupling differently to left- and right-handed particles³. Due to its weakness and short range, the weak interaction is negligible in everyday life. It is most prominent as the cause of radioactive beta decay.

The *photon* is the gauge boson of the electromagnetic interaction, modeled by the group $U(1)_Q$ left unbroken by EWSB. It has no mass and does not self-interact, making it the only known gauge boson with infinite range. As such, it governs nearly all physical effects down to the atomic scale.

²Only upper bounds exist on the neutrino masses [20, pp. 628, 1006].

³*W* bosons couple exclusively to left-handed particles. *Z* bosons, while preferring left-handed partners, also couple to right-handed particles, if they carry electric charge.




Figure 2.2: Particle content of the Standard Model before EWSB [19]. The fermions are arranged by family, generation, and chirality, the gauge bosons by symmetry group. All particles are annotated with the gauge bosons or the charges that they couple to (bottom). The scalar bosons and fermions are additionally annotated with their color C (left), chirality (top left or right), and weak hypercharge Y (right), if they carry any.

2.1.3 The Higgs Boson and Electroweak Symmetry Breaking

An issue of the Standard Model – as described so far – is that it is inconsistent: careful analysis reveals that the non-zero masses of W and Z gauge bosons break the symmetry whose requirement introduced them to the theory in the first place. In a similar manner, it is impossible to naïvely introduce non-zero fermion masses due to the way that parity symmetry violation of the W and Z bosons is modeled. The GWS model [6–8] solves both issues by unifying electromagnetic and weak interaction into a single *electroweak interaction*, whose combined symmetry is spontaneously broken by the ground state of the universe.

Figure 2.2 shows the Standard Model before symmetry breaking. It is similar to fig. 2.1, but differs in several notable ways:

- All symmetries are upheld exactly. The $W^{1,2,3}$ and B gauge bosons of $SU(2)_L \times U(1)_Y$ are not the gauge bosons that are physically observed after symmetry breaking.
- All gauge bosons are massless and instead of the real scalar field H , there is a complex $SU(2)_L$ doublet Φ .
- The fermions are massless as well. As a consequence, their left- and right-handed components may be regarded as independent particles.

- Parity violation is implemented by arranging the left-handed fermions into $SU(2)_L$ doublets and the right-handed ones into singlets.
- Right-handed neutrinos do not exist. If they did, they would be sterile with respect to all three interactions. (But see section 2.3 on Standard Model extensions that introduce them.)

In this model, the self-interaction of the Higgs field Φ is set up in such a way that the field has degenerate ground states with a non-zero vacuum expectation value. By picking any particular ground state and expanding Φ into a Taylor series around it, the electroweak gauge symmetry is broken. According to the Brout–Englert–Higgs mechanism, this causes Φ , $W^{1,2,3}$, and B to mix into new fields:

- three of the four degrees of freedom of Φ become longitudinal modes of the gauge bosons;
- the remaining degree of freedom becomes a real, massive scalar field called the Higgs boson H ;
- the gauge bosons $W^{1,2,3}$ and B can be rotated into new bosons $W^\pm$, Z and γ . While the W and Z bosons become massive from interacting with Φ , γ remains massless. It corresponds to the $U(1)_Q$ subgroup that remains unbroken by the ground state of Φ .

In this model, the gauge symmetry of the W and Z bosons is still fulfilled at the fundamental level. It is the choice of a ground state of the vacuum that breaks it and allows these bosons to appear massive. In an analogous manner, fermion masses are added to the theory by introducing Yukawa couplings between the fermions and the Φ field. After symmetry breaking, each fermion gains both mass and a coupling to the remnant Higgs boson H . Characteristically, this coupling is proportional to the fermion mass.

2.2 Verification of the Standard Model

The Standard Model has been verified experimentally in multiple ways over the course of its existence. This section presents a select few cases of the Standard Model predicting particles that have been found later.

The charm quark was first speculated about in 1964 [15, 21, 22], less than a year after the quark model itself had been proposed and raised some controversy [23, 24]. It was meant to create a symmetry between the proposed three quarks and the four leptons known at the time. In 1970, the charm quark was further motivated by the GIM mechanism, which required it in order to explain an observed suppression of flavor-changing neutral currents [25]. J/ψ , the first charmed hadron, was observed in 1974 [26, 27] and convinced a majority of physicists of the quark model, which would become integral part of the Standard Model [28, p. 415–418].

The top and bottom quark were first proposed in 1973 [29], at a time when the quark model was still a contentious topic in physics. They were necessary to explain the CP violation in neutral kaon decays, which was found nine years prior [30]. The discovery of the τ lepton in

1975 [16] made the existence of a third generation of quarks seem even more likely. In 1977, the Υ as the first bottom-containing meson was discovered [31]. Its isospin partner, the top quark, was not found until 1995 [32, 33] due to its unexpectedly high mass of 176 GeV.

The W and Z bosons were first proposed in Glashow’s unified electroweak theory in 1961 [6]. This model still left the origin of the vector boson masses unexplained and did not see significant recognition initially [34]. The boson masses were explained independently by Weinberg and Salam in 1968 [7, 8] and in 1972, the GWS model was proven to be renormalizable [35, 36]. The model predicted “weak neutral currents”, i.e. Z -mediated interactions, which were found only one year later at CERN [37, 38]. In 1983, direct evidence for the existence of the W and Z bosons was found, also at CERN [39, 40].

The Higgs boson, first proposed in the seminal symmetry-breaking papers [10, 12, 13], has been a crucial prediction of the combined GWS model from 1967 on. However, it has been much harder to detect than the weak gauge bosons due to its higher (and initially unknown) mass, smaller interaction cross-section, and the ambiguous final states of its decays. First evidence of it was seen at the Tevatron starting in 2011 [41, 42], but the statistical significance was too low for a conclusive discovery. In 2012, a boson compatible with the predictions for the Standard Model Higgs boson was discovered at CERN [43, 44]. Since then, the properties of this boson have been further investigated and it has been confirmed to be the Higgs boson [45, 46].

2.3 Beyond the Standard Model

Although the Standard Model has been verified repeatedly and to great accuracy, it is known that it describes certain phenomena insufficiently and fails completely to describe others. This section lists some of these limitations and describes searches for Physics Beyond the Standard Model (BSM).

One of the most promising leads is a discrepancy between the predicted and measured value of the anomalous magnetic moment of the muon. The most accurate experimental value known today [47] deviates from calculations by about three standard deviations [20, ch. 57]. The “Muon $g-2$ ” experiment E989 [48] currently ongoing at Fermilab aims to reduce the uncertainty on this value [49].

Another lead is the so-called *strong CP problem*. CP invariance is a symmetry under which left-handed matter and right-handed antimatter behave identically, and vice versa. So far, the strong interaction has been found to obey this symmetry exactly. The problem is that there is no reason for this within the Standard Model. The proposed IAXO experiment [50] will search for axions, hypothetical particles whose existence might explain this surprising CP invariance.

A closely related (though somewhat inverse) problem of the Standard Model is the lack of an explanation for observed CP violation. Although the Standard Model breaks CP invariance via the CKM quark mixing matrix, this effect is too small to explain the matter–antimatter imbalance seen in the universe. An additional CP-violating term in the PMNS neutrino mixing matrix might fill this gap. There are several experiments both ongoing and proposed [51–54] that search for such a term.

Related to this is the problem of the still unknown neutrino masses: it is not yet understood why they are so much smaller than those of the charged leptons. In the Standard Model as it stands, they are assumed to be exactly massless. One explanation for this large difference is the seesaw mechanism: it introduces sterile, right-handed neutrinos with a Majorana mass term to the theory. Such a term is gauge-invariant and naturally leads to small masses for the left-handed and enormous masses for the right-handed neutrinos. It would also allow two more CP-violating parameters, thus solving the matter–antimatter imbalance problem described above. This extension of the Standard Model predicts a neutrino-less beta decay, which is searched for by several experiments such as GERDA [55] and SNO+ [56]. However, this decay has not been observed yet and only lower limits on its lifetime exist [57].

Furthermore, the Standard Model does not explain gravitation in any way. Attempts to quantize gravitation in the same manner as electro- or chromodynamics yield an unrenormalizable theory without predictive power. Various, more complicated theories such as loop quantum gravity and superstring theory exist, but have not been able to make any verifiable predictions yet.

Finally, there is the naturalness problem of the Higgs field Φ . In the perturbative formalism that is used to derive quantitative predictions from the Standard Model, the vacuum expectation value of Φ (like all other parameters) receives corrections from higher-order terms in the perturbation series. However, because H is a scalar field, these corrections are on the order of the cut-off scale Λ at which one expects the perturbative approach to break down.

This leaves two options: the first is that the Standard Model works as-is up to the energy scale of gravitation $M_P \approx 10^{19}$ GeV and the corrections need to be countered by an extremely fine tuning of the bare value of the vacuum expectation value, on the order of 10^{-34} . This is possible, but makes it surprising that the value is neither close to zero nor close to M_P . The second option is that Λ is far lower, on the order of 1 TeV. In that case, one should expect to see BSM physics at contemporary particle accelerators such as the LHC⁴.

⁴For some of the goals of the LHC physics program, see sections 3.1.4 and 4.1.

Chapter 3

LHC, ATLAS, and the Liquid-Argon Calorimeters

This chapter describes the experimental setup that this thesis is concerned with. Section 3.1 describes the LHC, section 3.2 the ATLAS experiment. The latter also introduces the ATLAS coordinate system and various notational conventions. Section 3.3 focuses further on the ATLAS detector's LAr calorimeters and their readout electronics.

3.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is currently the most powerful particle accelerator in the world [58]. It is situated at the Franco-Swiss border near Geneva and operated by CERN, the European Organisation for Nuclear Research. It is installed in the tunnel that was originally constructed for the Large Electron–Positron Collider (LEP). The tunnel has a circumference of 26.7 km and lies between 45 m and 170 m below the surface on a plane that slopes towards Lake Geneva by 1.4 %.

The LHC is a synchrotron that collides two counter-rotating beams of particles of equal charge. The beams normally consist of protons, though fully ionized lead ions may also be used in special *heavy-ion runs*. Due to the tight space requirements in the LEP tunnel, the rings of both beams are situated in the same mechanical structure.

3.1.1 History

The LHC project was approved by the CERN Council in 1994. Two years later, in 1996, construction of a 14 TeV machine in a single stage was approved. Commissioning started on September 10, 2008 [59] with the first proton beam travelling around the full circle of the LHC. Only nine days later, a magnet quench incident occurred that damaged multiple magnets and cryostats and compromised the beam pipe vacuum [60, 61]. It delayed operations by more than year; ultimately, for LHC Run 1, data was taken from March 2010 until December 2012 [62, 63]. Because of the accident, the initial center-of-mass energy had been reduced to 7 TeV. This was increased to 8 TeV in April 2012 [64].

After the Long Shutdown 1 (LS1), the LHC operated in Run 2 from June 2015 until October 2018 [65, 66] at a center-of-mass energy of 13 TeV. In this time, the LHC delivered an integrated luminosity of 160/fb [66, p. 7], of which ATLAS recorded 139/fb with sufficient quality [67, p. 17].

Currently, the LHC is in Long Shutdown 2 (LS2), a period in which the accelerator is upgraded in order to increase the provided luminosity even further. At the same time, the detectors perform their Phase-I upgrade to adapt to the higher luminosity. This is described in more detail in chapter 4. Run 3 is expected to start in early 2021.

3.1.2 Experiments

There are four primary experiments located at the LHC:

ATLAS [68] is a general-purpose particle detector built with the objective to exploit the full discovery potential of the LHC. It is designed to be sensitive to a large variety of particle signatures and perform accurate calorimetry on the full range of possible particle energies. While it can take data in all of the LHC's collision modes, it is optimized for proton-proton (pp) collisions. It is described in more detail in section 3.2.

CMS [69] is the other general-purpose detector at the LHC that focuses on pp runs. It shares many of the physics and performance goals with ATLAS, but achieves them through different design decisions. The two experiments complement each other and are able to independently verify each other's findings.

ALICE [70] is the only experiment at the LHC that is optimized for heavy-ions runs. While it is also a general-purpose detector, its physics program puts strong emphasis on quantum chromodynamics (QCD) and the physics of quark-gluon plasmas.

LHCb [71] is a detector specialized on the measurement of B mesons. Its objective is the indirect search for new physics through the CP violation and rare decays of bottom and charm quarks.

Furthermore, there are three smaller experiments: LHCf [72], which measures particles in the far-forward regions around the ATLAS interaction point; TOTEM [73], which measures the total pp interaction cross-section and the cross-section of diffractive scattering to improve understanding of low-energy QCD processes; and MoEDAL [74], which searches for magnetic monopoles and other exotic particles.

3.1.3 Injection Chain and Bunch Filling Scheme

Proton beams are injected into the LHC through a chain of pre-accelerators that increase the beam energy in multiple steps [76, ch. 1]. Figure 3.1 shows the full CERN accelerator complex and the LHC injector chain. Proton beams are produced in the Linac 2 facility at an energy of 50 MeV. They are transferred to the Proton Synchrotron Booster (PSB) and accelerated to 1.4 GeV before being injected into the Proton Synchrotron (PS). The PS further boosts them to an energy of 25 GeV and transfers them to the Super Proton Synchrotron (SPS), which further increases their energy to 450 GeV. From the SPS, the beams are injected into the LHC, which




Figure 3.1: The LHC (dark blue ring) is the last in a complex chain of particle accelerators [75]. The smaller machines boost the particles successively to the LHC injection energy of 450 GeV. They also provide beams to a whole set of smaller experiments.




Figure 3.2: The standard filling scheme presently in use at the LHC [77, fig. 1]. The pattern is structured into *batches*, where each batch consists of 72 filled nominal buckets followed by at least 8 empty ones.

accelerates them to their final energy. This was 6.5 TeV per beam in Run 2, adding up to a center-of-mass energy of 13 TeV.

The proton beams at the LHC are not continuous, but instead focused longitudinally into radio frequency buckets [76, sec. 2.5]. Each bucket is either empty or filled with a *bunch* of up to 1.15×10^{11} protons. While the buckets have a time distance of 2.495 ns, only every tenth bucket is nominally used. The bunches in these *nominal buckets* are ultimately brought to collision at the LHC intersection points, giving rise to a nominal bunch crossing (BC) distance of $\Delta t_{BC} = 24.95$ ns.

Not every nominal LHC bucket is actually filled with a bunch. The gaps in the beam follow a well-defined pattern and account for the rise time of various kicker magnets in the LHC injection and beam-dump system. Figure 3.2 shows the “25 ns scheme”, which currently is the standard bunch-filling scheme.

3.1.4 Energy Scale and Luminosity

The pp collisions at the LHC occur with a maximal center-of-mass energy of $\sqrt{s} = 13$ TeV. For a given process P , the expected event rate in the collisions is given by:

$$\dot{n}_P(t) = \sigma_P \cdot L(t), \quad (3.1)$$

where σ_P is the *cross-section* of P and $L(t)$ is the accelerator’s *instantaneous luminosity*. L is independent of P but changes over time. For example, it is typically maximal at the beginning

of a data-taking run and decreases exponentially as collisions remove protons from the beams¹. Hence, the total expected number of events of process P is:

$$n_P = \sigma_P \cdot \int_{t_1}^{t_2} L(t) dt, \quad (3.2)$$

where $\int L(t) dt$ is the *integrated luminosity*.

The original design of the LHC foresaw a peak luminosity of $L_{\max} = 10^{34}/(\text{cm}^2 \text{ s})$. There are two motivations for both this value and the value of $\sqrt{s}$. One is given by several studies searching for BSM physics: $\sqrt{s}$ is on the same scale as the masses of many hypothetical particles predicted by extensions of the Standard Model², and $L_{\max}$ is high enough to produce even rare events at a rate that allows meaningful analysis [78, sec. 1.1].

The other, equally important motivation is a better understanding of the EWSB and the measurement of the Higgs boson mass m_H . The LHC has been designed to allow a Higgs boson search over the full mass range that is allowed by theory, but particular consideration has been given to the range below 300 GeV. This range was strongly favored by higher-order contributions to precisely measured Standard Model parameters and the then-recent measurement of the top quark mass m_t at the Tevatron [79].

3.1.5 Pileup

The instantaneous luminosity $L(t)$ is of great importance to all LHC experiments and must be known accurately. ATLAS, for example, measures it in several ways for cross-validation, both with its regular subsystems and with dedicated detectors [80, sec. 3].

A related quantity is $\langle \mu \rangle$, the average number of pp interactions per bunch crossing:

$$\langle \mu \rangle = \frac{\sigma_{\text{inel}} \langle L(t) \rangle_t}{f_r}, \quad (3.3)$$

where σ_{inel} is the pp inelastic interaction cross-section, $\langle L(t) \rangle_t$ is the average luminosity of a data-taking run, and f_r is the LHC revolution frequency. When taking a limited detector efficiency ε into account, $\langle \mu \rangle$ and σ_{inel} have to be replaced with their *visible* counterparts, $\langle \mu \rangle_{\text{vis}} := \varepsilon \langle \mu \rangle$ and $\sigma_{\text{inel,vis}} := \varepsilon \sigma_{\text{inel}}$. These two substitutions cancel out and do not change the equation qualitatively [80, sec. 2].

The quantity $\langle \mu \rangle$ is vital to filter-design studies, as it is a measure of *pileup*, an effect where multiple pp interactions occur in the same or subsequent BCs. *In-time pileup*, when multiple interactions occurring in the same BC, presents an irreducible background that can only be removed on average or by regional analysis. *Out-of-time pileup*, in which detector signals from pp interactions in subsequent BCs interfere with each other, can and should be suppressed as much as possible. For more information, see also section 3.3.5.

¹ This simple behavior is complicated by an technique called *luminosity leveling*, which is expected to be employed after the upgrade described in section 4.3. This technique begins an LHC run with suboptimal collider parameters that keep the luminosity below the theoretical maximum. Over the course of the run, as protons are removed from the beam, the parameters are tuned continuously so that the luminosity remains constant.

² Examples of hypothetical particles that could be discovered are the lightest stable supersymmetric particle (LSP) of supersymmetry, the MSSM Higgs bosons (A , $H^\pm$, and H_0), and heavy gauge bosons W' and Z' of an undiscovered gauge symmetry.




Figure 3.3: Cut-away view of the ATLAS detector [81]. The detector is 25 m high and 44 m long and it weighs approximately 7000 t [68].

3.2 The ATLAS Detector

ATLAS is one of the two general-purpose detectors installed at the LHC. It is situated in Underground Experimental Hall 15 (UX15) near CERN’s Meyrin site and surrounds Interaction Point 1 in a cylindrical, forward-backward symmetrical fashion. Figure 3.3 shows a 3D model of it.

Section 3.2.1 introduces the ATLAS coordinate system and nomenclature. Sections 3.2.2 to 3.2.5 describe the various subsystems of the detector. The liquid-argon calorimeters, as the focus of this thesis, are described separately in section 3.3.

3.2.1 ATLAS Coordinate System and Nomenclature

The nominal interaction point of the colliding beams is defined to be the origin of the ATLAS coordinate system. The x -axis points towards the center of the LHC ring and the y -axis points upwards. The z -axis points along the beam in such a way as to make a right-handed coordinate system³. Spherical coordinates are often preferred due to the detector’s symmetric design. For them, the azimuthal angle ϕ is taken from the positive x -axis towards the positive y -axis and the polar angle θ is taken from the positive z -axis.

A typical issue at hadron colliders is that at high energies, it is not the composite hadrons that interact, but their individual partons, i.e. gluons and quarks [82]. Each parton carries only a fraction of the entire hadron’s momentum and the fraction fluctuates considerably between

³The terms “side A” and “side C” are sometimes used to refer to the positive and negative half of the z -axis respectively.

collisions. Because of this, the collision's center of mass has a momentum that is not zero in the lab frame, but instead randomly distributed. More precisely, all measured momenta vary considerably in their component along the beam axis; the variance of their transverse component is small in comparison.

For this reason, it is often preferable to use quantities that are invariant under longitudinal Lorentz boosts [20, ch. 47]. For example, the full momentum p of a particle is usually eschewed in favor of its *transverse momentum*:

$$p_T := p \cdot \sin \theta, \quad (3.4)$$

and its longitudinal rapidity⁴:

$$y_L := \operatorname{artanh} \frac{p_L}{E}. \quad (3.5)$$

In the ultrarelativistic limit, p_T is approximated by the transverse energy E_T and y_L by the pseudorapidity η , where:

$$E_T := E \cdot \sin \theta, \quad \eta := \operatorname{artanh} \cos \theta = -\ln \tan \frac{\theta}{2}. \quad (3.6)$$

These quantities are easier to measure in the detector and the approximation is generally permissible due to the high center-of-mass energy of the pp collisions. In particular, η almost universally substitutes θ in the cylindrical coordinate system.

3.2.2 Inner Detector

The Inner Detector is the ATLAS subsystem closest to the interaction point. It is shown in fig. 3.4. It covers the pseudorapidity region $|\eta| < 2.5$ and serves to reconstruct tracks for the final-state particles of each pp collision. It also reconstructs each event's primary and secondary vertices, which is particularly important for b - and τ -tagging of hadron jets.

Three distinct mechanical units make up the Inner Detector: a *barrel* and two *end-caps*. The barrel covers the central region and surrounds the beam pipe in multiple cylindrical layers. The end-caps cover the forward and backward regions and consist of multiple concentric wheels perpendicular to the beam axis. This separation is employed in all ATLAS subsystems.

A solenoid magnet surrounds the Inner Detector and immerses it in a magnetic field. This field curves the tracks of charged particles to allow measuring their charge and momentum. In the barrel, it is homogeneous with a strength of 2 T; in the end-caps, it is weaker and bends outwards considerably.

Both the barrel and the end-caps consist of three sub-detectors, each one contributing location measurements in a distinct radial range from the beam pipe. Tracks are reconstructed from these measurements via a chain of sophisticated pattern-recognition algorithms.

⁴The longitudinal rapidity y_L differs from the regular rapidity y in that it is based on p_L instead of the full momentum p . The overwhelming majority of texts on hadron collider physics drop the distinction and call y_L simply *the* rapidity. This is approximately correct if particles are boosted, i.e. $p_T \ll p_L$.




Figure 3.4: Cut-away view of the ATLAS inner detector [83].

The Pixel Detector is the inner-most and (with approx. 80.4 million readout channels) most granular sub-detector. It consists of silicon pixel sensors and provides a resolution of $10\text{ }\mu\text{m}$ in the $R\phi$ -plane and $110\text{ }\mu\text{m}$ on the z -axis (barrel) or R -axis (end-caps). A single type of sensor module is used throughout the Pixel Detector, making it extremely homogeneous. The modules are arranged in such a manner that each particle crosses at least three layers of them.

In addition to the above, the Pixel Detector has been extended with the Insertable B-Layer (IBL) during LS1 [84, 85]. The IBL is a fourth pixel layer that is located within the inner-most pixel layer. To make space for it, the beam pipe has been replaced with a newer, smaller one.

The addition of the IBL has been motivated by three reasons: 1. it pre-empts future irreparable failures in the inner-most pixel layer; 2. it provides additional redundancy in track reconstruction, which is required when increasing the LHC's instantaneous luminosity beyond its design value; 3. it improves the performance of vertexing and b -jet tagging thanks to its location close to the interaction point.

The Semiconductor Tracker (SCT) surrounds the Pixel Detector and provides on average four measurements per track. It has approx. 6.3 million readout channels and employs silicon micro-strip sensors with a length between 7 cm and 12 cm. Each sensor has 698 active strips with a pitch of $80\text{ }\mu\text{m}$. This results in a resolution of about $17\text{ }\mu\text{m}$ in the $R\phi$ -plane.

The remaining coordinate (z in the barrel, R in the end-caps) is measured via the *small-angle-stereo* method: Each module actually consists of *two* layers of micro-strip sensors, which are glued back-to-back onto a carrier plate. While one layer is aligned to the beam pipe (parallel in the barrel, radial in the end-caps), the other is rotated by 40 mrad . The resulting moiré pattern of crossing strips allows measuring z and R respectively with a resolution of $580\text{ }\mu\text{m}$.




Figure 3.5: Cut-away view of the ATLAS calorimeter system [87].

The Transition Radiation Tracker (TRT) provides the majority of location measurements, typically 36 per track. It consists of many *straws*, where each one is a proportional drift tube detector. The cathode of each straw is kept at a potential of about -1530 V , the anode is kept at ground potential and connects directly to the readout electronics. The straws are interleaved with polypropylene, which causes crossing ultrarelativistic particles to emit transition radiation. The straws' gas filling is based on xenon⁵ and absorbs the transition radiation well, converting it to signals of much higher amplitude than that of ionizing particles. This allows discriminating between the two signals in the front-end electronics.

The TRT can only measure the $R\phi$ location of a hit. With a resolution of $130\text{ }\mu\text{m}$ and about 351 000 readout channels, it is also coarser than the previous two sub-detectors. However, this is compensated by the high number of hits per track.

3.2.3 Calorimeter System

The ATLAS calorimeters, shown in fig. 3.5, are designed to stop the particles that come out of a pp collision and to measure their energy as accurately as possible. Additionally, precise location information is highly desirable, too.

The calorimeters are categorized as either *electromagnetic* or *hadronic*. The former absorb photons, electrons, and positrons well, while the latter are optimized to absorb pions, kaons,

⁵In 2012, several inaccessible leaks have formed in the gas pipes between the detector and the cleaning and mixing stations. For Run 2, the xenon-based gas mixture has been replaced with an argon-based one in the modules most affected by these leaks. The reason for this is that argon, while absorbing transition radiation much less efficiently than xenon, is significantly cheaper [86].

and other hadrons. Neutral particles, such as neutrons and K_L^0 , also generally require the additional depth of the hadronic calorimeters to be stopped completely. Muons, while depositing energy in both calorimeters, cannot be stopped if their energy is greater than a few GeV. Neutrinos pass through both calorimeters undetected; their presence can only be inferred by the non-conservation of visible p_T .

All calorimeters used in ATLAS are *sampling calorimeters*: they consist of alternating layers of an absorber material and an active medium. The absorber decelerates the incoming particles and causes them to produce *showers*, but does not allow any measurements. The showers, in turn, excite the active medium and produce a signal that is proportional to the *visible energy*, i.e. the energy that is deposited in the active medium. If the *sampling fraction* of the detector, the ratio of a particle's visible to its total deposited energy, is known, the total deposited energy can be inferred from the signal.

The ATLAS detector has four calorimeters: three based on liquid argon (LAr) and one based on scintillator tiles. The LAr calorimeters are discussed in more detail in section 3.3.

The Electromagnetic Calorimeter covers the pseudorapidity range of $|\eta| < 3.2$ and consists of an Electromagnetic Barrel (EMB) and two Electromagnetic End-Caps (EMEC), each of the three having its own cryostat. The active medium is LAr. The absorber is lead, with additional stainless steel sheets for mechanical strength. For the readout, copper electrodes insulated with polyimide are inserted into the gaps between the absorber sheets.

Both the EMB and the EMECs are equipped with a *presampler*, a thin LAr gap without absorbers in front of the calorimeter proper. Its purpose is to estimate the energy that has been lost between the interaction point and the calorimeters. A slight overlap between the barrel and end-cap presamplers improves the energy resolution in a region where the amount of dead material is highly heterogeneous. For $|\eta| > 1.8$, the amount of material between the calorimeters and the interaction point is small enough that no presampler is necessary.

The Tile Calorimeter is a hadronic calorimeter that covers the pseudorapidity range $|\eta| < 1.7$. It consists of a *central barrel* and two *extended barrels* and surrounds the other calorimeters in a cylindrical manner. Steel is used as the absorber material and plastic scintillator⁶ tiles as the active medium. The three barrels are segmented azimuthally into 64 independent modules, each supported by a strong-back steel girder, which also houses the readout electronics. Each scintillator tile is read out on both ends via wavelength-shifting fibers that are coupled to photomultiplier tubes in the girder.

The Tile Calorimeter also contains several special modules that instrument the gap between the barrel and the end-cap calorimeters. Their measurements help correct for energy losses in the gap's dead material, e.g. the cryostat walls, the services for the inner detector, and the services and power supply for the other calorimeters. These special modules are either reduced sections of otherwise normal Tile calorimeter modules (then called *plugs*) or thin scintillator detectors attached to the extended barrel or end-cap cryostat (then called *gap* and *cryostat scintillators* respectively).

⁶The scintillator material is polystyrene doped with 1.5 % PTP and 0.044 % POPOP.




Figure 3.6: Cut-away view of the ATLAS Muon Spectrometer [88].

The Hadronic End-Caps (HEC) are two hadronic calorimeters, one being located on each side of the ATLAS detector. Each one is located within its respective end-cap cryostat, directly behind the EMEC. They cover the pseudorapidity region of $1.5 < |\eta| < 3.2$. This overlaps slightly with the η coverage of the Tile and Forward Calorimeters in order to soften the impact of the drop in material density in these two transition regions. Like the Electromagnetic Calorimeter, the HECs are LAr sampling calorimeters, but they use copper as absorber material.

The Forward Calorimeters (FCal) are the calorimeters closest to the beam pipe and exposed the most to radiation damage. They are embedded inside the HECs, using the same cryostats, and cover the pseudorapidity region $3.1 < |\eta| < 4.9$. Each FCal consists of three longitudinally arranged, cylindrical modules. All modules are LAr calorimeters, though their absorber material differs. The first one, FCal1, uses copper and is optimized for electromagnetic interactions. The other two, FCal2 and FCal3, use tungsten in order to efficiently absorb and contain hadronic showers.

3.2.4 Muon Spectrometer

The Muon Spectrometer is the outermost and largest subsystem of ATLAS. It is shown in fig. 3.6. Its purpose is to provide precise tracking of muons and reliable triggering on their detection. Similar to the Inner Detector, this is achieved by bending the muon trajectories via a magnetic field and measuring their location and momentum. The field is generated by three large superconducting air-core toroid magnets: one barrel magnet and two end-cap magnets

that are inserted into each end of the barrel magnet. This magnet design minimizes undesirable multiple-scattering effects, though its disadvantages are that the field is comparatively weak (0.5 T on average) and inhomogeneous [78, sec. 6.3].

The Muon Spectrometer uses four different kinds of detectors. Two of them are used for tracking, i.e. reconstruction of muon trajectories and momenta. This requires the η coordinate of the muon tracks to be measured with high precision, as this is the primary bending direction of the magnets. The other two serve multiple purposes: they provide the ATLAS trigger with fast and coarse muon track information, enable BC identification, and measure the complementary ϕ coordinate for the former two detectors.

Monitored Drift Tube (MDT) chambers are the primary tracking detectors of the Muon Spectrometer. Each chamber consists of drift tubes arranged in multiple layers and held in place by spacer bars. This design was chosen due to its high reliability in case of individual tube failure and to improve drift-time resolution. The tubes are operated with a voltage of 3080 V and an Ar/CO₂ gas mixture at 3 bar. The chambers are arranged in three layers around the interaction point, both in the barrel and the end-caps. They cover the full pseudorapidity range of $|\eta| < 2.7$. The only exception to this is the region $|\eta| > 2.0$ in the inner-most layer, where radiation is too high for their operation [89, sec. 6.1].

Cathode Strip Chambers (CSC) replace the MDT chambers in this region. A CSC is a multi-wire proportional chamber with parallel anode wires and cathode plates. These plates are lithographically segmented into strips. In contrast to conventional proportional chambers, the signal is not read out via the anode wires; instead, the track position is inferred by interpolation of the charges induced onto neighboring cathode strips.

Resistive-Plate Chambers (RPC) provide the muon trigger information in the barrel region $|\eta| < 1.05$. Each chamber contains two detectors, each detector in turn consisting of two electrode plates of constant distance with a high voltage and an operating gas between them. The high voltage causes any electrons from the ionizing muons' tracks to form avalanches on their way to the anode. Because these avalanches are formed nearly instantaneously⁷, the signal arrives in only a few nanoseconds. Pick-up strips located behind the electrode plates pick up the avalanches inductively and lead the signal to the readout electronics. Each detector's pick-up strips are orthogonal to each other, allowing to measure both η and ϕ .

Thin-Gap Chambers (TGC) finally provide the trigger information in the end-cap region $1.05 < |\eta| < 2.4$. They are multi-wire proportional chambers where the wire-cathode distance (1.4 mm) is smaller than the pitch between wires (1.8 mm). The operating gas, a mixture of CO₂ and *n*-pentane, is highly quenching and allows the chambers to be operated close to the Geiger region. In this mode, the signal amplitude not only arrives fast and is almost independent of the detected particle's energy and incident angle, but it is also insensitive to any mechanical deformations of the chamber. The ϕ coordinate is read out via the cathodes, which are segmented into radial strips, the η coordinate is read out via the anode wires.

⁷In wire-based detectors, avalanching occurs only in the region close to the anode wire. Outside of this region, electrons drift comparatively slowly.

3.2.5 Trigger and Data Acquisition System

At a frequency of 40.08 MHz, bunch crossings in the LHC happen more quickly than they can be written to permanent storage. Hence, a trigger algorithm must be applied that discards all events deemed “uninteresting”. Events selected by the trigger are transferred off the detector and processed further by the data acquisition system (DAQ) before ultimately being put into storage. Together, trigger and DAQ are referred to as TDAQ. In addition, the DAQ distributes the LHC clock signal to the various subsystems via the trigger, timing, and control system (TTC). Finally, it allows monitoring and configuring the detector’s run-time parameters via the Detector Control System (DCS).

Since its original design, the TDAQ has evolved significantly. To handle the expected increase in instantaneous luminosity in Run 2, its architecture has been streamlined and its hardware upgraded [90]. For the sake of relevance, the trigger as of Run 2 (2015–2018) is described in this section. Other sources describe the original design employed in Run 1 [68, 91].

The ATLAS trigger logic is split into two levels, each one further restricting the number of selected events. The higher level also has access to more information and more time to come to a trigger decision.

The L1 trigger is the first trigger stage in ATLAS. It reduces the event rate to 100 kHz (75 kHz in Run 1) and its latency is less than $2.5\ \mu\text{s}$. To accomplish this, it is implemented in custom hardware and located close to the detector in USA15⁸.

The L1 trigger is split into a calorimeter trigger and a muon trigger. The calorimeter trigger receives data at a reduced granularity and searches for clusters of energy depositions that are typical for high- p_T electrons, photons, jets or τ leptons. It also triggers on events with a high missing E_T and where the sum of E_T across all jets is above a given threshold. The muon trigger has access to RPC and TGC data and searches for high- p_T muons.

Both triggers forward their results (particle multiplicities and energy sums) to the Central Trigger Processor (CTP). The CTP compares them with a list of customizable *trigger conditions*. These conditions include thresholds on certain quantities, but also clocks and random triggers. The list of satisfied trigger conditions is then compared to a list of *trigger items*, where each condition may contribute to each item. Each trigger item is associated with a *pre-scale factor*, which further reduces the frequency with which they are accepted. Ultimately, the L1 accept signal is a list of all matched trigger items. If an event matches none, it is discarded.

Parallel to this selection step, the L1 trigger also marks a number of locations in $\eta - \phi$ space as regions of interest (RoI). These regions seed the algorithms of the next trigger stage, highlighting parts of the detector data that should be read out and investigated first.

The High-Level Trigger (HLT) is the second stage of the ATLAS trigger process. It combines what used to be two separate stages in Run 1: the L2 trigger and the Event Filter. In contrast to the L1 trigger, it is software-based and runs above ground on a cluster of commercially available computers.

⁸Despite the proximity, the finite signal propagation speed in the cables between detector and trigger alone introduces a latency of about $1\ \mu\text{s}$. It is not feasible to put the trigger any closer due to a lack of space and the cost of making its electronics radiation-hard enough.

For each incoming event, the HLT executes a sequence of algorithms. It starts with fast triggers for early rejection and continues with slower algorithms similar to those for offline reconstruction. A typical algorithm consists of one or more feature extraction steps – which request data fragments from the RoIs – and a hypothesis step, which check whether the data meet certain trigger conditions. For some algorithms (e.g. those that require the missing transverse momentum), data from the entire detector is requested. Overall, each algorithm is designed to terminate within a few hundred milliseconds.

Each accepted event is tagged based on the chain of L1 trigger item and HLT algorithms that it passed. Depending on this trigger chain, another pre-scale factor is applied and may discard the event⁹. Only after this stage is the event written to one or more files, the so-called *streams*. The streams, too, are selected based on the event's trigger chain. The most important ones are the *physics streams*¹⁰, which contain the full detector information for the event and commit it to permanent storage.

3.3 The ATLAS Liquid-Argon Calorimeters

At the ATLAS experiment, liquid-argon calorimeters are used for hadronic calorimetry in the end-caps and for electromagnetic calorimetry throughout the detector. They have been chosen due to their intrinsic radiation hardness and because their response to ionization signals is both highly linear and stable over time [68, ch. 5].

Liquid argon as a detection medium was first investigated in the early 1950s, when it was found that it didn't suffer from space charge accumulation effects like the contemporary crystal counters did [92]. Starting in 1968, when it became apparent that multi-wire gas chambers could not deliver the spatial resolution required by future experiments, the high-energy physics community became interested in noble liquids [93, 94]. Six years later, in 1974, three independently developed liquid-argon calorimeters were presented [95–97]. Soon after, liquid-argon calorimeters were used in a multitude of particle physics experiments, e.g. R807-808 at the ISR, DØ at Tevatron, and H1 at HERA [98].

This section first details the working principle of LAr calorimeters in general. Afterwards, the geometry and the readout and trigger electronics of the ATLAS LAr calorimeters are described. Finally, the most important sources of noise in the LAr calorimeters are enumerated.

3.3.1 Working Principle

LAr calorimeters work in an analogous manner to gaseous ionization detectors: an active detector volume of liquid argon is flanked by two electrodes and a high voltage is maintained between them. A particle passing through the detector will deposit some of its energy by ionizing the argon atoms along its track. The ions and free electrons thus produced are separated by the electric field and drift to the cathode and anode respectively. Because of the high density of liquid argon, particles have an extremely short mean free path in it and drift at a constant speed (approx. 225 ns/mm for electrons).

⁹A pre-scale factor of n means that the event is kept with a probability of $1/n$.

¹⁰The ATLAS physics streams are: electrons, muons, jets, photons, E_T^{miss} and τ leptons, and B physics.




Figure 3.7: Sketch of an EMB module with the detector cells in all four layers [68, p. 114]. Segmentation in ϕ is achieved by ganging multiple electrodes, in η and depth by etching the electrodes.

While the argon ions can usually be neglected due to their low drift speed, the electrons induce a current pulse on the anode while drifting. The pulse is approximately triangular in shape, with a near-instantaneous rise time and a slow, linear drop back to zero. This shape is derived in appendix A.

3.3.2 Detector Segmentation

The LAr calorimeter system consists of four subsystems EMB, EMEC, HEC, and FCal. Each one covers a different pseudorapidity region, with some overlap between them to avoid gaps and to lessen transition effects. In order to locate calorimeter hits, the detectors are segmented in three dimensions: radially, azimuthally, and along the beam axis. On two axes, the segmentation is achieved by etching the detector electrodes; on the third axis, multiple electrodes are *ganged*, i.e. their outputs are summed via analog electronics. The units of this segmentation are called *calorimeter cells*.

The detectors are segmented in a *pseudo-projective* manner towards the interaction point. This means their cells form a rectangular grid in the $\eta - \phi$ plane and are arranged in up to four layers perpendicular to this plane. Figure 3.7 visualizes this grid geometry for the EMB,

table 3.1 lists cell sizes for all detectors except of the FCal.

For various reasons, this geometry is only approximate and there are several deviations from it, in particular in the far-forward region. Some of these deviations are:

- The EMB consists of two half-barrels with a 6 mm gap between them at $\eta = 0$. This *crack* region is not instrumented and energy depositions in it are not read out [78, p. 121]. In the front layer, the gap amounts to two cell widths, one per half-barrel, and the corresponding readout channels are masked. In the other layers, the gap slightly reduces the energy resolution of the cells next to it.
- Due to its weight (approx. 110 t), the EMB sags slightly, shifting detector cells from their ideal positions in $\eta - \phi$ space [68, sec. 5.2.4]. This effect amounts to a shift on the η axis by up to one cell width in the front layer, as fig. 3.8 shows.
- Due to its mechanical structure, the EMEC has a 3 mm gap between its two wheels. However, due to its coarser segmentation, this does not impact the readout.
- Also for mechanical reasons, the back layer covers only the outer wheel of the EMECs (i.e. the region $1.5 < |\eta| < 2.5$), not the inner one.

Table 3.1: Segmentation of the LAr calorimeters into cells [68, p. 9]. The FCal does not follow the grid-like geometry of the other calorimeters and thus is not listed here.

Calorimeter	Layer	Region ($ \eta $)	Cell size ($\Delta\eta \times \Delta\phi$)
EMB	Presampler	0.0 ... 1.5	0.025 $\times$ 0.1
		1.5 ... 1.52	0.02 $\times$ 0.1
	Front	0.0 ... 1.4	0.025/8 $\times$ 0.1
		1.4 ... 1.475	0.025 $\times$ 0.025
	Middle	0.0 ... 1.4	0.025 $\times$ 0.1
		1.4 ... 1.475	0.075 $\times$ 0.025
	Back	0.0 ... 1.35	0.05 $\times$ 0.025
EMEC	Presampler	1.5 ... 1.8	0.025 $\times$ 0.1
	Front	1.375 ... 1.425	0.05 $\times$ 0.1
		1.425 ... 1.5	0.025 $\times$ 0.1
		1.5 ... 1.8	0.025/8 $\times$ 0.1
		1.8 ... 2.0	0.025/6 $\times$ 0.1
		2.0 ... 2.4	0.025/4 $\times$ 0.1
		2.4 ... 2.5	0.025 $\times$ 0.1
		2.5 ... 3.2	0.1 $\times$ 0.1
	Middle	1.375 ... 1.425	0.05 $\times$ 0.025
		1.425 ... 2.5	0.025 $\times$ 0.025
		2.5 ... 3.2	0.1 $\times$ 0.1
	Back	1.5 ... 2.5	0.05 $\times$ 0.025
HEC	all layers	1.5 ... 2.5	0.1 $\times$ 0.1
		2.5 ... 3.2	0.2 $\times$ 0.2




Figure 3.8: Cell misalignment due to sagging in the EMB front layer for $|\eta| < 1.4$. Each cell is given an index on the η -axis (horizontal) and on the ϕ -axis (vertical). The ϕ axis is oriented such that the bottom half of the diagram corresponds to the bottom half of the detector. Color indicates each cell's shift towards positive η in units of the cell width. The two-cells-wide vertical strip in the middle corresponds to the crack region between the two half-barrels.

- The FCal adheres to the grid geometry much less than the other detectors. One reason for this is that in the far-forward region, constant-size bins in η are physically smaller than in the central region. Another reason is the hexagonal arrangement of the FCal electrodes. Figure 3.9 shows an example of what typical FCal cells look like.

3.3.3 Accordion Geometry

The LAr calorimeters operate in ionization-chamber mode, i.e. no avalanche effect occurs at the anode. Hence, the signal pulse is rather small (around $2.5 \mu\text{A}/\text{GeV}$) and needs to be amplified. In order not to distort the pulse's sharp rising flank, the amplifier must have a short reaction time and the signal run time to the readout electronics must be short [98, p. 1259].

A short signal run time corresponds to a short cable length. In the EMB, this is achieved by arranging the absorbers and electrodes radially to the beam axis and reading them out on the inside and outside face of the barrel. They are bent into an *accordion shape* so that the azimuthal direction is covered homogeneously [100]. The accordion waves run radially and their folding angle decreases with increasing radius. As fig. 3.10 shows, this keeps the LAr gap and, by extension, the calorimeter response in $\mu\text{A}/\text{GeV}$ constant.

In the EMECs, a similar accordion geometry is employed, but the accordion waves run parallel to the z -axis to match the incident angle of particles. However, it is not possible to keep the calorimeter response radially constant this way. Rather, a combination of parameters is varied to minimize the change in response: the sampling fraction naturally decreases with




Figure 3.9: Example of an FCal trigger tower [99, p. 27], bounded here by the thick, straight outline. Its constituent detector cells are outlined by thinner, honeycomb-like zigzag lines. Each readout cell is made up of multiple electrode-equipped LAr gaps, here represented as circles.




Figure 3.10: The calorimeter response of the Electromagnetic Calorimeters [68, p. 148]. The jump at $|\eta| = 0.8$ is due to a change in absorber thickness, the jump at $|\eta| = 2.5$ due to the transition from the outer to the inner wheel of the EMEC.

increasing radius; the accordion folding angle decreases and the amplitude increases with increasing radius; and the high voltage across the LAr gap is tuned specifically to this end. The result of this effort is shown in fig. 3.10.

The accordion waves in the EMECs cannot be varied arbitrarily: the fabrication process of the electrodes puts a limit on the ratio of outer to inner radius of the EMECs. To circumvent this limit, the EMECs are split into two coaxial wheels, with a thin gap between them at $|\eta| = 2.5$.

3.3.4 Readout Electronics

All ATLAS LAr calorimeters share a common readout system. This system provides calorimeter information both to the L1 calorimeter trigger (at reduced resolution) and to the DAQ system (at full resolution). It can read out the calorimeter data at the full L1 trigger rate of up to 100 kHz and handles a dynamic range of energies from 10 MeV up to 3 TeV.

In order to reduce cost while also meeting its latency requirements, the readout is split into an off-detector *back end* and an on-detector *front end*. While the front end uses custom ASICs to tolerate the high radiation levels inside the detector, the back end is built from commercially available electronic components. The front end is further separated into *cold electronics*, which are installed within the cryostats, and *warm electronics*, which are situated in front-end crates (FEC) just outside of it. Figure 3.11 shows the general architecture common to all LAr calorimeters.

The cold front-end electronics comprise summing boards and motherboards. The former group neighboring electrodes into *calorimeter cells*, which is called *ganging*. The latter route the summing board signals to readout cables while minimizing cross-talk between cells. They also contain the point at which *calibration pulses* are injected into the readout chain.

The readout cables lead to the warm electronics via vacuum-sealed feedthroughs on the cryostat. On the entire way, the components' impedances have been matched to minimize reflections. Only the FCal electrodes, due to their unique shape, have a much smaller impedance than the other components. This causes significant signal reflections in them and gives the FCal signal a distinct pulse shape.

The HECs have a mechanical structure [101, sec. 10.1.3.2] that differs greatly from that of the electromagnetic calorimeters. This allows the preamplifiers to be mounded on the (cold) motherboards, greatly reducing the electronic noise. The other calorimeters, by contrast, amplify their signals in the warm front end.

The warm front-end electronics consist of multiple boards, the most crucial of them being the Front-End Board (FEB). It accepts the signals coming from the cold electronics and amplifies them via a transimpedance preamplifier [102, sec. 4.1] before splitting the electronics chain into a *L1 trigger path* and a *main readout path* (see fig. 3.12). In the HECs, where the preamplifier part of the cold electronics, the FEBs have a *preshaper*, which inverts, shapes, and adjusts the gain of its input to make the signal pulse shape more similar to that of the EMB and EMEC.

On the **main readout path**, the preamplifier output is split into three overlapping gain levels: low ($\times 1$), medium ($\times 10$), and high ($\times 100$). On each level, an *analog shaper* applies a CR-RC² band-pass filter, i.e. a high-pass filter followed by two low-pass filters. The former removes




Figure 3.11: Block diagram of the LAr readout electronics [68, p. 139]. From top to bottom, these are the back-end electronics, the warm front-end electronics, and the cold front-end electronics. Differently shaded boxes signify different chips.




Figure 3.12: Block diagram of the LAr Front-End Board [68, p. 141]. Detector signals arrive in the top-left. The trigger path leaves the board in the bottom left, the main readout path on the right.

the long tail of the triangular signal pulse, the latter reduce the amplified thermal noise of the cold electronics. The time constant $\tau := RC = 13$ ns of the filters has been chosen to minimize the total noise (see section 3.3.5 for details).

After shaping, the signals are sampled at the BC frequency of 40.08 MHz and buffered in a switched-capacitor array (SCA). If the L1 trigger accept signal arrives within the expected time window, a gain-selector chip chooses the optimal gain level and reads a configurable number of samples (typically four or five) from the SCA. The samples are passed through a 12-bit analog-to-digital converter (ADC), formatted, and transmitted optically to the back end.

On the **L1 trigger path**, the signal is only digitized in the back end. The front end reduces the detector resolution and makes the signal partially proportional to E_T (as opposed to E). The former is done by summing the signals of neighboring cells in multiple stages, the latter by tuning the gain of each stage depending on the location of the cells. This conversion happens only partially in the front end; the back end contains the final gain adjustment.

The FEB's *Linear Mixer* chip first sums the signals of four neighboring cells. It also shapes the sum in a similar manner to the main readout, but lacking the second low-pass filter. This is applied in the back end to suppress the additional thermal noise caused by the longer analog section.

The shaped signals are routed to the Layer Sum Board (LSB), which sums them to their final resolution on the η and ϕ axes. A Tower-Builder Board (TBB) receives these layer sums, combines them into *trigger towers* and sends these to the back end. In the HEC and FCal, due to their lower granularity, the layer sums are already produced on the FEB. Hence, the TBB is replaced by a Tower-Driver Board (TDB), which only transmits its input to the back end.

In addition to the boards described above, the FECs also contain: calibration boards with the necessary logic to inject calibration pulses into the cold electronics; controller boards which distribute the TTC and other signals to configure the front end; and baseplanes (two per FEC) which route the signals between the other boards.

The back-end electronics are located in the USA15 cavern, approximately 70 m away from the detector. It consists of three systems: The TTC system, which receives the global LHC clock signal and the L1 accept signal from the network and distributes it to the other systems; the L1 receiver system; and the Readout Drivers (ROD).

The L1 receiver system receives signals from the front end's trigger path and forwards them

to the L1 calorimeter trigger. It also adjusts these signals with the missing low-pass filter and the final, η -dependent gain to make it proportional to E_T .

On the L1 trigger side, a pre-processor module receives these signals and digitizes them. The sampling rate was 40.08 MHz in Run 1 and 80.16 MHz in Run 2 [91, sec. 3.3.1]. The resulting samples are processed by an Optimal Filter (OF, see section 5.6 for a description). Maxima in the OF output are used to identify BCs with energy deposits and to estimate the deposited E_T based on a look-up table.

If the deposited energy saturates the analog electronics, the OF is ignored. Instead, a rising-edge-detecting algorithm identifies the offending BC and a sentinel value indicating an out-of-range energy is produced. The algorithm has been adapted for Run 2, but still generally follows its original design [103, sec. 5.5].

The BC information and E_T are sent to downstream modules, which calculate various quantities and define RoIs. This information is eventually sent to the CTP, which ultimately decides whether to send the L1 accept signal.

The RODs receive the digitized signal samples from the front end's main readout path. They check their integrity, process them, and send the results to the DAQ system. As in the trigger, processing is based on an OF (see section 5.6 for a description).

3.3.5 Sources of Noise

The uncertainty of the reconstructed cell-level energy is commonly referred to as *noise*. Noise has three primary contributions:

- *quantization noise* in the ADCs,
- *thermal noise* in the analog readout electronics,
- *pileup noise* due to the overlap between signal pulses and the detector's high event rate.

The first two are collectively called *electronic noise* due to their similarity.

Quantization noise stems from the fact that the ADC introduces a rounding error into the signal. This process is not fully deterministic; rather, the least significant bit tends to fluctuate, even in the absence of input noise. This noise is fully uncorrelated.

Thermal noise occurs in the entire analog section of the electronics. However, the cold electronics dominate due to their amplification by the preamplifier. Though simple thermal noise is approximately uncorrelated [104, p. 69], the analog electronics shape it in many, non-trivial ways. As such, the thermal noise must be assumed to be correlated in some way.

Pileup noise originates in the detector proper and is caused by its high event rate, as described in section 3.1.5. In-time pileup can be understood as a randomly distributed energy added on top of every signal energy deposition. While various methods exist to remove it e.g. on the level of reconstructed jets [105, sec. 1], on the calorimeter cell level, it can only be subtracted on average.

Out-of-time pileup, on the other hand, is caused by the overlap between pulses of hits that closely follow each other. It could be removed completely from any cell hit in principle, if the timing and amplitude of all past hits of the same calorimeter cell were known.




Figure 3.13: Thermal and pileup noise σ over the pulse peaking time t_p [102, fig. 7b]. The peaking time is defined as the time between the points when the pulse reaches 5 % and 100 % of its amplitude. It depends only weakly on the amplitude and increases with the time constant τ of the shaper.

The magnitude of pileup noise increases both with the luminosity (see section 3.1.5) and with $|\eta|$. The latter is primarily due to the fact that cells in the forward region are wider in η and thus are hit by more particles per bunch crossing. For this reason, pileup noise dominates over electronic noise at high $|\eta|$ and the latter can be neglected. This effect will only increase in the future, when it is planned to increase the luminosity of the LHC even further.

The time constant τ of the shaper affects both electronic and pileup noise:

- Electronic noise is best suppressed with a high value of τ , which corresponds to a lower cut-off frequency of the shaper's band-pass filter.
- Out-of-time pileup noise is best suppressed with a low value of τ , which shortens the positive lobe of the signal pulses and thus reduces overlap.

Figure 3.13 shows this trade-off. The current time constant of 13 ns optimizes the combined noise [102, sec. 4.3]. As the LHC luminosity is scheduled to be increased, this choice is re-evaluated in section 7.3.

Chapter 4

Upgrades to the ATLAS Liquid-Argon Calorimeters

The LHC is the most powerful particle accelerator in the world and will remain so for at least the next two decades [106, sec. 1.1]. In order to extend its operability and increase its instantaneous luminosity L (see section 3.1.4), it is planned to be upgraded to the so-called High-Luminosity LHC (HL-LHC). This will increase L from its current peak value of $2.1 \times 10^{34}/(\text{cm}^2 \text{ s})$ [107, p.5] to a value between $5 \times 10^{34}/(\text{cm}^2 \text{ s})$ and $7.5 \times 10^{34}/(\text{cm}^2 \text{ s})$. To handle the increased amount of data, background interactions, and radiation, the experiments need to be upgraded as well.

The upgrade towards the HL-LHC is split into two stages: one is currently underway, while the LHC is in Long Shutdown 2 (LS2); the other will be applied in 2025 during Long Shutdown 3 (LS3). The respective upgrades of the ATLAS detector are referred to as *Phase-I Upgrade* and *Phase-II Upgrade*. Figure 4.1 shows the expected timeline for this process. While section 4.1 summarizes the physics goals of the HL-LHC, sections 4.2 and 4.3 describe the Phase-I and Phase-II Upgrade respectively.

4.1 Physics Goals

The physics program at the HL-LHC has two important goals: precision measurements of the Higgs sector and searches for BSM physics [109, sec. 8, 110, sec. XI.3].

To verify that the newly discovered particle at the LHC is indeed the Standard Model Higgs boson, its couplings to as many final states as possible must be measured and compared to Standard Model predictions. Of particular importance to this are: [110, sec. XI.3]:

- the decay $H \rightarrow \gamma\gamma$ as well as the vector boson fusion channels $H \rightarrow ZZ^{(*)} \rightarrow \ell\ell\ell\ell$ and $H \rightarrow WW^{(*)} \rightarrow \ell\nu\ell\nu$ for the gauge bosons;
- the Higgs decays to $b\bar{b}$, $\tau^+\tau^-$, and $\mu^+\mu^-$ for quarks and leptons.

Furthermore, the Higgs self-coupling must be measured to confirm the Brout–Englert–Higgs mechanism as the cause of EWSB. The best channel for this is $H \rightarrow HH \rightarrow b\bar{b}\gamma\gamma$. [109, sec. 8.2.2].

On the search for BSM physics, multiple leads will be pursued:




Figure 4.1: LHC plan for current and future upgrades as of January 2020 [108]. Also shown are the center-of-mass energy of collisions (upper red line) and luminosity (lower red line). The extended year-end technical stop (EYETS) was a seventeen-week period during which extended repairs and preparations for the LS2 were made.

- The vector boson scattering process $VV \rightarrow VV$ (with $V = W, Z$) will be probed for the possibility that there are other, unknown mechanisms that preserve unitarity of the cross-section in the longitudinal mode. [109, sec.8.2].
- The search for supersymmetric extensions of the Standard Model remains a high-priority task in the LHC program. One of the most promising channels for it is the associated production of charginos and neutralinos. [110, sec. XI.3.5].
- Finally, high-mass particles predicted by more exotic models are being searched for, e.g. Kaluza–Klein gravitons or dilepton resonances [110, sec. XI.3.6].

To accomplish these goals, the LAr calorimeters must operate at optimal performance, so that electrons and photons can be identified with high efficiency. In addition, the reconstruc-




Figure 4.2: Exploded representation of an EMB trigger tower in η - ϕ space [111, p. 17]. Each square represents an area of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. In the current system (left), calorimeter cells (black outline) are summed across layers into a single trigger tower. After the Phase-I Upgrade (right), calorimeter cells (gray outline) are summed into super cells (black outline).

tion of boosted jets in the far-forward region is crucial for the measurement of vector boson fusion and scattering processes.

4.2 Phase-I Upgrade

In Run 2, $L_{\max}$ reached a peak value of $2.1 \times 10^{34}/(\text{cm}^2 \text{ s})$. After the LS2, in Run 3, this value is expected to be held or further increased to $3 \times 10^{34}/(\text{cm}^2 \text{ s})$ [107, p.5]. This corresponds to a $\langle\mu\rangle$ of about 80. (see sections 3.1.4 and 3.1.5 for definitions.) Consequently, the number of background events that need to be suppressed increases as well.

The Phase-I Upgrade of the LAr calorimeters affects exclusively the L1 trigger path; the main readout path stays unmodified. The goal of this upgrade is to deploy more sophisticated filter algorithms to the L1 trigger. This is so that, despite the increased amount of background events, the trigger rate is kept below 110 kHz, the maximum supported by the main readout.

On a high level, the main change to the trigger path is the transition from trigger towers to *super cells* in the Electromagnetic Calorimeter, as shown in fig. 4.2. This means that summing no longer occurs across layers and that the trigger thus will receive a depth profile of energy depositions. Additionally, super cells have a finer granularity of $\Delta\eta \times \Delta\phi = 0.25 \times 0.1$ in the front and middle layer, thus generally providing higher-resolution information as well. The higher resolution makes it necessary to move energy reconstruction from the L1 trigger to the back-end electronics. The trigger, in return, can deploy algorithms that operate on a higher level of abstraction and perform simple particle identification.




Figure 4.3: Overview of the LAr readout electronics after the Phase-I Upgrade [111, p. 18]. New components are indicated by red outlines and arrows. There are two inaccuracies in this diagram: 1. The FEB has its own FELIX connection, communication is not routed through the back end. 2. In the LTDB, the TBB path splits off before shaping, not after.

Figure 4.3 shows how the Phase-I Upgrade will change the readout electronics. As can be seen, the legacy trigger path via the TBB and TDB will stay in place and active throughout Run 3. It will serve as a stable and well-understood physics trigger and as a baseline for the validation of the new trigger [91, sec. 3.9.1].

4.2.1 Front-End Electronics

Most of the changes to the front-end electronics accommodate the new super-cell-based geometry. The LSBs are replaced with new boards that provide finer-grained sums; the baseplanes, which route the signals between the different boards, are upgraded to handle the increased amount of data to be transferred; and new LAr Trigger Digitizer Boards (LTDB), which digitizes the signal on-detector, are added. The LTDBs have four tasks:




Figure 4.4: Draft of the LATOME firmware architecture [114, fig. 3].

- They perform the final shaping of the analog signal pulse. This also amplifies the signal with a gain that optimally uses the ADC's dynamic range.
- Their analog sections also sum the (unshaped) inputs and forward the sums to the legacy trigger path. Configurable delay lines ensure that the signals arrive in the TBBs and TDBs with the same timing as before the upgrade.
- The on-board ADCs sample and digitize the signal at a rate of 40 MHz. The least significant bit is different for each detector region, but generally lower than before the upgrade.
- The LTDBs also take over the functionality of the crate monitoring boards. The latter are removed in order to make more space on the baseplanes.

The digitized samples are ultimately serialized by LOCx2 chips [112] and sent to the back end and the monitoring system.

4.2.2 Back-End Electronics

The LAr Digital Processing System (LDPS) is added to the back end for the task of energy reconstruction. It follows the Advanced Telecom Computing Architecture (ATCA) [113, sec. 5.2.1], with LAr Digital Processing Blades (LDPB) as the primary processing units. Each LDPB is equipped with four mezzanine cards, the LAr Trigger Processing Mezzanines (LATOME).

As shown in fig. 4.4, the FPGA on each LATOME board is programmed in a modular fashion [114]. Each module solves one particular task:

- The **input stage** receives digitized samples from the LTDB and deserializes them.
- The **configurable remapping** groups data from super cells that belong to the same trigger tower.
- The **user code** calculates E_T , assigns it to the correct BC, and detects saturated pulses. It also transmits several quality bits that signal the success of each operation.
- The **output summing** sends the E_T values to the L1 calorimeter trigger. The values are partly summed and partly at full super-cell resolution (see section 4.2.3).
- The **TDAQ/monitoring** buffers data and sends it to the monitoring system upon request.

The user code consists of three blocks: the filter and selection block (consisting of two sub-blocks), the baseline correction block, and the combine block.

- The filter sub-block applies an OF to the signal samples to calculate E_T and $\xi := E_T \tau$ (see section 5.6 for details). The filter coefficients are scaled by a factor as detailed below.
- The selection sub-block uses simple cuts on these values to detect calorimeter hits and saturation of the analog electronics.
- In parallel, the baseline correction block estimates how much the baseline value of the signal is shifted for the current BC.
- This estimate is applied to E_T in the combine block. For saturated samples, this block also checks whether there are high-energy deposits in neighboring cells; if not, the sample is labeled as a misidentified saturated signal.

Strictly speaking, the filter applied in the user code is a *scaled* OF. The output sent to the L1 calorimeter trigger must be proportional to E_T , but the input received from the front end is only partially proportional, as described in section 4.2.1. The scaling factor necessary to make the data proportional is absorbed into the OF coefficients.

Besides the LDPS, the other back-end upgrade is the evolution of the Readout System (ROS). This upgrade starts in Phase I and is continued in Phase II. It is planned to replace the highly customized connections of the various sub-detectors to TDAQ and the TTC with a uniform interface: the Front End Link Interface Exchange (FELIX). This also replaces the custom, point-to-point connections from the detectors to nodes in the TDAQ network with switched connections that facilitate load balancing. Figure 4.5 shows the readout architecture before and after this upgrade.

4.2.3 Level-1 Calorimeter Trigger

Although the trigger system is being upgraded to exploit the super-cell information passed to it by the new back end, the current system (fig. 4.6a) will remain fully operational throughout Run 3 as a stable baseline. Figure 4.6b shows a diagram of the planned changes.

Three new types of modules, so-called Feature Extractors (FEX), are added to the trigger: the eFEX, the jFEX, and the gFEX. The eFEX is the only one that operates at super-cell resolution. It identifies isolated electrons, photons, and tau leptons. The jFEX, operating at trigger-tower resolution, identifies isolated jets of size $\Delta R = 0.4$. The gFEX has an even lower resolution




Figure 4.5: Evolution of the global readout architecture from the present situation (Run 2) to the system after the Phase-II Upgrade [91, fig. 59]. The Readout Buffers (ROB) are the part of the ROS that stores trigger data until the HLT requests or expunges it.

(2×2 trigger towers or more), but in return, the full calorimeter is available to it. It calculates the transverse energy imbalance E_T^{miss} and identifies large-radius jets.


The FEXs summarize their results as *trigger objects* and forward them to the L1 topological trigger. The topological trigger combines these objects with those of the L1 muon trigger and sends the combined objects to the CTP. The CTP, as before, makes the L1 trigger decision whether to accept or reject the event. The topological trigger has already been added during LS1; during Phase I, it is only modified by including FEX trigger objects.

4.3 Phase-II Upgrade


After LS3, the HL-LHC is expected to reach an instantaneous luminosity L_{max} of 5×10^{34} to $7.5 \times 10^{34}/(\text{cm}^2 \text{ s})$. This corresponds to 140 to 200 pp interactions per BC on average, which is more than twice the amount expected after LS2. Because the current hardware trigger likely cannot suppress the resulting amount of background events sufficiently, it will have to be upgraded.

4.3.1 Hardware Trigger

Figure 4.7 shows the architecture of the upgraded TDAQ system. Most notably, the L1 calorimeter trigger is going to be replaced by an L0 trigger. It will have a trigger rate of either 1 MHz



(a) The L1 calorimeter trigger as of Run 2.



(b) The planned Phase-I Upgrade of the L1 calorimeter trigger.

Figure 4.6: The L1 calorimeter trigger before and after the Phase-I Upgrade [91, figs. 18f.]. The components in light green have been added during LS1. The components in bright yellow are to be added in Phase I.




Figure 4.7: Baseline plan for the TDAQ system in Phase II [115, fig. 14.4]. The main readout occurs via FELIX. The trigger data line from the calorimeters to the L0 muon trigger only concerns the outer layer of the tile calorimeter and is out of scope for this thesis.




Figure 4.8: Connectivity diagram between the calorimeters and the FEXs [115, fig. 7.5]. LASP boards that process the EMB are not connected to the L0 calorimeter trigger. Though the LDPS also calculates energy sums for the HEC and FCal for redundancy, only the LASP sends this information to the trigger.

(baseline plan [115, sec. 4.2]) or up to 4 MHz (evolved scenario with a new L1 trigger [115, sec. 14.2]) and a latency of $10\ \mu\text{s}$ (both scenarios). The increased latency with respect to Phase I will allow for a more refined event selection at the lowest trigger level.

The FEXs, unchanged from Phase I, will still be part of the L0 trigger. However, a new fFEX will be added: it will reconstruct electromagnetic objects in a region $2.5 < |\eta| < 4.9$ and jets in a region $3.1 < |\eta| < 4.9$. To do so, it will receive data at full resolution from the FCal, HEC, and inner EMEC. Figure 4.8 shows the data flow from the calorimeters to the FEXs.

In addition, a new Global Trigger is inserted between the L0 calorimeter trigger and the CTP. It effectively replaces the L1 topological trigger. It receives not only the trigger objects calculated by the FEXs, but also reconstructed energies from all calorimeters at full resolution¹. These data are used for a series of algorithms similar to the offline analysis, e.g. topological clustering. Based on these results, new and refined trigger objects are sent to the CTP.

4.3.2 Readout Electronics

In addition to the trigger, the LAr calorimeter readout electronics are also planned to be upgraded. There are three major reasons for this effort:

- The upgraded L0 trigger will need a way to access the calorimeters at full resolution at a frequency of 40 MHz. The current readout does not have the necessary bandwidth to provide this information. (This is of no concern to the LDPS, which sends its data to the

¹To conserve bandwidth, only cells where the transverse energy exceeds twice the average noise ($|E_T| > 2\sigma$) are transmitted.




Figure 4.9: Overview of the LAr readout electronics after the Phase-II Upgrade [117, p. 3]. New components are underlaid with dark blue, sharp-cornered rectangles.

unmodified FEXs.) Furthermore, the current FEBs do not have the necessary data buffers to support the increased L0 trigger latency.

- The current FEBs have been qualified for 10 years of operation, with some safety factors included [116]. At the beginning of Run 4, they will be 15 to 20 years old and will not withstand the planned 8 years of HL-LHC operation without severely reduced performance due to radiation damage [110, sec. V.2.1]. This is of no concern to the LTDB and the cold electronics, which all have been qualified for the total radiation dose expected at the HL-LHC [111, sec. 2.5].
- There are several ASICs in the front end that have been fabricated using outdated technology and can no longer be replaced [109, sec. 3]. While spare parts exist, they will likely not suffice for the entirety of the HL-LHC program [111, sec 6.3.1].

Based on these considerations, it is planned to replace the FEBs, the RODs, and the calibration boards. Figure 4.9 shows a block diagram of the upgraded readout electronics.

4.3.3 Back-End Electronics

The RODs are replaced by Liquid Argon Signal Processors (LASP), an evolution of the LDPS developed for Phase I. Like the LDPS, they are based on the ATCA, with two LASPs per ATCA blade. However, the FPGAs to be used will have more processing power, allowing the implementation of more sophisticated algorithms.

The LASPs perform the following four tasks:

- They receive the calorimeter signal from the front end via optical links. The signal is digitized and arrives unbuffered at the rate of the LHC clock. It is available on two different, overlapping gain scales, one of which must be chosen on each tick.
- Using the digitized calorimeter signal, it identifies hits and reconstructs their deposited energy and timing. Due to the higher amount of available processing power, it is possible to use more sophisticated filters with active suppression of out-of-time pileup. (This is investigated in section 7.7.) It is also possible to implement different filters for the trigger and the main readout path; the former would be optimized to correctly identify hits, the latter to estimate the deposited energy and hit timing accurately.
- Data from all cells whose reconstructed energy exceeds a certain threshold are sent to the Global Trigger and, in case of the end-caps, to the fFEX. If too many cells pass the threshold, the LASPs must automatically discard the lowest-energy hits.
- The data must be buffered for at least $10\ \mu\text{s}$ and, if an L0 trigger accept signal arrives, sent to the DAQ network. Each LASP must be able to handle up to four accept signals in consecutive BCs.

4.3.4 Front-End Electronics

The upgraded Front-End Board 2 (FEB2) will have several improvements over its predecessor, though the general data flow stays similar. A combined preamp/shaper ASIC will handle both amplification and pulse shaping in one circuit. Its output drives an ADC that samples and digitizes the signal². The samples are sent to the back end, which further processes them. To guarantee synchronization between front end and back end, a bunch crossing identifier (BCID) is transmitted for each data packet. This is a simple integer number that increases with each BC and is reset after each LHC orbit via the global Bunch Counter Reset (BCR) signal.

The preamp/shaper chip not only amplifies and shapes the signal, it will also split it onto two overlapping gain scales. These scales are separated by a factor of 30. Together, they cover a dynamic range of 16 bits. This covers the entire range of energies possible at the LHC while keeping the quantization noise below the intrinsic calorimeter resolution. This is important to avoid degrading the total resolution by more than 5 % [111, sec. 3.7.1].

The LAr cryostat will not be opened during the Phase-II Upgrade. This means that for the HECs, amplification will continue to be handled by the cold electronics. To account for this, the

²The signal may be sampled at a rate of either 40 MHz or 80 MHz. It is foreseen that a sampling rate of 40 MHz will be used [111, sec. 3.7.2]. For a comparison of how both options impact the energy resolution see section 7.4.

HECs will be read out by modified FEBs, on which the preamp/shaper chip inverts the signal and adjusts its gain³.

The FEB2 has to stay compatible with the LTDB. Hence, it will also have to sum the calorimeter signals into super cells and transmit these onto the trigger path. It has not been decided yet whether this summing will be done within the preamp/shaper chip or on dedicated chips comparable to today's LSBs.

There are currently two prototypes for the preamp/shaper chip: the LAUROC [118], based on 130 nm CMOS technology; and the HLC1 [119], based on 65 nm CMOS technology. A comparison of both is given in [111, sec. 5.1.2] and in [120]. A third approach, based on bipolar silicon-germanium processes, is kept as a backup plan. It is disfavored because of its high power consumption.

The ADC is being pursued with three different approaches. The baseline solution is to use one 14-bit ADC per gain scale. The ADC is custom-made. It consists of a Dynamic Range Enhancement (DRE) block followed by a 12-bit ADC. The 12-bit ADC is the same that has been developed for the LTDB. The DRE block splits the signal onto two paths, one with gain 1 and one with gain 4. It determines the two most significant bits and, based on them, picks one path to present to the 12-bit ADC. The two results are combined in a way that preserves the required bit precision.

Besides the DRE, an alternative method based on a multiplicative digital-to-analog converter (MDAC) is currently being pursued. It does not change the architecture of the ADC, but promises to reduce the quantization noise significantly [121].

Besides this fully custom ADC, the other two approaches are a commercial off-the-shelf (COTS) ADC and a custom ADC chip that integrates a commercial intellectual-property (IP) core. The former would require an additional ASIC to interface the ADC with the other ATLAS electronics. It thus only serves as a backup solution in case the other two approaches don't meet the requirements.

³In addition to inversion and gain adjustment, the current preshaper also performs a pole-zero cancellation to equalize the signal pulse shapes of different HEC cells. It has been found that this has only negligible impact on the energy resolution, so it will not be done in Phase II [111, p. 79f].

Chapter 5

Noise Suppression With Digital Filters

Conceptually, the back end of the LAr calorimeters performs two important tasks: bunch crossing (BC) identification and energy reconstruction. The former means that for each cell and each BC, the back end has to determine whether there has been a hit by particles from the pp collision. The latter means that if there has been a hit, the back end should estimate how much energy has been deposited in that cell by the incident particles. Because the time of flight of the particles and the lengths of the readout cables vary considerably, the back end must also estimate the *timing* of each hit, i.e. the time between the hit and the last LHC clock signal.

In this chapter, various algorithms for these purposes are reviewed. In section 5.1, the quantities and symbols used in this chapter are introduced. Several terminology pitfalls and ambiguities in definition are pointed out as well. Section 5.2 serves as an introduction to digital filters, the algorithms used primarily for low-latency energy reconstruction. In each subsequent section, an algorithm is introduced and its properties investigated.

It must be noted that the final two approaches – sparse signal restoration in section 5.8 and artificial neural networks in section 5.9 – have not been implemented in AREUS nor investigated in chapter 7 due to time constraints. Nonetheless, they have been included in this chapter as both an inspiration and a reference to other physicists, who may wish to further investigate the problem of online energy reconstruction.

5.1 Terminology

In the study of digital signal processing, many concepts have multiple competing definitions. Hence, it is important to use a common subset of these definitions when comparing different algorithms. In this section, the conventions used throughout the chapter are summarized.

5.1.1 Continuous Formulation

When a particle crosses a calorimeter cell after a BC and deposits an energy E , the analog readout electronics respond with a voltage pulse w . Ignoring the finite flight time τ of the




Figure 5.1: A typical response of the readout electronics to a hit in a LAr cell. Notable features are the narrow positive lobe and the long, negative undershoot. The integral over the entire pulse is approximately zero. The circles mark points where the pulse is sampled by the ADC.

particle¹ and assuming the response scales linearly with E , it is:

$$\begin{aligned} w : \mathbb{R} &\rightarrow \mathbb{R}, t \mapsto w(t), \\ w(t) &:= E \cdot h(t) + u(t). \end{aligned} \quad (5.1)$$

The function h is the *response function* of the system and u describes the noise of the readout electronics. Various authors also call h the *analog pulse shape*, *pulse form* or *waveform*.

The function h describes an ideal and normalized pulse as produced by the readout electronics. A typical pulse is shown in fig. 5.1. Its domain is limited to an interval $[0; T]$ where T is the *pulse length*. Because the analog shaper of the readout electronics contains a high-pass filter, the integral of the response approximately vanishes:

$$\int_{-\infty}^{\infty} h(t) dt = \int_0^T h(t) dt = 0. \quad (5.2)$$

5.1.2 Discrete Formulation

The front-end electronics digitize the electric signals from the calorimeter and send a sequence of discrete voltage samples to the back end. This sample sequence is the input of the energy-reconstruction algorithm. The sampling period t_s of this sequence may be assumed to be any integer fraction of the BC period, 25 ns, though no sampling rate higher than two samples per BC is likely to be pursued.

¹It will be considered again in section 5.6.

A singular hit with energy E causes a sequence w of non-zero samples to be registered:

$$\begin{aligned} w &: \mathbb{N}_0 \rightarrow \mathbb{R}, n \mapsto w[n], \\ w[n] &:= E \cdot h[n] + u[n] \\ &:= E \cdot h(nt_s) + u(nt_s), \end{aligned} \quad (5.3)$$

where $h[n]$ is a sequence over the response function $h(t)$ and $u[n]$ over the noise $u(t)$. In practice, energy reconstruction needs to deal with not only one hit, but multiple hits in quick succession. This can be modelled by turning E into a sequence $E[t]$ of N hits, also with the period t_s . This sequence should be much longer than an individual pulse, i.e. $N \gg L$, where $L := \frac{T}{t_s}$ is the discretized pulse length. The sequence E should also be stationary in the wide sense, i.e. its average should not change too much².

Under these assumptions, the input to the energy reconstruction becomes:

$$w[n] := \sum_{k=0}^{\min\{n, L-1\}} h[k]E[n-k] + u[n]. \quad (5.4)$$

The sum over k is an acyclic convolution $(*)$ of E with h . This can be used to rewrite eq. (5.4) in terms of sequences instead of their elements:

$$w := h * E + u. \quad (5.5)$$

In practice, the acyclic convolution is often approximated by the cyclic convolution, defined as $(x \otimes y)[n] := \sum_{k=0}^{N-1} x[k]y[n-k] \pmod{N}$, where out-of-bounds indices wrap around. This is appropriate if $x[k]$ is only non-zero for indices $k \ll N$.

5.1.3 Autocorrelation and Cross-Correlation

E and u are stochastic sequences, i.e. each of their elements is itself a random variable. Because they are wide-sense stationary, each element has the same statistical properties:

$$\mathbb{E}[E[i]] = \mathbb{E}[E[j]] \quad \forall i, j \in \mathbb{N}_0, \quad (5.6)$$

where $\mathbb{E}[\cdot]$ is the expectation. This makes it possible to speak about quantities such as *the* population mean μ_X and population variance σ_X^2 of an entire sequence X .

For a stochastic sequence X , μ_X and σ_X^2 may be estimated by the sample mean $\bar{x}$ and the sample variance s_x^2 :

$$\bar{x} := \frac{1}{N} \sum_{n=0}^{N-1} x[n], \quad (5.7)$$

$$s_x^2 := \frac{1}{N-1} \sum_{n=0}^{N-1} (x[n] - \bar{x})^2 = \frac{N}{N-1} \left(\overline{x^2} - \bar{x}^2 \right). \quad (5.8)$$

²In practice, the sequence E is non-stationary due to both the bunch-train pattern employed in the LHC and the decreasing luminosity during a run.

where x is a concrete realization of X . (For simplicity, the distinction between both will be suppressed from here on out.) Generally, s_x^2 underestimates σ_x^2 ; the bias only vanishes if there is no correlation between the sequence elements.

The energy-reconstruction algorithms presented here not only need μ and σ^2 to describe E and u , but also a measure of *similarity* between two sequences. The simplest such measure is the *cyclic cross-correlation* R° . For two sequences x and y of length N , it is:

$$R_{xy}^\circ : (-N, N) \rightarrow \mathbb{R},$$

$$R_{xy}^\circ(n) := \frac{1}{N} \sum_{k=0}^{N-1} x^*[k]y[k+n] \pmod{N}, \quad (5.9)$$

where $(\cdot)^*$ denotes complex conjugation and $\pmod{N}$ means that indices outside the interval $[0; N-1]$ “wrap around”. The *lag* parameter n shifts one of the sequences by n samples and accounts for the fact that one sequence might replicate the other in a delayed manner.

The *autocorrelation* of a sequence x is defined as its cross-correlation with itself, R_{xx}° . As a measure of self-similarity, it is maximal for a lag of zero and converges to the quadratic mean with increasing n :

$$\mathbb{E}[R_{xy}^\circ(0)] = \mu_x^2 + \sigma_x^2, \quad (5.10)$$

$$\lim_{n \rightarrow \infty} \mathbb{E}[R_{xy}^\circ(n)] = \mu_x^2. \quad (5.11)$$

The speed with which R° converges to μ_x^2 can be regarded as a kind of “coherence time”.

For stochastic sequences, R° is a biased estimator of the population cross-correlation. Better results can be achieved using the *acyclic cross-correlation*, defined as:

$$R_{xy} : (-N, N) \rightarrow \mathbb{R},$$

$$R_{xy}(n) := \frac{1}{N - |n|} \sum_{k=\max(0, -n)}^{N-1-\max(0, n)} x^*[k]y[k+n]. \quad (5.12)$$

It is unbiased, but has a high variance for $|n| \approx N$. If one is only interested in the results for $n \ll N$, this presents no issue and $R(n) \approx R^\circ(n)$. As a note, the sum in eq. (5.12) is sometimes written as an operator “ $\star$ ”, i.e. $(x \star y)(n) := R_{xy}(n) \cdot (N - |n|)$.

Finally, a concept that will become useful is the *conjugated reversal* of a sequence x :

$$x^\dagger[n] := x^*[N-1-n]. \quad (5.13)$$

With it, the cross-correlation can be expressed as a convolution:

$$R_{xy}(n) = \frac{1}{N - |n|} (x^\dagger * y)[N-1+n], \quad (5.14)$$

where the shift by $N-1$ is necessary to make x and y “overlap”. This correspondence is important because it means that certain relationships between convolution and Fourier transform also hold for the cross-correlation.

5.1.4 The Fourier Transform

Transforming sequences into the frequency domain can make information visible that is difficult to discern in the time domain. However, there are several competing definitions of this transformation, sometimes with different normalization. Thus, this section will briefly introduce the conventions used in this chapter. Generally, the conventions field of digital signal processing are used [122].

For a finite sequence x of length N , the discrete Fourier transform (DFT) is:

$$\hat{x}[k] := \mathcal{F}[x][k] := \sum_{n=0}^{N-1} x[n] \cdot e^{-2\pi i \frac{nk}{N}}, \quad (5.15)$$

$$x[n] = \mathcal{F}^{-1}[\hat{x}][n] := \frac{1}{N} \sum_{k=0}^{N-1} \hat{x}[k] \cdot e^{2\pi i \frac{nk}{N}}. \quad (5.16)$$

The same equations also hold if x is infinite but periodic with period N . In either case, the sequence $\hat{x}$ is complex, infinite, and periodic with period N ; it is called the *spectrum* of x . Though the properties of the DFT are well summarized elsewhere [123], three are particularly important here:

1. If x is real, then $\hat{x}$ is symmetric under conjugated reversal: $\hat{x}^\dagger = \hat{x}$. This means that only half the values must be calculated for most practical sequences.
2. The DFT of $x^\dagger$ is $\hat{x}^*$.
3. For the circular convolution, it is: $\mathcal{F}[x \otimes y] = \mathcal{F}[x] \cdot \mathcal{F}[y]$.

The last two properties lead to the Wiener–Khinchin theorem [124, 125]:

$$R_{xx}^\circ(n) = \frac{1}{N} \mathcal{F}^{-1}[|\hat{x}|^2][n], \quad (5.17)$$

and Parseval's theorem [126] as a special case:

$$\mu_x^2 + \sigma_x^2 = \mathbb{E}[R_{xx}^\circ(0)] = \mathbb{E}\left[\frac{1}{N} \sum_{k=0}^{N-1} |\hat{x}[k]|^2\right]. \quad (5.18)$$

The sequence $|\hat{x}|^2$ is referred to as the *power spectrum* of x . An important consequence of Parseval's theorem is that the variance of a zero-mean sequence can be estimated by a sum over its power spectrum.

Stochastic sequences like the electronic noise are inherently aperiodic. In addition, the complex phases of their spectra are random, i.e. they contain no useful information. Consequently, the DFT does not describe them well. A better fit is the power spectral density (PSD):

$$S : \left[-\frac{f_s}{2}, \frac{f_s}{2}\right] \rightarrow \mathbb{R}_0^+, f \mapsto S(f),$$

$$S(f) := \lim_{N \rightarrow \infty} \frac{1}{N f_s} \mathbb{E}\left[\left|\sum_{n=0}^{N-1} x[n] e^{-2\pi i n \frac{f}{f_s}}\right|^2\right], \quad (5.19)$$

where $f_s := t_s^{-1}$ is the sampling rate. The PSD is based on the discrete-time Fourier transform (DTFT), which is a Fourier transform for aperiodic time-discrete signals. In contrast to the power spectrum, the PSD is a continuous function of f . If it were extended beyond its domain of $[-f_s/2; f_s/2]^3$, it would continue periodically due to aliasing [127, sec. 3.2]. It further differs from the power spectrum in units: if the signal is measured in V, the power spectrum has units V^2 , whereas the PSD has units V^2/Hz .

Calculating the PSD theoretically requires two averages over the measured data x : a time average over the index n and an ensemble average over all realizations of x . However, if x is strictly ergodic, both averages are the same and the PSD can be approximated with the power spectrum of x :

$$S(f = \frac{k}{N} f_s) \approx \frac{|\hat{x}[k]|^2}{N f_s}. \quad (5.20)$$

Variants of the Wiener–Khinchin theorem and Parseval’s theorem also hold for the PSD:

$$R_{xx}(n) = \mathcal{F}_{\text{DTFT}}^{-1}[S][n] := \int_{-f_s/2}^{f_s/2} S(f) e^{2\pi i n \frac{f}{f_s}} df, \quad (5.21)$$

$$\mu_x^2 + \sigma_x^2 = \int_{-f_s/2}^{f_s/2} S(f) df, \quad (5.22)$$

where the finite sum is replaced by the integral of the discrete-time Fourier transform.

Finally, *one-sided* spectral densities must be introduced. Like the power spectrum, the PSD is symmetric if the underlying stochastic process is purely real. In this case, it is customary to use only the positive half of the PSD and scale it by a factor of two:

$$S_{\text{os}} : \left[0; \frac{f_s}{2}\right] \rightarrow \mathbb{R}_0^+, f \mapsto S_{\text{os}}(f),$$

$$S_{\text{os}}(f) := 2S(f), \quad (5.23)$$

$$\mu_x^2 + \sigma_x^2 = \int_0^{f_s/2} S_{\text{os}}(f) df. \quad (5.24)$$


Whenever one is given a PSD on the interval $[0; f_s/2]$, it is important to know whether it is truly a one-sided PSD or simply one half of a two-sided PSD, as they differ by a factor two. Figure 5.2 shows a typical one-sided PSD of noise in the analog LAr readout electronics.

5.2 Digital Filters

Digital filters are employed in the LAr back end to perform energy reconstruction. Given a *signal* sequence:

$$w := h * E + u, \quad (5.25)$$

³ $f_s/2$ is known as the *folding* or *Nyquist frequency* f_N . For a real signal, the interval $[0; f_N]$ is the smallest interval that contains the full information about the PSD.



(a) The PSD at negligibly high sampling rate.

(b) The PSD at $f_s = 40$ MHz.

Figure 5.2: Typical one-sided PSD of electronic noise in the analog LAr readout. In the right diagram, the noise is under-sampled and thus distorted by aliasing. Nonetheless, in accordance with Parseval's theorem, its variance remains the same.

with some additive noise u , they should reconstruct the energy sequence E . There are many approaches for such filters, each with their own trade-offs, and they must be evaluated carefully.

Two important properties of filters are *stability* and *latency*. Stability is usually defined in terms of the Bounded-Input-Bounded-Output (BIBO) criterion: for any bounded input, the filter must produce a bounded output. Unstable filters are highly undesirable as they may produce run-away outputs that render them useless until their internal state is reset. The latency of a filter is the delay (in units of t_s) introduced by it to the readout chain. It can be estimated by applying the filter to the analog pulse h and taking the distance between the maximum sample of h and the response.

Linear filters are the category of digital filters that has been researched the most and is well understood. A filter is linear if its output $y[n]$ at time step n is a linear function of its previous outputs and inputs x :

$$y[n] = \frac{1}{b[0]} \left(\sum_{k=0}^{P-1} a[k]x[n-k] - \sum_{k=1}^{Q-1} b[k]y[n-k] \right), \quad (5.26)$$

where P and Q are the feed-forward and feed-back *filter depth* respectively and a and b are the corresponding filter coefficients. The stability of linear filters can be investigated via the Z-transform [128, sec. 13.2], a generalization of the discrete-time Fourier transform.

A special group of linear filters are those without feed-back terms. For them, eq. (5.26) simplifies to:

$$y[n] = \sum_{k=0}^{P-1} a[k]x[n-k], \quad (5.27)$$

or equivalently:

$$y = a * x, \quad (5.28)$$

i.e. a convolution of the input with the sequence of filter coefficients. The defining characteristic of such filters is their finite impulse response (FIR): when applied to the unit impulse, their response returns to zero in finite time. Consequently, filters of this category are called FIR filters.

One advantage of FIR filters is that they are stable by definition. In the frequency domain, they can be understood as band-pass filters without infinite resonances. Another advantage is that they are easy to implement in hardware: they consist of only delay terms, additions, and multiplications. As such, they can be used on their own or provide a basic building block for more complex filters.

Finally, non-linear filters are much less well understood and no general theory of them exists. Nonetheless, some of them – such as the extended Kalman filter [129] – have important applications in the field of digital signal processing. One simple, but important non-linear filter is the *Maximum Finder with threshold*:

$$y[n] := \begin{cases} x[n-1] & \text{if } x[n-2] < x[n-1] > x[n] \text{ and } x[n-1] > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (5.29)$$

It is trivially stable and has a latency of exactly one sampling period. At the LHC, it has been used during the Runs 1 and 2 on the LAr trigger readout as a fast peak-detection algorithm. Nonetheless, it is likely to be replaced in the Phase-1 Upgrade [114, sec. 5.2.2].

5.3 Wiener Filter

The Wiener filter was developed independently by Wiener [130] and Kolmogorov [131]. While it has been improved upon, e.g. by the Kalman filter [132, 133], it is still often used due to its simplicity. It is an FIR filter based on a minimization of the mean squared error.

5.3.1 Frequency Domain

Let w be the sequence of filter inputs and s be a sequence of *desired outputs*, both of length N . The simplest choice for s is $s := T_d E$, i.e. the true sequence of energy depositions delayed by d samples: $s[n] := E[n-d]$. The delay allows the filter to collect more information about the peak. In particular, energy resolution improves dramatically if the full positive lobe of h may be collected before a non-zero output is desired.

Let further c be the sequence of P Wiener filter coefficients and $\Delta := c * w - s$ the error sequence, i.e. the difference between desired and actual output. Then, the objective is to find a c that minimizes the expression:

$$\mathbb{E}[|\Delta|^2], \text{ i.e. } \mu_\Delta^2 + \sigma_\Delta^2.$$

A consequence of this objective is that the Wiener filter is generally biased, i.e. $\mu_\Delta \neq 0$.

When transformed into the frequency domain, this expression becomes:

$$\mathbb{E}[|\hat{c}\hat{w} - \hat{s}|^2],$$

or expanded:

$$\mathbb{E}[|\hat{c}\hat{w}|^2 + |\hat{s}|^2 - \hat{c}^* \hat{w}^* \hat{s} - \hat{c} \hat{w} \hat{s}^*].$$

This is a functional of the complex function $\hat{c}$. The derivative w.r.t. $\hat{c}$ can be calculated in the Wirtinger calculus [134, 135, sec. 0.4]. It is just what one would expect when naïvely treating $\hat{c}$ and $\hat{c}^*$ as independent variables. Setting it to zero gives the solution:

$$\hat{c}[k] = \frac{\mathbb{E}[\hat{w}^* \cdot \hat{s}][k]}{\mathbb{E}[|\hat{w}|^2][k]} = \frac{S_{ws}(kf_s/N)}{S_{ww}(kf_s/N)}. \quad (5.30)$$

S_{ww} is the PSD of the filter input w . S_{ws} can be regarded as a “cross-power spectral density” between w and s .

Equation (5.30) is well-suited to determine c via simulation, i.e. when the sequence E of true deposited energies is known. If one wants to derive c from actual measurements, one can insert the definitions of w and s and using the fact that E and the noise u are uncorrelated. The solution then is:

$$\hat{c} = T_d \frac{\hat{h}^* \cdot S_E}{|\hat{h}|^2 S_E + S_u}, \quad (5.31)$$

where S_E and S_u are the PSDs of E and u respectively. The arguments to $\hat{h}$, S_E , and S_u are analogous to eq. (5.30) and have been suppressed for clarity.

While this solution seems more complicated than eq. (5.30), all quantities can readily be measured at the ATLAS detector [111, sec. 6.1]:

- The Fourier transform of the time-shift operator is $T_d(f) = e^{-2\pi i d \frac{f}{f_s}}$.
- The analog pulse shape h can be obtained either directly from test beam data, or indirectly from delay runs [111, sec. 6.1.4].
- The PSD S_u can be estimated from electronic noise that is measured in pedestal runs.
- Energy depositions in different BCs are uncorrelated. Thus, $S_E(f) = \mu_E^2 \delta(f) + \sigma_E^2 / f_s$, with the Dirac distribution δ . The parameters μ_E and σ_E^2 can be determined from minimum-bias runs [136].

5.3.2 Time Domain and the Matrix Formalism

Because it is often inconvenient and inefficient to calculate filter coefficients in the frequency domain, it is useful to transform the above solutions back into the time domain analytically. By applying the inverse Fourier transform to eq. (5.30), one gets:

$$R_{ww}^\circ * c = R_{ws}^\circ, \quad (5.32)$$

where R_{ww}° and R_{ws}° should be understood to be correlation *sequences* for non-zero lag values. For sufficiently small filter depth P , this is approximately equal to the acyclic version:

$$R_{ww} * c = R_{ws}. \quad (5.33)$$

To solve this equation, one typically uses the following matrix formalism:

$$\begin{aligned} \mathbf{A} &= (a_{ij}) \in \mathbb{R}^{P \times P}, & a_{ij} &:= R_{ww}(j - i), \\ \mathbf{v}^\dagger &= (v_i^\dagger) \in \mathbb{R}^P, & v_i^\dagger &:= R_{ws}(-i), \\ \mathbf{c}^\dagger &= (c_i^\dagger) \in \mathbb{R}^P, & c_i^\dagger &:= c[P - 1 - i], \end{aligned}$$

where the autocorrelation matrix $\mathbf{A}$ is a real and symmetric Toeplitz matrix [137] and the cross-correlation $\mathbf{v}^\dagger$ and the coefficients vector $\mathbf{c}^\dagger$ are column vectors. The superscript dagger is a reminder that the elements in the vectors are ordered reverse w.r.t. to the sequences. This reversal is done in order to describe filter application as a scalar product:

$$y = x * a \Leftrightarrow \mathbf{y} = \mathbf{x}^\top \mathbf{a}^\dagger, \quad (5.34)$$

where $(\cdot)^\top$ denotes transposition.

In this formalism, the deconvolution of R_{ww} and c simplifies to a matrix inversion:

$$\mathbf{c}^\dagger = \mathbf{A}^{-1} \mathbf{v}^\dagger, \quad (5.35)$$

for which many efficient algorithms already exist.

Equation (5.31) can be transformed into the time domain in a similar manner. However, this process requires that the sequence E is uncorrelated and that the response function h has zero mean (which are both true). Under these assumptions, it is:

$$\mathbf{c}^\dagger = (\mathbf{H} \sigma_E^2 + \mathbf{U})^{-1} \mathbf{h} \sigma_E^2, \quad (5.36)$$

where:

$$\begin{aligned} \mathbf{H} &= (h_{ij}) \in \mathbb{R}^{P \times P}, & h_{ij} &:= (h \star h)(j - i), \\ \mathbf{U} &= (u_{ij}) \in \mathbb{R}^{P \times P}, & u_{ij} &:= R_{uu}(j - i), \\ \mathbf{h} &= (h_i) \in \mathbb{R}^P, & h_i &:= h[i - P - 1 + d]. \end{aligned}$$

The integer number d is the delay of the desired output. Both $\mathbf{H}$ and $\mathbf{U}$ are real and symmetric Toeplitz matrices. The term $h \star h$ is the unnormalized portion of the autocorrelation R_{hh} , as defined in eq. (5.12).

5.3.3 Peak Broadening and Narrowing

The Wiener filter as defined above ideally deconvolves the sequence E and the response function h , producing a sharp and narrow output whenever there was a hit. However, an ideal filter output would require infinitely many filter coefficients. Because the number of coefficients is limited in practice, the filter output will contain ripple-like artifacts right before and up to P samples after any considerable input. This is demonstrated in fig. 5.3. Additionally, while the Wiener filter minimizes the mean squared error $\mathbb{E}[|\Delta|^2]$ across all BCs, it tends to have a high variance when regarding only those BCs in which a hit is known to have occurred.




Figure 5.3: An example of the effect of peak broadening. This shows the original system response h (circle) as well as a regular Wiener filter's response to h (squares). Because the filter's depth ($P = 6$) is too small, it cannot model the high-frequency behavior correctly and there are visible ripples in its undershoot. A Wiener filter with broadened peak (triangles) does not suffer from this problem.

One solution to both problems is to modify s , the desired output sequence. Adding so-called *pre-peak* and *post-peak* values to s effectively broadens the response of the filter. As fig. 5.3 shows, this reduces both ripple effects and the variance in BCs where the filter has a strong response. In return, however, this naturally increases the error in BCs right before and after a strong response. Hence, a Wiener filter with broadened response gives bad results in cases where two hits occur in subsequent BCs.

The effect of *peak broadening* can be understood better in the frequency domain. Modifying the desired sequence s in the described manner is equivalent to multiplying the right-hand side of eq. (5.30) with a transfer function B :

$$B(f) = 1 + s_{\text{pre}} e^{2\pi i \frac{f}{f_s}} + s_{\text{post}} e^{-2\pi i \frac{f}{f_s}}, \quad (5.37)$$

where s_{pre} and s_{post} are the desired relative amplitudes of the pre- and post-peak sample. They are usually in the range $[-1; 1]$.

Figure 5.4 shows B for select values of s_{pre} and s_{post} . In particular, the most common choice ($s_{\text{pre}} = s_{\text{post}} = 0.5$) results in $B(f) = 2 \cos^2(\pi f / f_s)$, which suppresses high frequencies. The opposite effect, making the Wiener filter output even narrower, can be achieved by picking negative values for s_{pre} or s_{post} . This effect is conversely called *peak narrowing*.

5.4 Matched Wiener Filter

The Wiener filter as defined in section 5.3 minimizes the mean squared error across all BCs. Its output has a high variance, but it is well suited to separate the response function h from noise. In more practical terms, it provides a bad energy resolution, but excellent BC identification. This makes it well suited for the trigger path of the calorimeter readout, where only a rough estimate of the energy is needed but no event should be missed.




Figure 5.4: The absolute value of the transfer function B imposed on the Wiener coefficients' spectrum $\hat{c}$. Positive pre- and post-peak values suppress high frequencies while negative ones suppress low frequencies. If $s_{\text{pre}} \neq s_{\text{post}}$, B also acquires a non-constant complex phase, which is not shown here.

However, the situation is different on the main readout path: there, the BC in which a hit occurred is known and the energy should be reconstructed with the highest precision possible. Because of that, only the error $\mathbb{E}[|\Delta|^2]$ in these BCs should be minimized. This leads to a different minimization criterion and different filter coefficients. The difference to the regular Wiener filter is two-fold: While the variance in BCs with a cell hit decreases, the bias in surrounding BCs increases. As a consequence, multiple hits in quick succession cannot be separated as easily.

The solution of this minimization is superficially similar to eq. (5.35):

$$\mathbf{c}^\dagger = \mathbf{A}^{-1} \mathbf{v}^\dagger, \quad (5.38)$$

though the definition of $\mathbf{A}$ and $\mathbf{v}^\dagger$ differs:

$$\begin{aligned} \mathbf{A} &= (a_{ij}) \in \mathbb{R}^{P \times P}, & a_{ij} &:= \mathbb{E}[w[i+d-P+1]w[j+d-P+1]], \\ \mathbf{v}^\dagger &= (v_i^\dagger) \in \mathbb{R}^P, & v_i^\dagger &:= \mathbb{E}[w[i+d-P+1]s[d]], \\ \mathbf{c}^\dagger &= (c_i^\dagger) \in \mathbb{R}^P, & c_i^\dagger &:= c[P-1-i]. \end{aligned}$$

The primary difference to eq. (5.35) is the fact that the delay d is not averaged over and so it appears explicitly. It ensures that only samples around a peak in the input signal are considered.

In practice, $\mathbf{A}$ and $\mathbf{v}$ are calculated by simulating N individual pulses with true energies E_i and P samples w_{ij} around the peak. When these values are arranged into a vector $\mathbf{e} = (E_i) \in \mathbb{R}^N$ and a matrix $\mathbf{W} = (w_{ij}) \in \mathbb{R}^{N \times P}$, it is:

$$\mathbf{A} \approx \mathbf{W}^\top \mathbf{W}, \quad \mathbf{v} \approx \mathbf{W}^\top \mathbf{e}. \quad (5.39)$$

As in the previous section, eq. (5.38) can be analyzed in terms of $\mathbf{h}$ and $\mathbf{U}$. This leads to an alternative formula for the coefficients:

$$\mathbf{c}^\dagger = \frac{(\mu_E^2 + \sigma_E^2) \mathbf{U}^{-1} \mathbf{h}}{1 + (\mu_E^2 + \sigma_E^2) \mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}}. \quad (5.40)$$

By inserting eq. (5.40) into the definition of the reconstruction error Δ , the filter's bias and mean squared error can be calculated:

$$\mathbb{E}[\Delta] = \frac{\mu_E}{1 + (\mu_E^2 + \sigma_E^2) \mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}} \neq 0, \quad (5.41)$$

$$\mathbb{E}[|\Delta|^2] = \frac{\mu_E^2 + \sigma_E^2}{1 + (\mu_E^2 + \sigma_E^2) \mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}} \lesssim \frac{1}{\mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}}. \quad (5.42)$$

This filter does not have a universally accepted name. It is called *Wiener–Hopf filter* in some sources [138, 139], which ignores the important differences between this filter and the classic Wiener filter as described in section 5.3. In other sources, it is described generically as an “inverse filter based on least squares error” [140, sec. 3.2]. In this thesis, it will be referred to as *matched Wiener filter*, as belongs to the wider class of *matched filters*.

Matched filters are those linear filters whose coefficients $\mathbf{c}^\dagger$ are proportional to $\mathbf{U}^{-1} \mathbf{h}$. Common to all of them is that they maximize the signal-to-noise ratio of a sequence if the shape of the signal pulse is known. As such, they provide the best noise suppression among all linear filters. The scale of their coefficients is a free parameter and may be chosen freely – in this case to minimize $\mathbb{E}[|\Delta|^2]$.

5.5 Matched Wiener Filter Without Bias

The Wiener filter can be varied further by requiring it to be unbiased. This inherently must be traded off with a worse energy resolution. While such a filter does not see much use in practice, it is useful in order to better understand the filter presented in section 5.6.

The filter coefficients can be derived by minimizing the variance σ_Δ^2 under the constraint that $\mu_\Delta = 0$ using the method of Lagrange multipliers⁴. This leads to the coefficients:

$$\mathbf{c}^\dagger = \frac{\mathbf{U}^{-1} \mathbf{h}}{\mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}}, \quad (5.43)$$

which define a filter with the following bias and mean squared error:

$$\mathbb{E}[\Delta] = 0, \quad (5.44)$$

$$\mathbb{E}[|\Delta|^2] = \frac{1}{\mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}}. \quad (5.45)$$

The unbiased matched Wiener filter is highly similar to the biased one given by eq. (5.40). In particular, in the limit $\mu_E^2 + \sigma_E^2 \gg \sigma_u^2$ (the scale of signal hits is much higher than the scale of the noise), the biased filter tends towards the unbiased one. Furthermore, both are matched filters, so both provide the same degree of noise suppression.

⁴Minimizing the mean squared error $\mathbb{E}[|\Delta|^2]$ under the constraint $\mathbb{E}[\Delta] = 0$ leads to the same result.

5.6 Timing Reconstruction, Optimal Filtering, and Selection Criteria

A problem that has not been addressed so far is that of the hit timing. The timing $\tau \in [0; t_s]$ is used to describe variations in the phase of arriving signal pulses⁵. All derivations so far have assumed that the timing is constant and zero for all hits. In practice, this is approximately true for high-energy hits; but for lower energies (and especially for hits from soft pp collisions), the spread in time becomes non-negligible. In this case, the definition of the electronics' response must be adapted.

In the continuous formulation, eq. (5.1) becomes:

$$w(t) := E \cdot h(t - \tau) + u(t). \quad (5.46)$$

This can be transferred to the discrete case by assuming $\tau \ll t_s$ and approximating the Taylor expansion of h . Equation (5.4) thus becomes:

$$w[n] := \sum_{k=0}^{\min\{n, L-1\}} h[k] E[k-n] - \sum_{k=0}^{\min\{n, L-1\}} h'[k] \underbrace{E[k-n] \tau[k-n]}_{=: \xi[k-n]} + u[n],$$

$$w := h * E - h' * \xi + u, \quad (5.47)$$

where h' is a sequence over the derivative of h . Depending on the concrete definition of h' , τ may be given either in nanoseconds or in units of t_s .

5.6.1 Optimal Filtering

Equations (5.30), (5.35) and (5.38) use w and s directly and are agnostic regarding the precise nature of these quantities. For this reason, addition of the timing term $h' * \xi$ presents no issue when using them. For example, filter coefficients to estimate ξ can be calculated by replacing $s := T_d E$ with $s_\xi := T_d \xi$.

However, the analytic eqs. (5.31), (5.36) and (5.40) need to be adapted. Inserting the definition of w given by eq. (5.47) into the Wiener minimization problem leads to impractically complicated formulas. They depend on quantities such as σ_ξ^2 and $R_{\xi E}$, which are difficult to measure.

On the other hand, extending the unbiased matched Wiener filter of section 5.5 to include hit timing yields favorable results; the undesirable quantities cancel out exactly. The resulting Optimal Filter (OF) [141] consists of two sets of coefficients, commonly called $\mathbf{a}^\dagger$ and $\mathbf{b}^\dagger$, which estimate E and ξ respectively:

$$\mathbf{a}^\dagger := \frac{Q_2 U^{-1} \mathbf{h} - Q_3 U^{-1} \mathbf{h}'}{Q_1 Q_2 - Q_3^2}, \quad (5.48)$$

$$\mathbf{b}^\dagger := \frac{Q_3 U^{-1} \mathbf{h} - Q_1 U^{-1} \mathbf{h}'}{Q_1 Q_2 - Q_3^2}, \quad (5.49)$$

⁵Cf. section 6.2.2 for a more concrete definition of τ .

where:

$$Q_1 := \mathbf{h}^\top \mathbf{U}^{-1} \mathbf{h}, \quad Q_2 := \mathbf{h}'^\top \mathbf{U}^{-1} \mathbf{h}', \quad Q_3 := \mathbf{h}'^\top \mathbf{U}^{-1} \mathbf{h}.$$

One can confirm that in the limit $Q_3 \rightarrow 0$ of negligible interaction between E and ξ , $\mathbf{a}^\dagger$ converges to the coefficients of the unbiased matched Wiener filter.

5.6.2 Selection Criteria

Although the derivation of the OF assumed that the BC of a hit has already been identified, this filter may also be used when this is not the case. The procedure then is to apply the OF to all BCs:

$$\tilde{E}[n] := \sum_{k=0}^{\min\{n, P-1\}} a[k] w[n-k], \quad (5.50)$$

$$\tilde{\xi}[n] := \sum_{k=0}^{\min\{n, P-1\}} b[k] w[n-k], \quad (5.51)$$

and only to accept the result if the *estimated* sequences $\tilde{E}$ and $\tilde{\xi}$ pass certain selection criteria.

The Maximum Finder as the simplest selection criterion has been introduced in section 5.2. It accepts a BC whenever there is a peak in the sequence $\tilde{E}$; the sequence $\tilde{\xi}$ is not used at all. The criterion is trivial to implement and works well in low-pileup environments. However, it has several disadvantages: it is impossible, by construction, to select two hits in subsequent BCs. The Maximum Finder also introduces additional latency of one t_s . Furthermore, low-energy hits may be masked by the negative undershoot caused by high-energy hits, as shown in fig. 5.5. Thus, as the LHC moves towards higher luminosities and increasing pileup, the Maximum Finder is expected to become less relevant.

The $|\chi|$ criterion is based on the minimization of a χ^2 difference between the actual samples and the reconstruction result [142, p. 71]. Its definition has been optimized for calculation on an FPGA and it introduces no latency. Like the Maximum Finder, this criterion does not make use of $\tilde{\xi}$. It selects BCs where among the P most recent samples w_i , at least m satisfy:

$$|\chi_i| := |w_i - (\mathbf{w}^\top \mathbf{a}^\dagger) h_i| < \epsilon h_i. \quad (5.52)$$

The thresholds $m \in \mathbb{N}$ and $\epsilon < 1$ can be chosen freely.

The τ selection criterion requires that the reconstructed hit timing τ be within a certain range. Because the OF only estimates E and $\xi := E\tau$, implementing this criterion would normally require a division operation $\tilde{\xi}/\tilde{E}$, which is notoriously costly on FPGAs [143]. The division can be avoided by multiplying all sides of the comparison with $\tilde{E}$. The selection criterion then is:

$$\tilde{E}\tau_{\min} < \tilde{\xi} < \tilde{E}\tau_{\max}, \quad (5.53)$$




Figure 5.5: A pathological sample sequence (circles) and the response of the Maximum Finder with threshold (squares). The sequence is the OF response to two overlapping pulses: one at $t = 0$ ns with an amplitude of 100 % and one at $t = 250$ ns with an amplitude of 10 %. As the samples in the undershoot never exceed zero, the second pulse is not detected by the Maximum Finder.

where $\tau_{\min}$ and $\tau_{\max}$ are freely choosable constants. Setting them to powers of two ensures that even this multiplication is cheap on an FPGA. Like the $|\chi|$ criterion, τ selection introduces no latency and may avoid dead time after a selected BC. And because it is even more efficient to implement than the $|\chi^2|$ criterion, it is foreseen to be used after the Phase-1 Upgrade [114, sec. 5.2.2].

5.6.3 Filtering of In-Time Pileup

In each BC, there are dozens of soft pp collisions that produce low-energy hits in the detector. If the statistical properties of this background are known well enough, one might be interested in removing it on average from the reconstructed energy⁶.

This means that the sequence w is modified to contain a pileup term:

$$w = h * E + h * E_{\text{pu}} + u, \quad (5.54)$$

where E_{pu} is the sequence of the energies of all pileup hits within the same BC and calorimeter cell. It turns out that the effect of this pileup term can be subsumed into the noise autocorrelation in eqs. (5.31) and (5.36):

$$R_{uu}(n) \rightarrow R_{uu}(n) + \sigma_{\text{pu}}^2 (h \star h)(n), \quad (5.55)$$

where σ_{pu}^2 is the variance of the deposited energy per BC due to pileup events. The term $h \star h$ is the unnormalized portion of the autocorrelation R_{hh} , as defined in eq. (5.12).

⁶An exact removal per BC is impossible as hits from soft and hard pp collisions differ only in their energy distribution and frequency of occurrence. In other words, while the effects of *out-of-time* pileup can be suppressed, *in-time* pileup presents an irreducible background.




Figure 5.6: A pathological sample sequence (circles) and the response of the Maximum Finder with threshold (squares). The sequence is the OF response to two overlapping pulses of the same amplitude: one at $t = 0$ ns and one at $t = 375$ ns. The second pulse's energy is underestimated due to the first pulse's undershoot.

One peculiar effect of the bipolar response function h is that in-time pileup can only be reconstructed on average. That is, an attempt to optimize the OF (or any linear filter) to reconstruct the combined sequence $E + E_{\text{pu}}$ simply yields a scaled version of the original OF coefficients:

$$\mathbf{a}_{\text{incl pu}}^\dagger = \frac{\mu_E + \mu_{\text{pu}}}{\mu_E} \mathbf{a}^\dagger. \quad (5.56)$$

The formal reason for this is that $\mathbb{E}[h * E_{\text{pu}}][n] = \sum_k h[k] \cdot \mu_{\text{pu}} = 0$, i.e. the effects of in-time and out-of-time pileup cancel out on average. Another consequence is that using $\mathbf{a}_{\text{incl pu}}^\dagger$ necessarily increases the filter's noise $\mathbb{E}[|\Delta|^2]$.

5.7 Forward Correction

One issue of matched filters – such as the OF – is that their result is only reliable when applied to a BC that is known to contain a cell hit. Before and after the correct BC, they typically produce a long pulse similar to their input. An important consequence of this is that if two hits occur in quick succession, their pulses will overlap. In such a case, matched filters will underestimate the energy of the second hit. This is visualized in fig. 5.6. One possibility to avoid this is to enhance the filter with a *forward correction* [142, sec. 6.3.3, 144, sec. 5.6.3].

Forward correction requires a filter that estimates the deposited energy E for each BC, as well as a selection criterion that estimates whether a hit has occurred at all or not. It operates as a four-stage feedback loop, as shown in fig. 5.7:

1. The filter input $w[n]$ is corrected via element-wise addition of correction values $d[n]$, which have been calculated in step 4 of a previous iteration.
2. The corrected input w_c is passed to the linear filter, which responds with an energy estimate $\tilde{E}$.




Figure 5.7: The forward-correction algorithm. The bold numbers correspond to the steps described in the text.




Figure 5.8: An alternative forward-correction algorithm. In contrast to the version in fig. 5.7, the selection step operates on uncorrected samples.

3. The selection criterion is applied to $\tilde{E}$ to determine whether there was a cell hit in the current BC or not. If there was a hit, the over-all output of the algorithm, $\tilde{E}_{\text{sel}}$ is $\tilde{E}$. Otherwise, $\tilde{E}_{\text{sel}}$ is zero. Depending on the selection criterion, this step may introduce some additional latency.
4. In addition, $\tilde{E}_{\text{sel}}$ is used to estimate the undershoot that the linear filter would produce if there were no further hit. This estimate is written to a buffer to become the correction values $d[n+i]$ from step 1 for all subsequent BCs.

Forward correction has been shown to work well in low-pileup environments [142, sec. 8.5, 144, sec. 6.3.8]. Its usefulness at the HL-LHC still needs to be evaluated. The obvious difficulty with forward correction is that it introduces feedback to the energy reconstruction. One thus loses the guaranteed stability of FIR filters and has to take runaway output into account that quickly saturates the digital electronics.

Furthermore, any implementation of forward correction requires considerable resources on an FPGA: a buffer of approximately L samples must be provided in order to store the correction values d . Similarly, a model that estimates the undershoot must be stored on the FPGA, using either the full filter response $a * h$ or an approximation thereof.

An alternative way to apply forward correction is shown in fig. 5.8. The selection criterion is a regular Wiener filter (i.e. a deconvolution filter) with a threshold applied to it. Crucially, this filter sees the original input w and not the forward-corrected values w_c . If the Wiener filter identifies a hit, an OF is used to estimate its energy. The feedback loop of the forward correction only applied to the OF. The motivation for this is that, by taking the selection criterion out of the feedback loop, the algorithm may become more stable. However, due to the lack of theoretical models for non-linear filters, this hypothesis requires rigorous testing.

5.8 Sparse Signal Restoration

Sparse-representation theory provides an approach to energy reconstruction that is completely different from the more well-known filter-based approach. While its usage for *offline* deconvolution of noisy signals has been researched in the past [145, 146], little research has gone into *online* deconvolution. Part of the reason is that most algorithms are much more complex than linear filters. With computers growing more powerful and deconvolution tasks becoming more difficult, this trade-off is worth being reconsidered.

The theory is stated in terms of a known signal s , a dictionary $\mathbf{H}$, and an unknown sparse representation $\mathbf{x}$, where:

$$\mathbf{s} = (s_i) \in \mathbb{R}^P, \quad \mathbf{H} = (h_{ij}) \in \mathbb{R}^{P \times N}, \quad \mathbf{x} = (x_i) \in \mathbb{R}^N, \quad (5.57)$$

and $P < N$. The N columns of $\mathbf{H}$ are also called *atoms* and the N entries of $\mathbf{x}$ *atom weights*. For energy reconstruction, s are the P most recent samples, $\mathbf{H}$ is the convolution matrix of the system response h , $\mathbf{x}$ are the hit energies to be reconstructed. The atoms are time-shifted slices of h , each P samples long.

Given these quantities, the problem is to find the sparsest $\mathbf{x}$ that explains s :

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{subject to} \quad \mathbf{H}\mathbf{x} = \mathbf{s}, \quad (5.58)$$

where $\|\cdot\|_0$ is the ℓ_0 pseudo-norm that counts the non-zero entries of a vector. The equality can only be satisfied if the problem is noise-free. If s is contaminated with additive noise (as is the case here), the problem can be modified in three ways:

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{subject to} \quad \|\mathbf{H}\mathbf{x} - \mathbf{s}\|_2^2 < \epsilon, \quad (5.59)$$

$$\min_{\mathbf{x}} \|\mathbf{H}\mathbf{x} - \mathbf{s}\|_2^2 \quad \text{subject to} \quad \|\mathbf{x}\|_0 < K, \quad (5.60)$$

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{H}\mathbf{x} - \mathbf{s}\|_2^2 + \lambda \|\mathbf{x}\|_0, \quad (5.61)$$

where the squared ℓ_2 norm $\|\cdot\|_2^2$ is the usual sum of squares. The free parameters ϵ , K , and λ quantify the trade-off between sparsity and accuracy of the representation $\mathbf{x}$. Although the correspondence between these parameters is non-trivial, all three problems are ultimately equivalent.

Various algorithms exist to solve one or another of these problems. The important ones are reviewed briefly in this section, in particular regarding the possibility to use them in the LAr readout electronics.

5.8.1 Greedy Algorithms

An optimization algorithm is greedy if, at each iteration, it proceeds in the way that reduces the residual error the most. Such algorithms are easy to implement, but not necessarily optimal. Most importantly, they can converge to local minima far away from the global minimum.

The simplest among these algorithms is Matching Pursuit [147], shown in listing 5.1. It solves eq. (5.60) via pure *forward search*. This means that it starts out with $\|\mathbf{x}_0\|_0 = 0$ and a

correspondingly large residual error. In each iteration, it selects the *best-matching* atom, i.e. the atom with the greatest projection onto the residual. The weight of this atom is calculated and added to $\mathbf{x}$. Then, the residual is updated and the process is repeated until $\|\mathbf{x}\|_0 = K$. Once an atom i has been selected, its weight x_i is not modified anymore, unless the atom is selected again.

The algorithm is attractive due to its extreme simplicity. It requires only addition, comparison, and multiplication. In particular, the correlation calculation in line 8 of listing 5.1 is amenable to parallel computation.

The algorithm as described here has three disadvantages:

1. It may attribute hits to the wrong BC, especially where there is strong coherence between atoms. Figure 5.9 visualizes this effect for the undershoot of a typical LAr pulse, but this also concerns neighboring atoms close to the peak, as fig. 5.10 demonstrates.
2. It may reconstruct negative energies, resulting in a run-away sequence of ever-increasing weights that overcompensate their predecessors. This can be prevented by skipping the absolute value in line 9 of listing 5.1 and selecting an atom only if its overlap with the residual error is greater than zero.
3. It has to restart from scratch and requires multiple iterations to converge for each BC. For online usage, this is both prohibitively expensive and unnecessary, since there already is a good initial value for each BC: the result of the previous time step with all entries of $\mathbf{x}$ shifted by one place. In order to make use of this initial value, however, Matching Pursuit must be extended with a way to remove old, unimportant atoms.

Listing 5.1: The Matching Pursuit algorithm in the Python programming language, using the NumPy array routines [148]. The operator @ denotes the usual matrix product. $H.T$ is the transpose of H . $H.shape[1]$ is N , the number of atoms.

```

1  from numpy import *
2
3  def matching_pursuit(H, s, K):
4      r = s.copy()           # Residual error of the representation.
5      x = zeros(H.shape[1])  # H.shape[1]  $\equiv N$ .
6      active = []           # len(active)  $\equiv \|\mathbf{x}\|_0$ .
7      while len(active) < K:  # Stopping criterion.
8          a = H.T @ r        # Overlap of each atom with residual.
9          i = argmax(abs(a))  # Select atom with best overlap.
10         active.append(i)    # Increase  $\|\mathbf{x}\|_0$ .
11         x[i] = a[i]         # Weight of the selected atom.
12         r -= a[i] * H[:, i] # Subtract atom from the residual.
13     return x

```




Figure 5.9: A pathological sparse-signal restoration problem. Assume that at $t = 475$ ns, only the samples marked by dashed lines are available to the algorithm. At this point, it is impossible to determine with certainty whether the hit that produced them occurred at $t = 0$ ns, at $t = 75$ ns or at any time in-between.




Figure 5.10: Left: the dictionary H of a typical LAr pulse for $P = 5$. Each atom is a P -elements slice out of the pulse. Right: The mutual coherence of the atoms, i.e. their inner product with each other after normalization. Conventional wisdom is that the solution of sparse signal restoration becomes ambiguous if different atoms are highly coherent with each other. This is trivially the case for atoms in the pulse's undershoot, but also for the atom pairs (1, 2) and (6, 7).

Several extensions to Matching Pursuit exist. The two best-known ones are Orthogonal Matching Pursuit [149, 150] and Orthogonal Least Squares [151]. They differ only in subtle nuances [152]. Both extend the original Matching Pursuit by recalculating the weights of *all* selected atoms in each iteration. This turns the scalar product in line 11 into a matrix multiplication with the Moore–Penrose inverse of a *sub-matrix* of $\mathbf{H}$. Though this calculation can be made FPGA-friendly [153], it is not clear whether the algorithm can satisfy the strict latency requirements in the ATLAS readout.

Another direction in which Matching Pursuit can be extended is the addition of a backwards search step, i.e. allowing the algorithm to remove the least-important atoms. Many algorithms of varying complexity exist, and only two shall be described here.

Subspace Pursuit [154] is based on Orthogonal Matching Pursuit, but adds the K best-matching atoms to the solution set simultaneously, not one by one. Only then, the residual is updated. The algorithm then switches multiple times between adding K more best-matching atoms (for a total of $2K$ atoms) and pruning the K atoms with the *smallest* projection, updating the residual after each step. This may be repeated until some stopping criterion is reached.

Forward–Backward Pursuit [155] improves upon this by introducing an asymmetry between forward and backward steps. It starts out with $\|\mathbf{x}_0\|_0 = 0$ and has two hyperparameters: α and $\beta < \alpha$. In each iteration, it adds the α best-matching atoms and calculates their weights. It then updates the residual and removes the β worst-matching selected atoms. This process is repeated until K atoms have been selected. It has been reported that due to this modification, Forward–Backward Pursuit is more robust w.r.t its hyperparameters than Subspace Pursuit.

Many more algorithms exist and cannot be fully examined here. Numerous variations of Orthogonal Matching Pursuit are named in the presentation of Fusion Forward–Backward Pursuit [156]. Single Most Likely Replacement [157, 158] and Single Best Replacement [159] improve upon Orthogonal Least Squares based on a stochastic analysis of sparse signal restoration. Despite their complexity issues, it seems possible to transfer insights from these algorithms back to the original Matching Pursuit, e.g. avoiding the update of already-calculated atom weights.

5.8.2 Convex Optimization

It is possible to replace the ℓ_0 pseudo-norm in eq. (5.61) with the ℓ_1 norm, a process called *convex relaxation* [160–162]. This turns the problem into an unconstrained convex quadratic optimization problem:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{H}\mathbf{x} - \mathbf{s}\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad (5.62)$$

which can be solved by many general-purpose optimization algorithms. Equation (5.62) is referred to both as Basis Pursuit Denoising [163] and Lasso [164, 165]. Both terms only refer to the problem statement; they do not suggest any algorithm to *perform* the optimization.

Two common algorithms to solve this problem are interior-point methods [166, 167], such as the primal–dual method [168], and iterative-shrinkage algorithms [169]. However, any algorithm for non-linear convex problems may be used, e.g. the BFGS algorithm [170–173]. Some of these algorithms have been implemented on FPGAs [174], so it may be possible that they fulfill the strict latency requirements at ATLAS.

5.9 Artificial Neural Networks

The artificial neural network (ANN) is a machine learning algorithm that has recently seen an immense increase in popularity. “Machine learning” means that, just like the Wiener filter of section 5.3, a generic algorithm is adapted to a concrete problem via a *training* procedure. For the Wiener filter, this is the calculation of its coefficients based on an exemplary input $\mathbf{w}$ and the desired output s . For ANNs, training is the numeric minimization of a non-linear and high-dimensional error function.

Research into *perceptrons*, i.e. artificial neurons in the modern sense, started in 1957 [175]. An introduction to ANNs in general as well as their history until 2015 are given elsewhere [176, ch. 3]. The most notable development since then has been the publication of the TensorFlow library [177] and its adoption in machine-learning libraries like Keras [178]. This made deep neural networks more approachable for beginners and greatly accelerated their adoption both in industry and academics [179].

ATLAS, like other experiments, has adopted ANNs to solve a variety of problems [180]. They are currently being used for track reconstruction [181], particle identification [182], and in the L1 calorimeter trigger [183]. Furthermore, research is going on to extend their usage in the trigger [184, 185] and use them for jet energy calibration [186], calorimeter cell clustering [187], and pileup suppression in the calculation of missing E_T [188]. For energy reconstruction at the calorimeter cell level, investigations have been performed in 2010 [189]. More, independent studies are currently underway [190].

5.9.1 Feed-Forward Neural Networks

The fundamental unit of an ANN is the neuron. A neuron is a mapping $y : \mathbb{R}^n \rightarrow \mathbb{R}$ with $n + 1$ adjustable parameters:

$$y(\mathbf{x}|\mathbf{w}, b) = \varphi(\mathbf{w}^T \mathbf{x} + b), \quad (5.63)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the neuron input. The adjustable parameters $\mathbf{w} \in \mathbb{R}^n$ and $b \in \mathbb{R}$ are the *neuron weights* and the *bias* respectively. For the *activation function* $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, many different choices exist, each useful in its own context. The most common ones are the hyperbolic tangent $\tanh$, the logistic function σ , the identity I , and the rectified linear unit (ReLU) R , where:

$$\sigma(x) := \frac{1}{2} (1 + \tanh x), \quad (5.64)$$

$$I(x) := x, \quad (5.65)$$

$$R(x) := \begin{cases} x & \text{if } x > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (5.66)$$

Although a single neuron can solve linearly separable problems [191], the technique only realizes its full potential by joining them into neural networks. The simplest way to do so is the multi-layer perceptron (MLP) or feed-forward network [175]. It organizes neurons into layers and the output of one layer is the input to the next (see fig. 5.11). Its simplicity stems from the fact that eq. (5.63) can be adapted in a straight-forward manner to describe an entire layer of M neurons with N inputs:

$$\mathbf{y}(\mathbf{x}|\mathbf{W}, \mathbf{b}) = \boldsymbol{\phi}(\mathbf{W}^T \mathbf{x} + \mathbf{b}), \quad (5.67)$$




Figure 5.11: A feed-forward network or MLP. Inputs are received by the input layer, which does no computation of its own. Each subsequent layer performs a linear transformation followed by element-wise application of its activation function. The final layer’s output is the over-all result of the network. This architecture is stateless: the result for one input (x, y, z) does not depend on previous inputs (x', y', z') .

where $W \in \mathbb{R}^{N \times M}$, $b \in \mathbb{R}^M$ and ϕ is the element-wise application of ϕ .

In order to apply ANNs to a problem, their weights and biases must be calibrated via training. Training is the minimization of an error function $e(y, s)$ that, for each input x_i , measures the distance between the desired output s_i and the actual ANN output $y(x_i)$. Many training algorithms exist [192–194], but virtually all are based on *back-propagation* [195], an algorithm that allows to efficiently calculate $\partial e / \partial w_{ij}$.

5.9.2 Time-Series Analyses with Neural Networks

Energy reconstruction can be understood as a pattern-identification problem on time series. There are two architectures of neural networks that are used for such problems: recurrent neural networks (RNNs) and convolutional neural networks (CNNs). For time-series analyses, both can be understood to operate on an input sequence $x : \mathbb{Z} \rightarrow \mathbb{R}$ and produce an output sequence $y : \mathbb{Z} \rightarrow \mathbb{R}$.

In RNNs, connections between neurons are not only “feed-forward” as in MLPs, but may also lead backwards. Hence, a neuron’s output $y[i]$ at time step i depends not only on $x[i]$, but also on its previous output $y[i - 1]$. This allows the ANN to exploit the autocorrelation of x .

Research into RNNs started in the late 80s [196]. For a while, the Elman network [197] was the most common choice for RNNs. Its main disadvantage is that the gradient computed via back-propagation tends to diverge or to go to zero when propagated sufficiently far. The long short-term memory (LSTM) architecture provided a solution to this [198]; by arranging multiple neuron layers into an LSTM unit (see fig. 5.12a) and combining their outputs in a particular manner, the “vanishing-gradient problem” can be circumvented.




Figure 5.12: Typical LSTM and GRU layers. At time step t , the output $y[t]$ depends on the input $x[t]$, the previous output $y[t-1]$, and – for the LSTM – the internal state $c[t-1]$. Round rectangles denote feed-forward layers with the given activation function. Circles denote element-wise operations (addition, multiplication, and $x \rightarrow 1 - x$).

In 2014, the gated recurrent unit (GRU) [199] as a simplification of LSTMs has been proposed (see fig. 5.12b). While it has been proven unable to solve certain problems solvable by LSTMs [200], it offers comparable performance on others [201]. RNNs are a topic of active research and various approaches exist to improve their performance [202, table 4].

CNNs follow a different approach; they don't have any internal state and are better understood as non-linear digital filters. A convolutional neuron is equivalent to an FIR filter composed with an activation function φ . Like linear filters, they model correlations in their input through the weights associated with different time lags; see fig. 5.13 for a visualization. Like in the MLP, convolutional neurons can be arranged into layered networks. Neurons in higher layers are able to pick up higher orders of correlation, improving pattern recognition beyond what linear filters can achieve.

Research of CNNs began in parallel in the context of image classification [205] and time-series classification [206, 207]. In the latter context, they are also called time-delayed neural network (TDNN) and time-lagged feed-forward network (TLFN)⁷. Like RNNs, they are a topic of active research. The technique of *dilation* [208, 209], shown in fig. 5.14, is a recent development that has been shown to dramatically reduce the number of adjustable parameters in deep CNNs without decreasing their performance. Dilated CNNs have been shown to out-perform classical RNNs on several time-series analysis problems [202]. The difference to state-of-the-art RNNs is much smaller and both approaches should be considered competitive.

⁷The term TLFN implies that the first layer is convolutional, while higher layers are fully connected.




Figure 5.13: An exemplary CNN with two hidden layers. Conventionally, the weights of each layer are called *kernel*. The kernel size K corresponds to the filter depth P of FIR filters.




Figure 5.14: A dilated CNN. For clarity, neurons of one layer are overlapped and their scalar output sequences joined into vectors. While all layers have the same kernel size, the distance between samples grows exponentially.

Chapter 6

Simulation of the ATLAS Liquid-Argon Calorimeter Readout Electronics

This chapter describes the setup of the simulation studies that are presented in chapter 7. In particular, section 6.1 introduces AREUS, the simulation program with which the studies of this thesis have been carried out. In sections 6.2 to 6.4, the data used in the simulation and their origin are detailed. Sections 6.5 and 6.6 document other crucial implementation details of the simulation. Finally, the statistical treatment of the results is explained in section 6.7.

For the purpose of reproducibility, the data, AREUS configuration files, and computer scripts that are explained here have been made available online [210]. They are open to all CERN users and contain detailed usage instructions.

6.1 AREUS

AREUS [142, 144, 211] is a program for the simulation of the LAr readout electronics. More concretely, it simulates the readout chain from the analog electronics just outside of the calorimeter up to – and including – the cell-level reconstruction in the back-end electronics.

Such simulation has already previously existed in Athena [212], the common ATLAS offline software¹. However, the Athena simulation is *event-based*: It has a notion of “interesting” events and only simulates a certain time interval around such events (32 BCs before and 32 after it). In contrast, AREUS simulates arbitrarily long time intervals. This allows researching complex and non-linear filters, whose behavior might be correlated at time frames larger than 32 BCs.

AREUS performs discrete-time simulation at various resolutions. Calorimeter hits are simulated with a fixed frequency of 40 MHz, approximately the LHC bunch crossing frequency.

¹In this context, *offline software* means software that runs independent of the LHC clock. Athena is used e.g. for event reconstruction in the high-level trigger, for monitoring of the *online* software, for event simulation, and for data analysis.

The digital electronics are simulated at a sampling rate that is an integer multiple of that. This makes it possible to study the impact of a sampling rate of 80 MHz on the energy resolution of the readout electronics. Finally, the analog electronics are simulated at a fixed sampling rate of 2 GHz ($\Delta t = 0.5$ ns). This fine resolution is necessary to accurately take the hit timing τ into account, due to which pulses might be digitized with a slight phase shift.

The AREUS program is written in a modular manner. The processing steps are formally independent of each other, and they only communicate through dedicated channels. These channels adhere to the *Observer* pattern [213, pp. 293–303]: objects that represent a later processing step observe objects that represent an earlier step and are notified whenever the earlier step produces some output. Strongly typed message objects are used for communication, thus it is guaranteed that each object only observes subjects that it understands.

Figure 6.1 shows a typical setup in which various processing steps are arranged into a chain of subjects and observers. The setup models the LAr readout electronics and is similar to the one used for the simulation studies in this thesis. Concretely, it performs the following processing steps:

1. The *Event Loop* is the root subject and drives the simulation by notifying its observers repeatedly until a given termination condition is met. It is also where the user may specify a *BC pattern*, i.e. a pattern in which the LHC nominal buckets are filled with bunches (see fig. 3.2 and the corresponding text).
2. The *Hit Samples* read hit n-tuple files of detector events from disk and arrange them into a sequence of hits. This process and the format of hit n-tuple files is described in more detail in section 6.2.
3. The *Cell Map* receives the hits and sorts them by calorimeter cell. This produces one hit sequence for each calorimeter cell. At this step, AREUS also applies a *cell position mask* to the hits and discards all hits outside of the mask. This limits the simulation to a certain region of interest and speeds it up significantly.
4. Each *Super Cell* receives the hit sequences of multiple cells and merges them into one sequence. This counter-intuitive split makes the mapping between cells and super cells configurable. AREUS calls this the *position map*.

This flexibility was important for Phase-I Upgrade studies, as the mapping was not finalized at the time. Nonetheless, it is also possible to use this system for Phase-II Upgrade studies, where super cells no longer exist. All that is necessary is to pass a position map to AREUS that defines exactly one “super cell” for each calorimeter cell.

- a) The *Hit Sample Maker* receives the hit sequences of all simulated super cells and writes them into a hit n-tuple file. This way, these intermediary results can be reused in later simulations to save computation time.
- b) The *Test Sequence Maker* does almost the same, but produces its output in a different format. It observes only a single super cell and writes only this cell’s sequence of hits into the output file. This *test sequence* may later be read via the *Test Cell* to reproduce the hit sequence. Such simulations are useful for debugging and if one is interested in the behavior of only a single cell.




Figure 6.1: An example of a processing chain defined in AREUS [142, fig. 7.1]. Objects in light blue are instantiated once per simulation, objects in dark green multiple times. For example, AREUS instantiates one *LAr Filter* object per simulated super cell. Cylinders represent output files that are produced by AREUS.

5. The *Digitization* performs two tasks at once: it simulates the analog electronics and the ADC. The former includes signal shaping, pulse superpositioning, and the addition thermal noise; it is described in sections 6.3 and 6.4. The latter includes the conversion from analog to digital quantities and the addition of quantization noise; it is described in section 6.5.

Up to this point, the processing chain in AREUS is rigid and the steps must be performed in the correct order². After the digitization, by contrast, processing steps can be arranged in a nearly arbitrary manner. This makes it possible to create chains that solve complex problems using only a few configurable building blocks.

6. The *Filters* are a whole class of different processing steps, where each represents a specific filter algorithm. Their common feature is that they all describe transformations from one sequence of digitized samples to another. For example, the OF, the Wiener filter, and the Maximum Finder all are implemented as filters in AREUS. Filters can be run in sequence or in parallel. A *Muxer* makes it possible to combine the output of several parallel filters into one. Section 6.6 describes the filters implemented in AREUS in more detail.
7. At the end of the simulation chain, there are sink objects that only observe subjects but are not observable themselves. AREUS supports two different kinds of sinks:

- a) The *Filter Analysis* records all data produced by the processing chain it observes and performs some simple analyses on it. For example, it records the energy reconstructed by the digital filters for each BC and calculates the *reconstruction error*, i.e. the difference to the true deposited energy. The reconstruction error is presented as a function of the BC, of the deposited energy, as a histogram and more. All data is written to a file at the end of the simulation.

The filter analysis records not only the output of the filter it observes, but of each object in the processing chain, starting at the digitization. This works because whenever a filter object creates a message for its observers, it attaches its own input message as metadata. The filter analysis follows this linked list for each BC.

- b) The *Tree Maker* (called *N-Tuple Maker* in AREUS) collects a more restricted set of data and writes them into an n-tuple file. Among this data are e.g. the true deposited energy, the number of hits, and the reconstructed energy. In contrast to the filter analysis, which produces human-readable output, this is intended as the input to a dedicated external analysis. This is useful for higher-level analyses that e.g. use the simulated output of multiple cells to calculate shower-shape variables. At least one such analysis has been carried out so far [144, sec. 6.2.4].

Figure 6.2 shows an example of several filter algorithms working cooperatively. There are three digitization objects, each with a different least significant bit: High Gain, Medium Gain, and Low Gain. Medium and Low Gain are each observed by one *saturation filter*. A muxer

²Being written in C++, AREUS combines the strong typing of the language and its runtime-type-information capabilities to ensure processing steps are in the correct order. If any object observes a subject it is not compatible with, the program aborts with a corresponding error message.




Figure 6.2: An example of several filter algorithms working cooperatively. A priority is associated with each object and determines which signal is forwarded by a muxer. The muxer only considers the priority of objects it observes directly. The saturation filter differs from most AREUS objects in that it can change their priority dynamically.

observes High Gain as well as the two saturation filters. At each time step, it forwards the output of the subject with the highest priority.

By default, the muxer forwards High Gain output, as it has the highest priority. However, if the Medium Gain output exceeds a certain threshold (e.g. due to a high-energy event), it is assumed that the High Gain has saturated. In this case, the Medium Gain saturation filter increases its own priority for a certain number of time steps. The muxer then switches to Medium Gain in the next time step. If the Low Gain determines that the Medium Gain has saturated as well, the Low Gain increases its priority to make the muxer switch to it.

This way, six AREUS objects, each configured slightly different, can be combined to simulate a complex gain-selection stage. To a subsequent filter, like the OF, the resulting sample sequence is indistinguishable from any other sequence.

6.2 Hit Generation and Sampling

Hit sequences are generated in three independent phases:

- Hit files containing *events* are created. This is done outside of AREUS.
- An AREUS helper script *slims* the hit files into smaller, simpler *hit n-tuple files*.
- The hit n-tuple files are used to produce hit sequences. This is done by AREUS itself.

6.2.1 Generating Hit Files at ATLAS

The fundamental unit of hit files is the *event*, i.e. a single pp collision [136, sec. 3]. An event consists of a *hard-scatter* process and the *Underlying Event*. The former is an inelastic collision between two partons of the interacting protons and is characterized by a large momentum transfer. The latter is the collection of all the soft, i.e. low-momentum interactions that accompany the hard-scatter process. Though the soft interactions are the most numerous at the LHC, it is the hard interactions that are actually interesting.

Event simulation is often done by dedicated Monte-Carlo generators such as Pythia [214]. Their results are written into *event files*³ in the HepMC [215] format. More details on this can be found in the ATLAS Physics Analysis Workbook [216].

The next step is the detector simulation, i.e. a simulation of how the particles emerging from the pp collision interact with the detector material. This is done using Geant4 [217, 218] and results in *hit files*⁴. These hit files describe which calorimeter cells have been hit and how much energy has been deposited by each event. See section 6.2.5 for the hit files that have been used for this thesis.

6.2.2 Converting Hit Files to Hit N-Tuple Files

To make hit files readable by AREUS, they have to be *slimmed* to hit n-tuple files. This means that all data irrelevant to AREUS is discarded and the data format is flattened. A specialized script is used to perform this operation [219]. The hit n-tuple files are defined as follows:

- Each entry corresponds to one event.
- There is one group of related branches for each calorimeter: EMB, EMEC, HEC, FCal, Tile Calorimeter.
- For each calorimeter, there is one scalar branch that contains N , the number of cell hits in the calorimeter for the given event.
- The other branches contain vectors of length N that contain information about each hit in the given calorimeter and event:
 - a numeric ID of the calorimeter layer or module in which the hit occurred,
 - the η and ϕ coordinates of the cell that has been hit,
 - the total deposited energy in MeV,
 - the timing τ of the hit in ps.

The conversion process is complicated by a few incompatibilities and legacy issues that are worth being pointed out separately:

1. The layer/module ID follows the CBNT format [220]. No official documentation about it exists beyond the code that generates and reads it [221–223].

³At ATLAS, event files can be recognized by the tag EVNT in their name.

⁴At ATLAS, hit files can be recognized by the tag HITS in their name.

2. The η and ϕ coordinates are *raw*, i.e. they ignore misalignment, sagging of the detector, and similar effects (cf. fig. 3.8).
3. Geant4 only calculates the *visible* energy, i.e. the energy that a particle deposits in the liquid argon. The total deposited energy is defined as the visible energy divided by the calorimeter's *sampling fraction* [224, 225].
4. The timing τ is defined as the time span between the BC and the hit, *reduced by* the time of flight of a particle moving from the interaction point to the calorimeter cell at the speed of light in vacuum [226, 227]. As such, it is always non-negative.

6.2.3 Generating Hit Sequences

AREUS generates hit sequences in a straight-forward manner. For each simulated BC, the event loop propagates whether the given nominal LHC bucket is filled according to the BC pattern. If it is empty, no hits are simulated for this BC. Otherwise, a number of events are sampled from the hit n-tuple files and each of their hits is propagated to the correct cell based on its coordinates. If a cell is hit multiple times within one BC, AREUS keeps these hits separate and simulates them separately⁵. This is important if both hits vary considerably in their timing and so produce voltage pulses that are shifted with respect to each other.

The way in which AREUS reads the hit n-tuple files is configurable. Any number of files may be declared and each of them may be sampled in one of two modes:

Poisson: Given a positive parameter μ , AREUS samples a random number of events from the file every BC. The number is Poisson-distributed with mean μ .

fixed_distance: Given a positive integer parameter N , AREUS samples one hit from the file every N BCs. Empty buckets due to the BC pattern are not counted and so may increase the distance by an arbitrary amount.

In order to evaluate energy reconstruction algorithms, a workload must be picked. The most common choice is a series of rare, high-energetic hits (the *signal*) on a *background* of extremely frequent low-energy hits.

- The signal hits may come from simulated events with high-energy particles such as $Z \rightarrow ee$ events. If one is only concerned with the performance of energy reconstruction, these hits may also be created arbitrarily. It often makes sense to inject them in *fixed_distance* mode to make their number predictable and to ensure they do not interfere with each other.
- The background hits should be as faithful to the true background at ATLAS as possible. As explained above, it is dominated by soft parton-parton interactions with little momentum transfer. Such interactions are named *minimum-bias interactions* [136, sec. 3] after the corresponding L1 trigger. For this reason, they're typically the *only* background that needs to be considered and their average number per BC can be considered equal to $\langle\mu\rangle$. As such, *Poisson* mode is a good fit for these hits.

⁵This functionality is implemented in the cell map. The objects it receives are of type `CLArHit`, but the objects it propagates are of type `CLArCellHit`. Each `CLArCellHit` is implemented as a collection of `CLArHits`.




Figure 6.3: Fraction of events in an LHC bucket that stays the same when using AREUS’ *random skip* feature. The simulation has been repeated 10 000 times for each data point. The points show the mean value, the error bars the standard deviation.

6.2.4 Sample Looping and Random Skips

A problem when simulating high-pileup scenarios is the limited number of available simulated minimum-bias events. For example, using a hit file with one million events, only 5000 BCs can be simulated with a pileup of $\langle\mu\rangle = 200$. If AREUS exhausts a hit file during simulation, it either terminates simulation or restarts reading from the beginning. When simulating high-pileup scenarios, only the latter is a reasonable choice.

However, reusing the same hit file runs the risk of introducing a long-range periodicity into the hit sequence. Assuming that AREUS exhausts a hit file after simulating N BCs and restarts reading from the beginning, the bucket of BC $N + 1$ will share approximately 50 % of its events with the first BC’s bucket⁶.

One method to avoid this overlap would be to sample events from the hit files in a random order. However, this method is extremely inefficient, as hit files are based on the ROOT file format, which is optimized for sequential reading. Random access would slow the simulation down by several orders of magnitude.

As a compromise, AREUS provides a feature called *random skip*. Whenever it reads an event from a hit file, it first skips a random number of events. The number is uniformly distributed between 0 and a parameter M . Setting $M = 0$ disables the feature. This kind of randomization takes the way ROOT files are read into account while also reducing the similarity of BCs in subsequent loops over the same hit file. This is shown in fig. 6.3.

6.2.5 Hit N-Tuple Files Used in This Thesis

The studies presented in this thesis are based on one file with “signal” hits and two files with “background” hits.

⁶It is not 100 % because the number of events per BC is random. This means it is random at which point in the hit file AREUS starts reading events for BC $N + 1$. Thus, a full overlap isn’t guaranteed.

The “signal” hits are artificial in nature and designed to cover the energy range of the readout electronics uniformly and not to interfere with each other. The signal hits occur at a constant rate of once every 30 BCs in all observed calorimeter cells. The number has been chosen because pulses from the readout electronics have a maximal length of 625 ns, equivalent to 25 BCs. The remainder accounts for the fact that the FIR filters have a typical depth of $P \approx 5$ and might prolong the pulse by as many BCs.

The energy of each hit is random and sampled from a uniform distribution. The minimum of the distribution is 1 μeV , an arbitrary small number that avoids division by zero when calculating some ratios. The maximum of the distribution depends on the calorimeter cell and has been chosen so as to avoid gain switching, since this would introduce non-trivial effects (see section 6.5). It is 59 GeV in the presampler, 7 GeV in the front layer, and 35 GeV in the middle and back layers.

The “background” hits are derived from simulated minimum-bias events. Table 6.1 lists the hit files that have been used. For statistical reasons⁷, the full simulation is split into a simulation of low- p_T interactions and a simulation of high- p_T interactions. When using both samples together, they must be weighted according to their interaction cross-section (79.310 mb for both) multiplied by their filter efficiency (99.652 % for low- p_T , 0.345 32 % for high- p_T) [230, 231]. This can be achieved by multiplying the μ parameter of each hit file’s sampling Poisson distribution (see section 6.2.3).

6.3 Pulse Shapes

AREUS does not simulate the individual parts of the analog electronics. Instead, it treats them as a black box and uses a database that maps cell locations and deposited energies to pulses. The energy dependence of the pulses is taken into account so that saturation effects and non-linearity can be simulated, if necessary. For each hit that occurs, AREUS looks up the corresponding pulse from the database. If there is no entry for an energy, it linearly interpolates between the two entries with the closest energies.

Table 6.1: Samples used to simulate background hits. Both samples have been created in the campaign *MC15*, 14 TeV using Pythia 8. *A2* is the set of tuned generator parameters [228], *MSTW2008LO* is the parton distribution function, and *minbias inelastic* is the physics process. *Low* and *high* refer to the low- p_T and high- p_T approximation respectively. The tags are described on AMI [229].

Dataset	Short physics description	AMI tags
119995	Pythia8_A2MSTW2008LO_minbias_inelastic_low	e1133, s3142, s3144
119996	Pythia8_A2MSTW2008LO_minbias_inelastic_high	e1133, s3142, s3144

⁷The energy distribution of minimum-bias events has a strong peak at low energies and a long tail at high energies. The latter domain is sampled separately to collect a sufficient number of events without drowning it out with the corresponding amount of low-energy events.

Additionally, AREUS uses the pulse database to build a conversion function between deposited energy and induced voltage. This function is piecewise linear and its breakpoints are given by each pulse's associated energy and peak voltage. AREUS uses this function to convert electronic quantities, such as noise, into the physical domain of energies. This calibration effectively removes any non-linearity effects of the analog electronics from quantities that are measured in units of energy.

6.3.1 Pulse Databases Provided by AREUS

By default, AREUS provides three different pulse databases:

analytical: Pulses are described by an analytical model with 8 parameters. The model was originally designed for the trigger readout path in Run 2 [232] and later adapted for Run 3 [144, sec 4.3.4]. As such, it is likely unsuitable for studies of the Phase-II Upgrade. Furthermore, the model neglects saturation and the pulses always scale linearly with the deposited energy.

analytical-readout: This is similar to *analytical*, but uses a model [233] that is tuned to the main readout path of the HEC. Since the analog readout electronics of the HEC are not expected to change in a significant manner⁸, this model can be used for Phase-II Upgrade studies.

spice: This is a collection of numeric pulses that have been calculated via a SPICE simulation [234]. Like *analytical*, this has been originally designed for the trigger readout path of Run 2 [232] and adapted for Run 3 [144, sec. 4.3.4]. In contrast to *analytical*, this takes non-linearity and saturation effects into account.

6.3.2 Pulse Databases Used in This Thesis

For this thesis, two custom databases of numeric pulses have been used. They contain pulses of various amplitudes for six representative cells of the electromagnetic calorimeters. Both databases are based on a simulation of the preamp/shaper chip LAUROC 1 [210, 235], so they are tuned to the main readout path after the Phase-II Upgrade. One simulation covers the pulses at the chip's high-gain output, the other the pulses at its low-gain output. Since there has been no major design change in version 2 (the current prototype), this simulation is expected to still be accurate (but see also section 6.4.2).

Because this thesis also investigates the impact of the shaping stage on the energy resolution, the shaper is not covered by the simulation. Instead, it is approximated with ideal low-pass and high-pass filters (see listing 6.1). These filters ignore parasitic effects and the active amplification in the real electronics, but suffice to make comparisons between different shaper parameters. Figure 6.4 shows examples of the pulses that have been used.

Figure 6.5 shows the linearity of the energy-to-voltage conversion implied by the pulses and their associated energies. Ignoring saturation of the high gain, the conversion is linear to an accuracy of about 1 % in both databases. Since these simulations have been performed, the non-linearity of the LAUROC 1 has been measured and shown to be less than 0.5 % [236].

⁸The FEB2 preshaper is virtually identical to the current one, see section 4.3.4.




Figure 6.4: The ideal pulse shapes used in this thesis, based on a simulation of the LAUROC chip. All pulses have been shaped with an ideal CR-RC² band-pass filter with $\tau = 13$ ns. The first (at $\eta = 2.0$) belongs to an EMEC cell, the other five to cells of the EMB.




Figure 6.5: Linearity of the simulated analog electronics. Shown are the results for a simulated cell in the EMB front layer at $\eta = 0.5$ (points) and the ideal, linear behavior (line). The relative integrated non-linearity (INL) is defined as $(V_{\text{peak}} - Ea)/Ea$, where a the slope of the line. The high-gain electronics saturate for $E \gtrsim 70$ GeV and thus are expected to be strongly non-linear in this domain.

6.4 Thermal Noise

The thermal noise of the analog electronics has been thoroughly investigated in the past. In particular, its variance and autocorrelation coefficients are measured on a regular basis [237]. From this, it is known that the noise closely follows a Gaussian distribution.

There is correlation between the noise of neighboring channels (“coherent noise”) due to common grounding and the proximity to the high-voltage power supply. It has been shown to be smaller than 2 % and is ignored in AREUS.

The thermal noise is also correlated in time. For one, this is because of the band-pass filters employed in the analog electronics, which suppress high-frequency noise. Another reason is that the transistors that are part of the analog electronics contribute *flicker noise*, which is dominant at low frequencies. Figure 6.6 shows typical power spectral densities (PSD) of the thermal noise.

Listing 6.1: The band-pass filter applied to the LAUROC pulses, in the Python programming language. The parameter Δt is the time step of the simulation, τ is the filter’s time constant.

```

1  def ideal_shaper(samples, Δt, τ):
2      α = Δt / Δ(t + τ)          # α is a useful short-hand.
3      samples = _highpass(samples, α)
4      samples = _lowpass(samples, α)
5      samples = _lowpass(samples, α)
6      return samples
7
8  def _lowpass(samples, α):      # The first-order low-pass filter
9      out = []                  # is defined recursively:
10     x = 0.0                    #  $x_n = \alpha s_n + (1 - \alpha)x_{n-1}$ .
11     for sample in samples:
12         x = α*sample + (1-α)*x
13         out.append(x)
14     return out
15
16  def _highpass(samples, α):    # The first-order high-pass filter
17     out = []                  # is defined recursively:
18     prev = 0.0                #  $x_n = \alpha(s_n - s_{n-1}) + \alpha x_{n-1}$ .
19     x = 0.0                    # Note that  $(1 - \alpha)$  does not appear.
20     for sample in samples:
21         x = α*(sample-prev) + α*x
22         prev = sample
23         out.append(x)
24     return out

```




Figure 6.6: Typical thermal noise in the LAr readout electronics. These are the results for the high-gain output of a simulated cell in the EMB front layer at $\eta = 0.5$. The noise PSD, as it would be measured at the preamplifier output (circles) is passed through an ideal CR-RC² band-pass filter with $\tau = 13$ ns (rectangles). It is then folded onto itself to reduce the sampling rate to 40 MHz (up-pointing triangles). The result of this is used to simulate noise with AREUS. The PSD of the simulated noise is averaged over 1000 runs (down-pointing triangles). See page 89 for details on the artifacts in the last curve.

6.4.1 Noise Simulation

AREUS provides three different noise models for simulation. Each model generates noise with a different PSD.

- *White noise* has a flat PSD, shown in fig. 6.7a. It is the simplest form of noise and is easily generated by sampling numbers from a pseudo-random number generator.
- *Pink* or *flicker noise* has a PSD proportional to $1/f$, shown in fig. 6.7b. AREUS generates it via the Voss–McCartney algorithm [238–240, 241, sec. 1.4.4], which is shown in listing 6.2, lines 3 to 12. It is accurate within the band $[f_s/2^{N+1}; f_s/2]$, where N is the order of approximation. Outside of this band, the PSD is flat.
- The *spectral-noise model* generates noise of an arbitrary spectral density. An example of this is shown in fig. 6.7c. The model uses the method of randomized phases [242]. It calculates Fourier coefficients from the PSD and rotates each coefficient by a random complex phase. Applying the inverse DFT to these coefficients produces Gaussian noise. An implementation of this method is given in listing 6.2, lines 14 to 22.

There is not much literature on the problem of generating noise of arbitrary spectral density. As such, the following paragraphs present the method used by AREUS in more detail and compare it to an alternative.

The method of randomized phases has the clear advantage that the generated noise sequence is consistent: its variance is *exactly* what is expected based on the given PSD and

Listing 6.2: Various noise-simulation algorithms in the Python programming language, using the NumPy array routines [148]. The algorithms are implemented as *generators*, i.e. functions that pause execution every time they *yield* a value and may resume execution from these points. Shown are the Voss–McCartney algorithm, the method of randomized phases, and the method of shaped noise. The parameter σ_{tot} is the standard deviation of the generated noise; N is the order of the approximation; S is an array of equidistant points on a PSD; f_s is its sampling rate.

```

1  from numpy import *
2
3  def voss_mccartney( $\sigma_{\text{tot}}$ , N):
4       $\sigma$  =  $\sigma_{\text{tot}}$  /  $\sqrt{N}$                 # Split noise across  $N$  uncorrelated
5      rand = random.normal(0.0,  $\sigma$ , N)    # generators.
6      periods = 1 << arange(N)             # [1,2,4,8,..., $2^{N-1}$ ]
7      pos = periods >> 1                   # [0,1,2,4,..., $2^{N-2}$ ]
8      while True:
9          pos = (pos + 1) % periods         # Update generator  $i$  every
10         mask = (pos == 0)                 #  $\text{pos}[i]$  steps.
11         rand[mask] = random.normal(0,  $\sigma$ , count_nonzero(mask))
12         yield sum(rand)
13
14  def method_of_randomized_phases(S,  $f_s$ ):
15      n_s = (len(S) - 1) * 2              # Number of samples from an irfft call.
16      ps = S * n_s *  $f_s$  / 2             # Use eqs. (5.20) and (5.23) to convert
17      X = sqrt(ps, dtype=complex)         # from PSD to power spectrum.
18      while True:
19           $\phi$  = random.uniform(size=X.shape) # Rotate each amplitude by a
20          X *= exp(2j * pi *  $\phi$ )           # random complex phase in  $[0;2\pi]$ .
21          for sample in fft.irfft(X):
22              yield sample
23
24  def method_of_shaped_noise(S,  $f_s$ ):
25      n_s = (len(S) - 1) * 2
26      while True:
27          u = random.normal(size=n_s)
28          X = fft.rfft(u) * sqrt(S/2)      # The 2 is from eq. (5.23).
29          for sample in fft.irfft(X):
30              yield sample

```




Figure 6.7: PSDs of noise simulated with AREUS (dark blue) and the ideal curves (light orange) at $f_s = 2$ Hz. The simulated PSDs are averaged over 50 independent runs. The pink noise has been simulated to an order of $N = 5$ and hence is accurate down to $f_s/64$ (dashed line). The spectral noise shows *spectral leakage*, an effect caused by discontinuities between batches of simulated noise.

eq. (5.22) (Parseval's theorem). However, an important limitation of this method is that it is based on the Fourier transform and the number of available Fourier coefficients is finite.

Consequently, the inverse DFT only produces a finite number of noise samples. Most implementations circumvent this issue by generating the entire noise sequence once in the beginning, using as many Fourier coefficients as necessary. For AREUS, which may require millions of samples and for which the thermal noise is only one of many things to simulate, this would be prohibitively expensive.

Instead, AREUS generates spectral noise in multiple batches, repeating the phase randomization and inverse DFT every time a batch of noise is exhausted. While this approach solves the computational issue, it introduces discontinuities in the noise sequence. These discontinuities appear between each pair of noise batches and distort the PSD of the noise in two characteristic ways:

- They cause *spectral leakage*, i.e. a broadening of sharp peaks in the original PSD. This effect is visible in fig. 6.7c and can be modeled by convolving the original PSD with an appropriately scaled sinc function.
- They cause *batching artifacts* at $f = 0$ and $f = f_s/2$. These artifacts force the PSD of the simulated noise towards zero, independent of the values of the original PSD at these points. This is because, according to the laws of the DFT, the Fourier coefficients that correspond to these two points should always be real. Multiplying them with a random complex phase causes destructive interference between their contributions to different noise batches⁹. The effect can be observed in fig. 6.6.

⁹It is possible to exempt these two coefficients from phase randomization; however, this would lead to unnat-




Figure 6.8: The difference between PSDs produced via the method of randomized phases and the method of shaped white noise. The simulated PSDs on the left are averaged over $N = 50$ independent runs. The right figure shows that, as N increases, the difference ϵ_{rel} between both methods decreases like $1/\sqrt{N}$. This is compatible with the hypothesis that the differences are statistical in nature. The difference is calculated according to eq. (6.1).

The method of shaped noise [243] is another way to generate noise with an arbitrary PSD. The algorithm is shown in listing 6.2, lines 24 to 30. The procedure is to generate an appropriately sized batch of white noise, transform it into the frequency domain, multiply each Fourier coefficient with a factor proportional to the target PSD, and transform the result back into the time domain.

In contrast to the method of randomized amplitudes, this method produces noise whose variance itself is random. This effect removes the batching artifact at $f = 0$, but the artifact at $f = f_s/2$ remains. Figure 6.8 show that the noise produced with both methods is asymptotically indistinguishable in the frequency domain. In particular, fig. 6.8b shows that the difference between both methods vanishes if averaged over enough runs. The difference is defined as:

$$\epsilon_{\text{rel}} := \sqrt{\int_0^{f_s} df \left(\frac{S_1(f) - S_2(f)}{\sigma^2} \right)^2}, \quad (6.1)$$

where S_1 and S_2 are the PSDs to compare and σ^2 is the noise variance.

Finally, there is a known method to improve spectral-noise simulation by sampling the PSD at random, non-equidistant points [244]. However, this makes it impossible to apply the Fast Fourier Transform algorithm, so it has not been explored further.

6.4.2 Noise Simulation for This Thesis

In this thesis, the spectral-noise model implementing the method of randomized phases has been used. The method of shaped noise has been forgone because it is computationally more

usually coherent artifacts in the generated noise sequence.

expensive (requiring two Fourier transforms per batch instead of one) and both methods produce noise that is sufficiently similar.

The noise PSDs are based on the same simulation as the pulse databases described in section 6.3.2. As for the pulses, it is expected that the simulation describes version 2 of the LAU-ROC as accurately as version 1. However, there is one caveat: it has been shown that the detailed simulation of the electronics underestimates the measured thermal noise by about 20 %. The reason for this is not fully understood at this point, but it is suspected that the simulation of the MIM capacitors used in the preamplifier is too inaccurate [245, pp. 8–12].

6.5 Quantization

The simulation of ADCs in AREUS is implemented in a straight-forward manner. Its input is the pulse that has been acquired from the pulse database, possibly by interpolation (see section 6.3 for details). This pulse is sampled in time and the thermal noise is added (see section 6.4 for details). The pulse is then overlaid with pulses from previous BCs and its amplitude is quantized.

Pulse sampling occurs at a fixed sampling rate f_s . The rate is configurable, but always is a multiple of 40 MHz, the approximate LHC clock rate. The sampling points may be shifted by a constant *static phase* ϕ_{stat} . This does not simulate jitter. However, each pulse may additionally be shifted by a random *dynamic phase* ϕ_{dyn} to simulate variation in the hit timing. This is useful in cases when the hit n-tuple file does not contain useful timing information. Because the hit timing is strictly non-negative (see section 6.2.2), ϕ_{dyn} follows the half-normal distribution [246].

In summary, the sample sequence $s : n \rightarrow s[n]$ is a superposition of the K last pulses $g_k : t \rightarrow g(t)$ and the thermal noise $u : n \rightarrow u[n]$ as follows:

$$s[n] := \sum_{k=0}^{K-1} g_{n-k} \left(\frac{k}{f_s} + \phi_{\text{stat}} + \phi_{\text{dyn}}[n-k] \right) + u[n]. \quad (6.2)$$

The result of sampling is a sequence of time-discrete voltages. Amplitude quantization is done by applying a scaling factor to these voltages and rounding them to a configurable number of bits. Numbers are rounded towards the nearest integer, exact halves are rounded away from zero¹⁰. The scaling factor – which may be supplied in V, MeV or MeV of transverse energy¹¹ – is a measure of how many volts one ADC count corresponds to. Hence, the factor is informally called the *least significant bit*.

6.5.1 Simulation of Quantization Noise

As detailed in section 3.3.5, amplitude quantization adds additional noise to the sample sequence. An ideal ADC with a bit width b has two extreme voltages V_{min} and V_{max} that can be

¹⁰This method has no positive/negative bias and only a slight bias away from zero. This is usually negligible in comparison with the quantization noise.

¹¹If the scaling factor is given in regular or transverse energy, it first needs to be converted to volts. This is done using the piecewise linear conversion function defined in section 6.3.




Figure 6.9: Visualization of ENOB simulation. Shown are the PSDs of a sine wave sampled and quantized in various ways. They correspond to the wave before quantization (circles), after quantization with 12 bits (rectangles), 5 bits (up-pointing triangles), and 12 bits again (down-pointing triangles). For the last curve, the wave has been smeared before quantization to simulate an ENOB of 5.

encoded. With these, the ADC's full-scale range D and resolution Δ are defined as:

$$D = V_{\max} - V_{\min}, \quad (6.3)$$

$$\Delta = \frac{D}{2^b}. \quad (6.4)$$

Quantizing a sample introduces an error e_b . The error is uniformly distributed in $\left[-\frac{\Delta}{2}; \frac{\Delta}{2}\right]$ and so has a variance

$$\sigma_{\text{ideal}}^2 = \frac{\Delta^2}{12} = 2^{-2b} \frac{D^2}{12}. \quad (6.5)$$

A real ADC consists of electronics that are themselves noisy. Therefore, its noise σ_{real} is always greater than σ_{ideal} . To quantify this, b_{eff} , the effective number of bits (ENOB) is defined such that:

$$\sigma_{\text{real}}^2 = 2^{-2b_{\text{eff}}} \frac{D^2}{12} = 2^{2(b-b_{\text{eff}})} \frac{\Delta^2}{12}. \quad (6.6)$$

AREUS simulates this effect by adding Gaussian-distributed noise to the samples before quantizing them [144, sec 5.1]. The magnitude of this noise is chosen so as to bridge the gap between σ_{ideal} and σ_{real} :

$$\sigma_{\text{smear}}^2 = \sigma_{\text{real}}^2 - \sigma_{\text{ideal}}^2 = (4^{b-b_{\text{eff}}} - 1) \frac{\Delta^2}{12}. \quad (6.7)$$

An example of this simulation is shown in fig. 6.9. A sine wave is quantized once with an ideal ADC ($b = 5$) and once with a real ADC ($b_{\text{eff}} = 5$, $b = 12$). The resulting noise has the same magnitude and PSD (though the exact distribution differs).

6.5.2 ADC Model Used in This Thesis

The readout electronics installed during the Phase-II Upgrade will have to cover a dynamic range of 16 bits. It is planned to do this with two overlapping gain stages: A high gain ($\times 30$) and a low gain ($\times 1$). Each stage is digitized separately with a 14-bit ADC. This ADC in turn is planned to consist of two 12-bit ADCs with dynamic switching between [111, sec. 3.7].

For the purpose of simulation, this readout chain can be approximated as four distinct gain stages, where two stages correspond to the low gain and two to the high gain. This has been summarized in table 6.2, together with the value of the least significant bit that has been used in each stage and detector region. The values have been calculated using the formula:

$$\text{LSB} = \frac{I_{\max}}{k_{\min} \cdot 2^b \cdot G \cdot s_{\min}}, \quad (6.8)$$

where $I_{\max}$ is the maximum current allowed by the analog electronics; $k_{\min} := \min_{\eta} k(\eta)$ is the minimum of the energy-to-current conversion factor across each layer; $b := 12$ is the bit width of the simulated ADC; $G \in [1, 4, 30, 120]$ is the gain of the given stage. The factor $s_{\min} \approx 0.75$ is the positive fraction of the dynamic range; the remainder is reserved to represent the negative undershoot of the bipolar signal pulse without saturation.

For lack of a more appropriate value, it is assumed that the ADCs installed in the Phase-II Upgrade will have an ENOB of 11.3 bits. This is the middle between the minimum required by the specification (11 bits) and the value measured on a prototype ADC (11.6 bits) and compatible with the current electronics [111, sec. 7.1.2].

6.6 Digital Filters

The filters implemented in AREUS are generally as described in chapter 5. However, in a few cases, they either provide some additional features or only a subset of what is possible. These deviations are documented here.

One important point is that the comparative studies presented in this thesis are focused on the main readout. In this context, it may be assumed that the BCs in which a hit has occurred have already been identified. This means that the sole purpose of the digital filters is to reconstruct the energy deposited in these hits as accurately as possible.

Table 6.2: Least significant bits in the three layers of the electromagnetic calorimeter for all four simulated gain stages. Equation (6.8) has been used to calculate them based on $I_{\max}$ and $k_{\min}$ [111, tab. 3.2f].

Layer	$I_{\max}$ (mA)	$k_{\min}$ (mA/TeV)	Least significant bit (MeV)			
			HG $\times 4$	HG $\times 1$	LG $\times 4$	LG $\times 1$
Presampler	2	0.28	19.4	77.5	581	2325
Front	2	2.26	2.4	9.6	72	288
Middle	10	2.26	12.0	48.0	360	1440
Back	10	2.33	11.6	46.6	349	1397

This focus not only limits the applicability of the conclusions drawn here, but it also has an impact on the choice of filter algorithms to compare. For example, there is no point in investigating the performance of a naïve Wiener filter as defined in section 5.3, since it is explicitly *not* optimized for the problem of energy reconstruction.

6.6.1 Optimal Filter

The OF is implemented in AREUS exactly as described in section 5.6. However, it should be noted that the numbering of filter coefficients follows the convention of eq. (5.27) and *not* of eq. (5.34). This means that the sequence of OF coefficients a is proportional to $h^\dagger$, the *reversed* sequence of ideal pulse samples h . This is opposite to the ordering commonly used in particle physics [141].

Because the BCs with cell hits have already been identified, no selection criterion must be applied. Hence, the filter coefficients $b^\dagger$ are ignored here. Nonetheless, AREUS also provides variants of the base OF algorithm which use b to estimate ξ and e.g. apply τ selection.

6.6.2 Wiener Filter

The Wiener filter is implemented in AREUS according to eq. (5.35):

$$\begin{pmatrix} c[P-1] \\ c[P-2] \\ \vdots \\ c[0] \end{pmatrix} = \begin{pmatrix} R_{ww}[0] & R_{ww}[1] & \cdots & R_{ww}[P-1] \\ R_{ww}[-1] & R_{ww}[0] & \cdots & R_{ww}[P-2] \\ \vdots & \vdots & \ddots & \vdots \\ R_{ww}[1-P] & R_{ww}[2-P] & \cdots & R_{ww}[0] \end{pmatrix}^{-1} \begin{pmatrix} R_{ws}[0] \\ R_{ws}[-1] \\ \vdots \\ R_{ws}[1-P] \end{pmatrix}, \quad (6.9)$$

which, since all involved quantities are real, can be simplified to:

$$\begin{pmatrix} c[0] \\ c[1] \\ \vdots \\ c[P-1] \end{pmatrix} = \begin{pmatrix} R_{ww}[0] & R_{ww}[1] & \cdots & R_{ww}[P-1] \\ R_{ww}[1] & R_{ww}[0] & \cdots & R_{ww}[P-2] \\ \vdots & \vdots & \ddots & \vdots \\ R_{ww}[P-1] & R_{ww}[P-2] & \cdots & R_{ww}[0] \end{pmatrix}^{-1} \begin{pmatrix} R_{ws}[0] \\ R_{ws}[1] \\ \vdots \\ R_{ws}[P-1] \end{pmatrix}. \quad (6.10)$$

Because AREUS has access to the true deposited energy, it can calculate R_{ww} and R_{ws} during simulation.

In a deviation from the description in section 5.3, the Wiener filter in AREUS is *unbiased*. This is implemented via a renormalization of the filter coefficients c :

$$c' = c \cdot \frac{\max_n h[n]}{\max_n (c * h)[n]}, \quad (6.11)$$

where h is the normalized system response of the readout electronics. Consequently, if a normalized pulse is passed through the Wiener filter, the response is normalized as well. For this reason, AREUS's Wiener filter can be used for energy reconstruction.

Additionally, AREUS supports both peak broadening and narrowing as described in section 5.3.3. Peak *narrowing*, unless disabled, is done automatically during filter calibration. After calculating the coefficients c , AREUS calculates the ideal filter response $g := c' * h$ to and locates its peak sample $g[i_{\max}]$. It then verifies that $g[i_{\max}-1]$ and $g[i_{\max}+1]$ are below a certain

threshold, zero by default. If not, it decreases the desired pre-peak and post-peak samples, s_{pre} and s_{post} , and restarts the calibration process.

Peak *broadening*, in contrast, is implemented by setting both the threshold on $g[i_{\text{max}} - 1]$ and $g[i_{\text{max}} + 1]$ as well as the initial values for s_{pre} and s_{post} to 0.5. This can be enabled for the pre-peak sample and the post-peak sample separately. (Note that during calibration, the automatic peak-narrowing process might reduce the initial values again.)

6.6.3 Wiener Filter with Forward Correction

Due to its complexity, forward correction is not implemented as a mere add-on to other filters in AREUS, but as a filter of its own. In particular, it is tightly coupled with the Wiener filter, as the Wiener filter is well suited for detecting whether a hit has occurred.

The Wiener filter with forward correction (WFFC) is based on the algorithm visualized in fig. 5.7. By default, it uses a peak-broadened Wiener filter with $s_{\text{pre}} = 0$ and $s_{\text{post}} = 0.5$. The selection algorithm is:

$$\tilde{E}_{\text{sel}}[n] = \begin{cases} \tilde{E}[n-1] & \text{if } \frac{\tilde{E}[n]}{\tilde{E}[n-1]} > \frac{g[i_{\text{max}} + 1]}{g[i_{\text{max}}]} - m, \\ 0 & \text{otherwise,} \end{cases} \quad (6.12)$$

where $\tilde{E}$ is the Wiener filter output sequence, g is defined as above, and m is a variable margin that allows to relax the selection criterion as necessary. The algorithm triggers for the previous sample whenever the current sample is greater than the post-peak sample one would expect of a signal pulse.

In this thesis, a more restrictive selection algorithm is used. It uses a peak-broadened Wiener filter with $s_{\text{pre}} = s_{\text{post}} = 0.5$ and its criterion is

$$\frac{\tilde{E}[n-2]}{\tilde{E}[n-1]} > \frac{g[i_{\text{max}} - 1]}{g[i_{\text{max}}]} - m \quad \wedge \quad \frac{\tilde{E}[n]}{\tilde{E}[n-1]} > \frac{g[i_{\text{max}} + 1]}{g[i_{\text{max}}]} - m, \quad (6.13)$$

i.e. it additionally applies an analogous threshold to the sample before the previous one. Analogous criteria exist for the other two combinations of pre-peak and post-peak samples. If no peak broadening is used at all, the peak candidate must only be greater than zero.

The forward correction is as described in the original publication [142, p. 77, 144, p. 102]:

$$d[n+i] = -\frac{\tilde{E}[n-1]}{g[i_{\text{max}}]} \cdot \begin{cases} g[i_{\text{max}} + i + 1] & \text{if } i \in [1; P_1 + 1), \\ \bar{g}_- & \text{if } i \in [P_1 + 1; N - P_2), \\ g[i_{\text{max}} + i + 1] & \text{if } i \in [N - P_2; N), \end{cases} \quad (6.14)$$

where N is the number of samples in the undershoot of g . It is determined by finding the last sample in g that is below a configurable threshold m_{last} , -0.01 by default. P_1 and P_2 are the depths of two FIR filters that reproduce sections of the undershoot, $\bar{g}_-$ is an estimate of the undershoot's flat section:

$$\bar{g}_- := \frac{1}{N - P_2 - P_1 - 1} \sum_{k=P_1+1}^{N-P_2} g[i_{\text{max}} + k + 1]. \quad (6.15)$$

In addition to this forward correction, the *current* sample is also corrected:

$$\tilde{E}[n] \rightarrow \tilde{E}[n] - \tilde{E}[n-1] \cdot \frac{g[i_{\max} + 1]}{g[i_{\max}]}. \quad (6.16)$$

This allows the WFFC to trigger correctly in multiple subsequent BCs.

The WFFC in AREUS is equipped with two more features to improve its run-time behavior:

- A *swing-up protection* checks in each iteration n whether the absolute forward correction $|d[n]|$ is above a certain threshold. If this is true for a number of consecutive samples, the filter is reset. The motivation for this is that a large correction value indicates that the filter has diverged due to its non-linear behavior. In such a case, it is unable to return to finite values without outside intervention.
- A *saturation lock* checks for each iteration n whether the incoming sample $w[n]$ is above a certain threshold. If it is, all samples are blocked from entering the algorithm and replaced by zeros for a number of samples. The reason for this is to avoid processing saturated pulses; due to their irregular shape, the Wiener filter would likely give a response that does not match the selection criterion. It would hence fail to correct for the pulses, which would bias subsequent results.

6.6.4 Other Algorithms

Due to time constraints, it has not been possible to implement ANNs or algorithms for sparse signal restoration in AREUS. Both methods have been studied on small, artificial problems [247, 248], but a comprehensive optimization and comparison with linear filters has yet to be done.

6.7 Statistical Analysis

The data produced by the AREUS simulation must be interpreted statistically. For the task of energy reconstruction, the most important figure of merit is the *reconstruction error* Δ , i.e. the difference between the reconstructed energy E_{reco} and the true deposited energy E_{true} for a given BC. Across all BCs, this error can be regarded as a random variable. It is Gaussian-distributed for low luminosities, where electronic noise dominates; as the luminosity and $\langle\mu\rangle$ increase, it becomes significantly non-Gaussian and the effect of outliers becomes significant. This is shown in fig. 6.10.

Although the mean $\mathbb{E}[\cdot]$ and the standard deviation $\sigma[\cdot]$ are often used to describe unimodal distributions like that of Δ , these measures are highly susceptible to outliers. More robust measures are the median Q_2 and the interquartile range (IQR). They are defined as:

$$Q_2 := F^{-1}(0.5), \quad (6.17)$$

$$\text{IQR} := F^{-1}(0.75) - F^{-1}(0.25), \quad (6.18)$$

where F is the cumulative distribution function of the random variable (here Δ) and F^{-1} is its generalized inverse, the *quantile function*.




Figure 6.10: Typical frequency distributions of the inclusive (Δ_{tot} , circles) and exclusive (Δ_{sig} , squares) reconstruction error at a luminosity corresponding to $\langle\mu\rangle = 200$. The non-Gaussian tails of the error distribution are clearly visible. The distribution has been obtained by simulating 50 000 signal pulses in a cell in the EMEC middle layer at $\eta = 2.0$.

For a Gaussian distribution $N(\mu, \sigma)$, it is $Q_2 = \mu$ and $\text{IQR} \approx 1.35\sigma$. The latter is somewhat undesirable – ideally, the measure of spread would be comparable to σ . For this reason, the IQR is replaced with a bespoke quantity, the *one-sigma half range*:

$$\varsigma := \frac{F^{-1}(0.84135) - F^{-1}(0.15865)}{2}, \quad (6.19)$$

i.e. half the width of the interval that contains 68.27 % of the data. It has the same statistical properties as the IQR, but it is $\varsigma \approx \sigma$ for Gaussian-distributed variables.

There are two variants of the reconstruction error Δ that may be considered as a measure of energy resolution:

1. the *inclusive* error $\Delta_{\text{tot}} := E_{\text{reco}} - E_{\text{true, sig}} - E_{\text{true, pu}}$ compares the filter's reconstructed energy to the total deposited energy, including in-time pileup;
2. the *exclusive* error $\Delta_{\text{sig}} := E_{\text{reco}} - E_{\text{true, sig}}$, by contrast, only considers energy depositions due to the hard-scatter process.

The advantage of Δ_{tot} is that it is intuitively correct: when a digital filter reconstructs energy based on a signal pulse, it is impossible to tell how much of it is due to soft and how much due to hard pp collisions. Especially in the rare case of a high-energy hit from the Underlying Event, Δ_{tot} is expected to stay small, as it takes the hit's energy into account. This is not the case when using Δ_{sig} .

However, though there is no difference between hits from soft and hard collisions on an individual basis, they do differ in their statistics. While for the hard-scatter process, signal pulses rarely overlap, the same is almost guaranteed for pulses from soft collisions. Because the signal pulse approximately integrates to zero, pileup hits cancel out on average. As described in section 5.6.3, matched filters without additional calibration reconstruct Δ_{sig} on average, not Δ_{tot} .

Finally, as the focus of the studies presented here lies on the main readout, it is assumed that bunch crossing identification has already happened. In other words, it is known for each BC whether a hit from a hard-scatter process has occurred or not (see section 6.6).

Hence, although the filters are applied to each BC, their output is only considered in the analysis if it corresponds to a non-zero signal energy deposition. In particular, the filters are allowed to produce false positives, i.e. non-zero output in BCs without energy deposition due to a hard collision. This allows for broad filter responses, as they are typical for e.g. the OF.

Chapter 7

Results of the Readout Electronics Simulation Studies

In this chapter, comparative simulation studies of the ATLAS LAr readout electronics are presented. The overarching goal is to optimize their energy resolution for the Phase-II Upgrade. To do so, parameters of both the analog and the digital electronics are varied and their impact on the energy resolution is inspected. For a detailed description of the simulation setup, refer to chapter 6.

First, a statistical analysis of the simulation setup is presented in section 7.1. A figure of merit representative of the energy resolution is picked for the subsequent comparisons, and its uncertainty is estimated.

In section 7.2, the electronic-noise simulation of AREUS is compared to data measured during Run 2. This serves to verify the accuracy of the simulation. In a similar vein, the dependence of the energy resolution on the time constant τ of the shaper is investigated in section 7.3. This is both instructive and verifies that the simulation works as intended when including pileup. The influence of τ – and several other parameters of the analog electronics – is analyzed in more detail in section 7.4 and they are tuned for optimal performance.

In the following sections 7.5 to 7.7, three filter algorithms are chosen (the Optimal Filter, the Wiener filter, and the Wiener filter with forward correction) and individually tuned for optimal energy resolution. In section 7.8, these filters are compared with each other. Their behavior is investigated and conclusions drawn from the study are presented.

7.1 Statistical Treatment

The purpose of this section is to investigate the figures of merit introduced in section 6.7 and to select one for the subsequent analyses. To do so, the same simulation is executed multiple times, using different seeds for the pseudo-random number generator each time. For each run, statistical quantities are calculated for the reconstruction error Δ . These figures of merit fluctuate statistically themselves; this section also calculates and discusses these meta-statistics.

The simulation follows the description in chapter 6. While it has been run for all six reference calorimeter cells chosen for this thesis, results are presented for the cell in the EMEC




Figure 7.1: The sample mean of the exclusive (Δ_{sig} , circles) and inclusive (Δ_{tot} , squares) reconstruction error, and the median for both (up- and down-pointing triangles resp.), after application of the OF, for five different pileup scenarios. The results are shown with a horizontal shift to avoid overlap. For the mean, error bars indicate the sample standard deviation; for the median, they indicate the one-sigma interval, i.e. the interval that contains the most central 68.27 % of the data. Simulation target is a cell in the EMEC middle layer at $\eta = 2.0$.

middle layer at $\eta = 2.0$, except where indicated otherwise. The analog shaper is an ideal CR-RC² band-pass filter with $\tau = 13$ ns. The digital filter for energy reconstruction is an OF with a depth of $P = 5$. Each run simulates 50 000 signal pulses at a constant distance of 30 BCs, for a total of 1 500 000 simulated BCs. Of these, 40 000 BCs are used to calibrate the OF, the rest for evaluation. (See the end of this section for a motivation of this number.)

The simulation has been run 250 times, 50 times for each of five pileup scenarios:

- $\langle\mu\rangle = 0$: no pileup noise, only electronic (thermal and quantization) noise;
- $\langle\mu\rangle = 20$: about as much pileup as in Run 2;
- $\langle\mu\rangle = 80$: the maximum amount of pileup expected in Run 3;
- $\langle\mu\rangle = 140$: the amount of pileup expected in Run 4, after the Phase-II Upgrade;
- $\langle\mu\rangle = 200$: the maximum amount of pileup expected in the HL-LHC era.

The inclusive and exclusive reconstruction errors, Δ_{tot} and Δ_{sig} , have been calculated for each signal pulse after application of the OF. Their sample mean $\bar{\Delta}$, sample standard deviation $s[\Delta]$, median $Q_2[\Delta]$, and one-sigma half range $\varsigma[\Delta]$ have been calculated for each simulation run separately (see section 6.7 for their definition). They have then been averaged across 50 simulations for each pileup scenario. The results are shown in fig. 7.1.

For $\langle\mu\rangle = 0$, the four data points show no differences. This is to be expected, as the electronic noise is the only source of error, i.e. $\Delta_{\text{sig}} = \Delta_{\text{tot}}$. Electronic noise is approximately Gaussian and in this case, the one-sigma half range matches the standard deviation by construction.




Figure 7.2: Comparison of the sample standard deviation (s , filled markers) and the one-sigma half range (ς , empty markers) of the inclusive reconstruction error Δ_{tot} . The left shows frequency distributions of both quantities for five pileup scenarios with 50 simulations per scenario; The right shows each distribution's coefficient of variation. Simulation target is a cell in the EMEC middle layer at $\eta = 2.0$.

As the error bars in fig. 7.1 show, the measures of noise increase with $\langle \mu \rangle$. Two effects can be seen: One is that ς is smaller than s in all cases, underlining its robustness against non-Gaussian tails. The other is that Δ_{sig} fluctuates considerably more than Δ_{tot} . This can be explained by the fact that linear filters cannot distinguish between signal and pileup hits, and so inherently reconstruct E_{tot} , the energy of both (see section 6.7). For this reason, the variance $s^2[\Delta_{\text{sig}}]$ is effectively the sum of $s^2[\Delta_{\text{tot}}]$ and $s^2[E_{\text{true,pu}}]$, the energy deposited due to pileup.

While this is an argument in favor of using Δ_{tot} , the mean shows another effect to be aware of. While $\overline{\Delta_{\text{sig}}}$ stays close to zero, independent of $\langle \mu \rangle$, $\overline{\Delta_{\text{tot}}}$ shows a trend towards the negative with increasing $\langle \mu \rangle$. The reason for this ultimately lies in the bipolar shape of the signal pulse. By construction, its positive lobe and negative undershoot integrate to zero. Hence, even though linear filters reconstruct Δ_{tot} rather than Δ_{sig} , the pileup hits cancel each other out on average and Δ_{sig} is the quantity that remains unbiased¹.

Ultimately, it has been decided to use Δ_{tot} for the following analyses. The reason is that it more closely describes the behavior of the OF in absence of statistical effects; the pileup-dependent bias of $\overline{\Delta_{\text{tot}}}$ is largely independent of the filter (see section 7.8) and is corrected in later stages of the readout chain anyway [111, sec. 2.2.3].

To decide between the sample standard deviation and the one-sigma half range, $s[\Delta_{\text{tot}}]$ and $\varsigma[\Delta_{\text{tot}}]$ resp., it has been investigated by how much their values vary between simulation runs. This is visualized in fig. 7.2. It shows both the actual frequency distributions and each

¹Curiously, the median $Q_2[\Delta_{\text{sig}}]$ also shows a negative bias, though it is much smaller. The cause for this is the asymmetrical distribution of Δ_{sig} : a slight shift of the maximum towards the negative is balanced out by a long, non-Gaussian tail towards the positive. The median is ignorant of this tail and thus makes the shift visible.




Figure 7.3: Coefficients of variation of $\varsigma[\Delta_{\text{tot}}]$ for the Optimal Filter (OF), the Wiener filter (WF), and the Wiener filter with forward correction (WFFC) as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing.

distribution's *coefficient of variation*. The latter is defined as:

$$c_V[\Delta] := \frac{s[\Delta]}{\Delta}, \quad (7.1)$$

i.e. the distribution's standard deviation divided by its mean.

As can be seen, c_V is of the same order of magnitude for both $s[\Delta_{\text{tot}}]$ and $\varsigma[\Delta_{\text{tot}}]$. However, while it increases with $\langle \mu \rangle$ for $s[\Delta_{\text{tot}}]$, it decreases for $\varsigma[\Delta_{\text{tot}}]$. Furthermore, fig. 7.2a shows several outliers in the distributions of $s[\Delta_{\text{tot}}]$, e.g. near $\Delta_{\text{tot}} = 200$ MeV. As a robust estimator, $\varsigma[\Delta_{\text{tot}}]$ does not suffer from this problem. For this reason, it is used for the subsequent analyses, except where noted otherwise.

The above procedure has been repeated for a baseline configuration² of the Wiener filter and the Wiener filter with forward correction (WFFC). The values of c_V for all three filters are presented in fig. 7.3. They are used to approximate the uncertainty of $\varsigma[\Delta_{\text{tot}}]$ for all subsequent analyses. The underlying assumption is that, while modifications of the readout electronics may change the energy resolution, they do not significantly change its uncertainty.

Finally, the training size, N , has been investigated. This parameter of AREUS determines for how many BCs it collects data to estimate the noise autocorrelation before it calibrates the OF. The simulation has been run for several different sizes and repeated five times for each one. The purpose is to also estimate roughly how the training size impacts the variance of the energy resolution. The results are shown in fig. 7.4.

²The filter depth is $P = 6$, the pre- and post-peak samples are $s_{\text{pre}} = s_{\text{post}} = 0.5$. Automatic peak narrowing is employed for both filters. They are calibrated according to eq. (5.36), using the ideal pulse shape and ignoring electronic noise. The WFFC selection criterion is eq. (6.15), its margin parameter is $m = 0.25$. The corrective FIR filters have a depth of $P_1 = 5$ and $P_2 = 3$ resp. The threshold determining the end of the undershoot is $m_{\text{last}} = -0.015$.




Figure 7.4: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different training sizes. The points are shifted horizontally to avoid overlap. The bottom figure shows the relative difference between each scenario and the baseline solution.

All three variants have the same energy resolution within their respective uncertainty. This is expected, as they only differ in the accuracy of their estimate of the noise autocorrelation. There is one outlier in the EMB at $\langle \mu \rangle = 20$, but it is likely a statistical fluctuation.

It is also apparent that the uncertainty of the energy resolution decreases with increasing training size. The reason for this is that the noise autocorrelation matrix generally is ill-conditioned, meaning that it must be estimated with high precision to avoid overly large numerical errors when inverting it. A training size of 40 000 BCs has been chosen for the baseline simulation, as it keeps the uncertainty below 1 % in all scenarios. Because resolution improvements smaller than 1 % are often considered not worth the engineering effort associated with them, this limit is good enough for the comparisons presented here.

7.2 Simulation Verification Using Run-2 Data

In this section, the simulation of electronic noise is verified by comparing it to data that has been measured during Run 2. This has been done by configuring AREUS with data from a calibration run and simulating a run without pileup, i.e. at $\langle \mu \rangle = 0$. The noise after OF application is then compared to the corresponding noise measured at ATLAS.

The data have been taken from the calibration runs 322 719, 322 723, and 322 724 of the high-gain readout, which have been recorded on May 10, 2017 [249]. They comprise:

- σ_{elec} , the electronic noise standard deviation before OF application in ADC counts;
- g_i , the normalized signal pulse shape of the calorimeter electrodes;




Figure 7.5: Electronic noise in six reference calorimeter cells. The left figure shows data from the calibration run (circles) and the table used for conversion (squares). The right figure shows the standard deviation of measured (circles) and simulated (squares) noise.

- $\rho(n)$, the autocorrelation coefficients.

The latter are related to the autocorrelation proper via:

$$\rho(n) := \frac{R(n) - \mu_{\text{elec}}^2}{\sigma_{\text{elec}}^2}. \quad (7.2)$$

See section 5.1 for details on the notation. Because trivially $\rho(0) = 1$, the first coefficient is not actually stored in the calibration files. Furthermore, the noise is unbiased, i.e. $\mu_{\text{elec}} = 0$.

The pulse shapes have a length of 800 ns and a resolution of approx. 1 ns. Because they are normalized, it is impossible to determine the energy–voltage conversion function that they imply (see section 6.3). To circumvent this issue, this analysis compares the noise in units of energy. This way, the conversion from *deposited* energy to voltage and the conversion from voltage to *reconstructed* energy cancel out. The conversion function can thus be chosen arbitrarily, e.g. $V(E) = E \cdot 1 \text{ mV/MeV}$.

The noise PSD, $S(f)$, can be reconstructed from $\rho(n)$ via the Wiener–Khinchin theorem, eq. (5.21). To do so, $\rho(n)$ must first be converted back to the autocorrelation $R(n)$ via eq. (7.2). This requires σ_{elec} to be converted from ADC counts to volts, which is done in two steps:

1. It is first converted to MeV using a reference table [250]. Figure 7.5a shows that the values in the table agree well with the calibration data.
2. It is then converted to volts via the above conversion function $V(E)$.

One problem of this approach is that calibration data only contains $\rho(n)$ for $n < 32$. This causes $S(f)$ to take on complex values in a manifestation of Gibbs' phenomenon [251, 252]. Though it introduces a small numeric error, the solution to this is to take the absolute value $|S(f)|$.

With this data, an AREUS simulation has been run. It has been configured to use an ideal 12-bit ADC with a least significant bit of 10 MeV. This value [253, sec. 2] has been chosen because the calibration data has been taken for the high-gain readout. The ADC is assumed to be ideal because quantization noise has already been absorbed by σ_{elec} .

The OF has been configured to a depth of $P = 5$ and calibrated on 10 000 simulated BCs. It has then been applied to 10 000 more BCs and the electronic noise σ_{OF} has been calculated as the sample standard deviation of the OF output sequence. The uncertainty of σ_{OF} due to the pseudo-random number generator has been estimated by repeating the process 50 times. It is $\lesssim 0.05$ MeV and negligible in comparison to the approximations taken for this study.

The reference noise measurements have been taken from the ATLAS conditions database using the *CaloNoiseTool* [254, sec. 1.3] and the conditions tag *OFLCOND-MC15c-SDR-14-03*. These measurements are taken regularly during pedestal runs [111, sec. 6.1.1]. The results of the comparison are presented in fig. 7.5b.

They show that the simulation is accurate to a precision of 1 % in the front layer and 5 % in the middle layer. In the presampler, a large deviation by about 16 % is visible. The exact reason for this is not clear; the most likely cause is the reduced high voltage employed in the presampler since 2012 [255, 256]. Another reason may be that due to its noise autocorrelation, the presampler's σ_{OF} is more sensitive to the above approximations than in the other layers.

7.3 Dependence of the Noise on the Shaping Time

In this section, the claims made in section 3.3.5 are verified empirically – in particular the statement that the electronic noise is suppressed best if τ is small and pileup noise if τ is large. Another goal is to qualitatively verify fig. 3.13, which shows the dependence of both noise sources on the signal pulse peaking time.

This has been achieved in multiple steps: First, the relationship between the pulse peaking time t_p (see below for a definition) and τ has been examined. In a second step, the readout electronics have been simulated with only electronic noise enabled and the noise after OF application is calculated; this is expected to decrease as t_p increases. The third step has been to disable thermal noise (but not quantization noise³) and simulate only pileup noise; this is expected to increase with t_p . In a fourth and final step, the full noise simulation (electronic and pileup) has been run. The resulting noise is expected to be approximately equal to the square sum of its components and to have a local minimum.

To determine the relationship between t_p and τ , the ideal CR-RC² band-pass filter that is applied to the pulse shapes introduced in section 6.3.2 is varied in τ and the pulses' peaking times are measured as a function of τ . The peaking time t_p is defined as the time between the points at which a pulse reaches 5 % and 100 % of its amplitude. For each simulated calorimeter cell, only the lowest-energy pulse has been used, as the electronics are assumed to be per-

³It has been observed that disabling the quantization noise actually increased $s[A_{\text{tot}}]$ for large t_p . While this has not been investigated in detail, the most likely cause is a systematic rounding effect due to the size of the least significant bit, see section 6.5.2.




Figure 7.6: The numeric results (circles) and a rational-power model (line) for the dependence of the peaking time t_p on the time constant τ of the shaper.

factly linear up to this point. This removes the dependence of the peaking time on the pulse amplitude.

The results are shown in fig. 7.6 for two exemplary calorimeter cells in the middle layer of the EMB and EMEC respectively. For both cells, the relationship between τ and t_p can be modelled well by a rational-power function of the form $f(x) = ax^b$. This is also true for other cells, albeit with less accuracy. These models are used below to map from τ to t_p .

To simulate only thermal noise, the noise PSDs described in section 6.4.2 have been shaped with varying values of τ , much in the same manner as the pulse shapes. Then, the AREUS simulation has been run using these pulse shapes and noise PSDs. No deposition of energy has been simulated⁴, only the thermal noise of the analog electronics and the quantization noise of the ADC.

In this configuration, 10 000 BCs have been simulated and the ADC output has been used to calibrate an OF of filter depth $P = 5$, as described in section 6.6.1. Afterwards, 190 000 more BCs have been simulated (for a total of 200 000 BCs) and the calibrated OF has been applied to them. To avoid any short-term correlations due to the thermal-noise simulation, only every 30th sample of the OF output has been recorded. Because $E_{\text{true}} \approx 0$ for all samples, $s[\Delta_{\text{tot}}]$ has been calculated as the sample standard deviation of all recorded samples. The results of this analysis are shown with black circles in fig. 7.7.

To simulate only pileup noise, the simulation of thermal noise has been effectively disabled: the *noise scale* parameter, a factor all simulated thermal-noise values are multiplied with, has been set to zero. At the same time, energy deposits from pileup samples have been added to the simulation (see section 6.2.5 for details). Three different luminosity scenarios have been studied: $10^{33}/(\text{cm}^2 \text{ s})$, $10^{34}/(\text{cm}^2 \text{ s})$ and $10^{35}/(\text{cm}^2 \text{ s})$. The former two because they are shown in

⁴Every 30 BCs, a negligible energy of $1 \mu\text{eV}$ has been deposited; otherwise, AREUS would produce no output whatsoever.




Figure 7.7: Sample standard deviation of the inclusive reconstruction error Δ_{tot} as a function of the peaking time t_p when simulating thermal noise (circles), pileup noise (empty triangles and squares), and both (filled triangles and squares). Quantization noise is simulated in all scenarios. The one-sigma half range $\varsigma[\Delta_{\text{tot}}]$ shows the same behavior as $s[\Delta_{\text{tot}}]$, but depends on t_p in a less smooth way.

fig. 3.13, the latter one because it corresponds roughly to the maximal luminosity expected at the HL-LHC. The three scenarios have been parameterized approximately as $\langle\mu\rangle = 2$, $\langle\mu\rangle = 20$, and $\langle\mu\rangle = 200$ respectively.

The modified simulation has been run as before and Δ_{tot} has been calculated. Since $E_{\text{true}} \neq 0$ here, the energy deposited by pileup has to be taken into account. Figure 7.7 shows the results of this analysis with empty triangles and squares.

For the full noise simulation, the above simulation of pileup noise is run with thermal noise re-enabled. The results are shown in fig. 7.7 with filled triangles and squares. Crucially, because the total noise has been calculated from an independent run of the simulation, it needs not come out as exactly the square sum of thermal and pileup noise.

The graphs in fig. 7.7 qualitatively confirm the conclusions made from fig. 3.13. The curves differ slightly in shape – the pileup noise curve, in particular, flattens towards the two ends of the t_p interval. The most likely cause for this is that the signal is shown after filtering here, but before filtering in fig. 3.13.

This analysis has two important results:

1. the peaking time that minimizes the total for $\langle\mu\rangle = 20$ is confirmed to be close to 45 ns;
2. close to the minimum, the total noise curve is nearly flat, which means that a wide range of τ values give near-optimal results.

This suggests that there is little benefit to optimizing τ during the Phase-II Upgrade. This is confirmed more concretely in the following section 7.4.

7.4 The Analog Readout Electronics and the ADC

The purpose of this section is to study how certain parameters of the readout electronics influence the energy reconstruction. Concretely, these parameters are the transfer function of the analog shaper, the time constant τ of the shaper, and the sampling rate of the ADC. The simulations for this section are set up in the same manner as for section 7.1. Where there are differences, they are pointed out.

The analog shaper, as explained in section 3.3.4, is a CR-RC² band-pass filter. This is a shorthand notation for one high-pass followed by two low-pass filters. The high-pass filter gives the triangular LAr pulse a bipolar shape to facilitate estimating its amplitude; the low-pass filters reduce the electronic noise by cutting off high frequencies.

The exact number of low-pass filters presents a trade-off that is highly similar to the optimization of τ . Increasing their number reduces electronic noise, but it also makes the positive lobe of the pulse wider, hence increasing the effect of out-of-time pileup. With the expected increase in pileup after the Phase-II Upgrade, it stands to question whether the current number still is optimal.

To answer this question, the pulse shapes and noise PSDs described in sections 6.3 and 6.4 have been shaped with three different transfer functions:

$$H_n(f) = \frac{2\pi i f \tau}{(1 + 2\pi i f \tau)^{n+1}}, \quad n \in \{1, 2, 3\}. \quad (7.3)$$

They describe a shaper with one, two, and three low-pass filters respectively. When applying them to a PSD, care must be taken to take the magnitude squared of the transfer function: $S'(f) = |H(f)|^2 S(f)$.

These three sets of pulse shapes and noise PSDs have been used in simulation runs that are otherwise as described in section 7.1. The results are shown in fig. 7.8. Using a single low-pass filter, though it theoretically reduces out-of-time pileup, is not worth the resulting increase in electronic noise. While this is trivially true at $\eta = 0.5$, where pileup is low, the difference is significant even at $\eta = 2.0$. Conversely, adding a third low-pass filter gives a modest improvement, most notable at $\langle\mu\rangle = 0$.

In a related simulation study, also using AREUS, the same comparison has been performed for a cell in the HEC at $\eta = 2.35$, where pileup is much higher [111, fig. 3.2b]. It has been found that a third low-pass filter degrades the performance and using a single low-pass filter is at best equivalent to the current shaper. Thus, it must be concluded that, while better options exist for individual subsystems, only the current CR-RC² filter is optimal for both calorimeters and will remain so after the Phase-II Upgrade. For the FCal, no comparable study is known at this point and no definitive statement can be made.

The time constant of the shaper has already been studied extensively in section 7.3. It is nonetheless useful to present these results in a different manner. For one, it is τ and not t_p that can be configured in the electronics. Furthermore, τ cannot be modified to arbitrary precision, so it is more realistic to investigate only a few distinct values of it.




Figure 7.8: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different band-pass filters realized in the analog shaper. Shown are the results for one high-pass filter followed by one (circles), two (squares), and three (triangles) low-pass filters. The bottom figure shows the relative difference between each scenario and the baseline solution.




Figure 7.9: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different values of the time constant τ of the shaper. The bottom figure shows the relative difference between each scenario and the baseline solution.




Figure 7.10: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different sampling rates. The filter depth P of the OF has been adapted to each sampling rate. The bottom figure shows the relative difference between each scenario and the baseline solution.

For this reason, a study has been carried out that is similar to the above; but instead of varying the shaper's transfer function, a CR-RC² band-pass filter has been used throughout and the time constant τ has been varied between 8 ns, 13 ns and 18 ns. The results are shown in fig. 7.9. They largely mirror the results for the transfer function: lowering τ does not decrease pileup noise enough to compensate the increase in electronic noise. Conversely, increasing τ does decrease the total noise, but only slightly.

In a related simulation study (see above), the same comparison has been performed in the HEC at $\eta = 2.35$ [111, fig. 3.3b]. There, it has been found that choosing $\tau = 8$ ns reduces $s[\Delta_{\text{sig}}]$ significantly by 11 %. For this reason, it is being considered to make τ programmable in the FEB2 and adapt it to each detector region [111, sec. 3.6].

The sampling rate of the ADC is another parameter that promises improvement of the energy reconstruction. Doubling it from 40 to 80 MHz gives more data points of the signal pulse and makes its approximation by linear filters more accurate.

To test this hypothesis, another simulation study has been carried out; it has been set up as the ones before, but the shaper has been left in its default configuration (CR-RC², $\tau = 13$ ns). Instead the simulation has been run once with the default sampling rate of 40 MHz and once with two samples per BC, corresponding to 80 MHz. To ensure that the OF actually makes use of the additional resolution, its filter depth has been doubled to $P = 10$ in the latter case. The results of this study are shown in fig. 7.10.

As expected, increasing the sampling rate consistently improves energy reconstruction. The improvement is about 5 % in the EMB and about 10 % in the EMEC. A related study in the HEC at $\eta = 2.35$ found a similar improvement to that in the EMEC [111, 3.7b].




Figure 7.11: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different values of the OF depth P . The bottom figure shows the relative difference between each scenario and the baseline solution.

While this is the strongest improvement over the baseline so far presented in this thesis, it comes with a high cost: doubling the sampling rate means that double the bandwidth is necessary to transmit the sampled signal to the back end. Similarly, it requires double the bandwidth of the ADC, which is difficult to realize in the chip technologies that are required due to the radiation levels in the detector. Ultimately, it has been decided that the improved energy resolution is not worth the additional cost. Hence, the Phase-II Upgrade is foreseen to use a sampling rate of 40 MHz.

7.5 The Optimal Filter (OF)

In this section, the OF and the influence of its parameters on the energy resolution are studied. The simulation is as described in section 7.1 in all parameters except the ones being varied.

The filter depth P is the main hyperparameter of the OF that is not given by the experimental data. With the Phase-II Upgrade bringing more efficient electronics to the back end, it is possible that a computationally more expensive OF, using more coefficients to perform its calculations, improves energy resolution.

Hence, three different scenarios have been simulated: $P = 5$ (the number used during Run 1), $P = 9$, and $P = 13$. The results are shown in fig. 7.11. As can be seen, $P = 9$ gives a modest improvement of 2 to 3 %. $P = 13$ improves energy resolution further in the EMB, but not in the EMEC. A related study for a HEC cell at $\eta = 2.35$ gives results that are qualitatively similar to those in the EMEC [111, fig. 4.3b].




Figure 7.12: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing. Shown are the results: 1. of an OF whose coefficients are tuned to each pileup scenario (circles); and 2. of an OF that has been calibrated once for $\langle\mu\rangle = 140$. The bottom figure shows the relative difference between each scenario and the baseline solution.

Fixed coefficients are not an optimization, but rather a necessity. The OF coefficients depend on the autocorrelation R of the noise (see section 5.6), but R changes over time. The reason for this is that the instantaneous luminosity decreases exponentially over the course of an LHC run⁵. This changes the shares with which electronic and pileup noise contribute to the total noise. Consequently, while a given set of OF coefficients may be optimal for the beginning of a run, this may not be the case close to the end of it.

To investigate this issue, the baseline simulation, as described in section 7.1, has been compared to a scenario in which the OF has been pre-calibrated at $\langle\mu\rangle = 140$. In the latter case, since the first 40 000 BCs have not been used for OF calibration, they have been used for evaluation. The impact on the statistics is negligible, as they make up only 2.7 % of the full dataset.

The results of the comparison are presented in fig. 7.12. They show that energy resolution indeed deteriorates when going from $\langle\mu\rangle = 140$ to lower luminosities. However, at $\langle\mu\rangle = 200$, no significant difference can be found in the performance of both filters.

The reason for this is that at $\langle\mu\rangle = 140$, pileup already is the dominant contribution to the total noise. If its magnitude increases even further, its *share* of the noise autocorrelation increases only negligibly. These results are consistent with a related simulation study of a HEC calorimeter cell at $\eta = 2.35$ [111, fig. 4.6b].

⁵See, however, the footnote on luminosity leveling on page 15.




Figure 7.13: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different pre- and post-peak samples of the Wiener filter. The bottom figure shows the relative difference between each scenario and the baseline solution.

7.6 The Wiener Filter

In this section, the Wiener filter and the influence of its parameters on the energy resolution are studied. The simulation is as described in section 7.1 in all parameters except the ones being varied.

The desired output sequence can be varied in several ways, the most common of which are peak broadening and narrowing (see section 5.3.3). While peak narrowing is done automatically in AREUS, there are multiple choices for manual peak broadening (see section 6.6.2). Both the pre-peak and the post-peak sample can be set to either 0 or 0.5, giving four scenarios in total. All four have been simulated to find the optimal choice. The results are shown in fig. 7.13.

As expected, the Wiener filter with a delta-impulse-like response performs worst. This is because its response contains a lot of high frequency contributions in order to form its sharp peak. Consequently, a lot of high-frequency electronic noise passes through.

The peak-broadened filters, by contrast, remove high-frequency contributions better and hence reduce the impact of electronic noise. Among them, the variant with the broadest peak ($s_{\text{pre}} = s_{\text{post}} = 0.5$, “Both” in fig. 7.13) shows the best energy resolution. While this suggests that one could improve the results further by broadening the peak even further, this approach eventually meets the same fate as using more low-pass filters in analog shaping: broadening the peak increases out-of-time pileup effects and thus deteriorates energy resolution in a setting where signal pulses may overlap.

The filter depth P can be optimized for the Wiener filter just like for the OF. For this study, P has been varied between 4 and 9 to find the optimum. A summary of the results is shown




Figure 7.14: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different values of the Wiener filter depth P . The bottom figure shows the relative difference between each scenario and the baseline solution. The difference between the curves for $P = 6$ and $P = 8$ is smaller than 0.5 % and within their uncertainty.

in fig. 7.14. As can be seen, the Wiener filter reaches optimal resolution at $P = 6$ and does not improve further beyond this point.

The issue of fixed coefficients, as it applies to the OF, has not been investigated deeply for the Wiener filter. As the calibration procedure of the Wiener filter does not actually make use of the electronic noise, it is trivially independent of the calibration luminosity. This has been verified by the fact that the calibration arrives at exactly the same filter coefficients in all five pileup scenarios.

7.7 The Wiener Filter with Forward Correction (WFFC)

In this section, the WFFC and the influence of its parameters on the energy resolution are studied. The simulation is as described in section 7.1 in all parameters except the ones being varied.

Peak broadening and narrowing have been investigated for the WFFC in the same manner as for the Wiener filter in the previous section. Figure 7.15 shows the results. They generally agree with fig. 7.13. However, there are a few cases where the WFFC performs worse than the bare Wiener filter, e.g. in the EMB by about 7 %, independent of $\langle\mu\rangle$. Forward correction brings a strong improvement only in the EMEC at high pileup, also amounting to approximately 7 %.

The reason for this effect can be seen in fig. 7.16. It shows how often the WFFC falsely detects energy depositions, so-called false positives. This generally happens because the elec-




Figure 7.15: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different pre- and post-peak samples of the WFFC. The bottom figure shows the relative difference between each scenario and the baseline solution.




Figure 7.16: Histograms of the false-positive rate (FPR) of four WFFCs with different peak-broadening settings. The FPR is defined as the number of false positives in a simulation run divided by the number of simulated BCs. A false positive is a BC in which no energy was deposited ($E_{\text{true}} = 0$), but the filter reconstructed some energy ($E_{\text{reco}} \neq 0$). The simulation has been run at $\langle\mu\rangle = 0$.




Figure 7.17: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different selection algorithms of the WFFC. See the text for details. The bottom figure shows the relative difference between each scenario and the baseline solution.

tronic noise occasionally peaks in a manner similar enough to the expected Wiener filter response pulse to pass the selection algorithm. Each false positive triggers a forward correction. In absence of the expected pulse, this distorts the signal seen by the Wiener filter in later BCs.

The effect is strongest without any peak broadening, as the selection criterion is the most permissive in this case. The false-positive rate can be reduced by making the selection criterion more restrictive, as fig. 7.16 shows. It also naturally decreases with increasing pileup, as the probability of a cell not being hit in a BC eventually converges to zero. This is the reason why the WFFC performs best at high $\langle\mu\rangle$ and in the far-forward region.

More restrictive selection criteria have been investigated to solve the issue of false positives. To do so, the baseline WFFC first has been simulated at $\langle\mu\rangle = 0$ and the false positive with the highest reconstructed energy has been recorded. This energy, E_{max}^f , is 600 MeV in the EMB and 360 MeV in the EMEC. Then, the simulation has been run again with E_{max}^f as a threshold: only peaks with a higher reconstructed energy pass the WFFC selection. Peaks below the threshold are not selected and do not cause forward correction.

Another method to reduce the number of false positives that has been investigated is a so-called “max check”. It is an additional restriction on the selection criterion, requiring that the pre-peak and post-peak sample not only exceed a given threshold, but also stay below the value of the peak sample. The motivation for this is that electronic noise is of the order of magnitude of the least significant bit. As such, the most common way by which a false positive occurs if three subsequent samples have a value of one or two ADC counts.

Figure 7.17 shows the results for both selection restrictions, as well as for the baseline solution. The “max check” restriction does not improve energy resolution significantly. The




Figure 7.18: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different values of the WFFC depth P . The bottom figure shows the relative difference between each scenario and the baseline solution.

threshold actually deteriorates energy resolution by up to 2 % in the EMB and non-significantly by up to 1 % in the EMEC. This effect can be traced back entirely to BCs in which the true deposited energy is below the threshold and thus does not pass the selection criterion. It is weaker in the EMEC as its threshold is lower. In summary, neither restriction of the selection criterion improves energy resolution.

The filter depth P has been optimized for the WFFC in a manner analogous to the bare Wiener filter. P has been varied between 4 and 9; the results for a selection of these values are shown in fig. 7.18. Like the Wiener filter, the WFFC reaches an optimal resolution at $P = 6$. While $P = 5$ gives slightly better results in the EMEC, the difference is less than the uncertainty and thus likely not significant.

The selection margin, m , influences the energy resolution by either relaxing or restricting the selection criterion (see section 6.6.3). To investigate its impact, simulations have been run in which it has been varied between 0.05 and 0.5 in steps of 0.05. The value $m = 0.5$ is the theoretical maximum. The value $m = 0.0$ has not been used for technical reasons⁶.

Figure 7.19 shows the results. For visual clarity, only every other value of m is presented. In both the EMB and the EMEC, the baseline solution of $m = 0.25$ gives nearly optimal results. Though higher values perform slightly better (by 1 % for $m = 0.45$ and $\langle \mu \rangle = 200$), the difference is within the range of uncertainty and not necessarily significant.

⁶If $m = 0$, the selection criterion becomes too strict and the ideal Wiener filter response in the EMEC no longer passes the selection. The ideal WFFC response thus becomes a flat zero, which AREUS cannot handle.




Figure 7.19: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different values of the WFFC margin m . The bottom figure shows the relative difference between each scenario and the baseline solution.

Interestingly, although the margin m effectively relaxes the selection criterion, false positives do not seem to pose a problem in this case. One possible reason is that placing any restriction whatsoever on the samples before and after the peak candidate sample suffices to suppress the majority of false positives (see fig. 7.16).

7.8 Final Comparison and Conclusions

Figure 7.20 shows the energy resolution of all three filters after optimization. In the EMB, forward correction actually *deteriorates* the energy resolution of the Wiener filter by 0.5 to 1.0 %, depending on $\langle \mu \rangle$. The OF outperforms both Wiener filters by a significant fraction: at $\langle \mu \rangle = 200$, its noise is only 74 % of that of the WFFC.

In the EMEC, the picture is similar for $\langle \mu \rangle \leq 20$. For higher pileup, however, the relation changes: at $\langle \mu \rangle = 80$, OF and WFFC give the same resolution; at $\langle \mu \rangle = 140$ and $\langle \mu \rangle = 200$, the WFFC improves energy resolution by up to 10 %. The bare Wiener filter is equivalent to the OF at $\langle \mu \rangle = 140$ and outperforms it at $\langle \mu \rangle = 200$ by about 3 %. This is comparable to the improvement seen when increasing the OF filter depth to $P = 9$.

The performance of the WFFC differs strongly between EMB and EMEC. This can be explained using fig. 7.21, which shows the probability of a cell to be hit in any given BC. Since pileup increases with $|\eta|$, this probability is much higher in the EMEC than in the EMB. For $\langle \mu \rangle \geq 80$, it even approaches one. In such a scenario, it becomes impossible for the WFFC to find false positives (see section 7.7), as every detected peak is a true positive. Hence, the negative impact of false positives on the energy resolution vanishes.




Figure 7.20: One-sigma half range of the inclusive reconstruction error Δ_{tot} as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing, for different filters. The bottom diagram shows the relative difference between each filter and the WFFC.




Figure 7.21: True-positive rate (TPR) of two different simulated cells as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing. The TPR is defined as the number of simulated BCs in which some energy was deposited in the cell ($E_{\text{true}} \neq 0$) divided by the total number of simulated BCs.




Figure 7.22: Median of the inclusive reconstruction error Δ_{tot} as a function of $\langle\mu\rangle$, the average number of pp interactions per bunch crossing, for different filters.

Figure 7.22 shows the median of Δ_{tot} for all three filters. It shows the pileup-dependent bias that has been explained in section 6.7. In the EMEC, OF and Wiener filter show nearly exactly the same behavior; in the EMB, the Wiener filter seems to be less biased by approximately 2 MeV at $\langle\mu\rangle = 200$.

This difference, however, is incidental: in the EMB middle layer, the signal pulse’s undershoot has a minuscule “dip” at its rising end (see fig. 6.4) that makes it slightly longer than the expected 625 ns. Due to this, the analyzed signal pulses overlap slightly and each pulse is biased by its predecessor’s undershoot. The undershoot of the Wiener filter response is less sensitive to this effect than that of the OF and does not suffer from this problem.

Among the three filters, the WFFC has the smallest bias. The reason for this is a better suppression of out-of-time pileup on the one hand, and an overestimation of the deposited energy due to false positives on the other. The two effects cannot be separated easily. For this reason, even if the WFFC is employed, it will remain necessary for later processing stages to subtract the average amount of pileup⁷.

Figure 7.23 gives a comparison of the OF and the WFFC in all six reference calorimeter cells. The middle-layer cell at $\eta = 1.0$ shows roughly the same behavior as the one at $\eta = 0.5$, though the WFFC performance is much worse than expected. The most likely cause for this is that the pulse shape is slightly different and requires the WFFC hyperparameters to be tuned differently. It has been found that performance is particularly sensitive to the depth and positioning of the two corrective FIR filters.

In the two front-layer cells, both filters are competitive for $\langle\mu\rangle \geq 80$. There is the unusual effect at $\eta = 1.0$, where WFFC performance drops for $\langle\mu\rangle = 200$. As far as can be discerned, the reason is that the WFFC occasionally assigns the wrong bunch crossing identifier (BCID) to high-energy hits. This is possible because the WFFC, as implemented in AREUS, performs

⁷This average amount of pileup may be parameterized by the position of a BC in its bunch train; this technique is referred to as *baseline correction* [114, sec. 5.2].




Figure 7.23: One-sigma half range of the inclusive reconstruction error Δ_{tot} of the OF and WFFC as a function of $\langle \mu \rangle$, the average number of pp interactions per bunch crossing. Shown are the results for all six reference calorimeter cells. The bottom diagram shows the relative difference between the OF and the WFFC. Positive values mean that the OF energy resolution is worse.

its own BC identification even if the correct BC is known. Unfortunately, this issue could not be explored further due to time constraints.

The reason that the WFFC performs better in the front than in the middle layer is that pileup noise is much more dominant there, as fig. 7.24 shows. In the presampler, where the level of pileup is comparable to the front layer, the WFFC outperforms the OF by about 10 %. It additionally profits from the fact that the average signal energy to reconstruct is much higher than in the front layer (see section 6.2.5), meaning that fewer hits from the signal sample file deposit energies that are of the order of magnitude of the electronic noise.

In summary, the WFFC improves upon the OF energy resolution, though this result depends highly on the scenario. The WFFC works best in high-pileup environments, where both the occurrence and the average energy of hits is high. A prospective filter design that would improve on these results, and be applicable in the central detector region as well, would have take this into account and take measures against false positives due to electronic noise.




Figure 7.24: Ratio of pileup to electronic noise after digitization and before filtering. The data has been acquired by simulating 15 000 BCs with negligible signal energy depositions ($E_{\text{true, sig}} = 1 \mu\text{eV}$) at varying pileup conditions $\langle \mu \rangle$. The small, but non-zero signal energy avoids division by zero when AREUS calculates certain ratios. Given the sample standard deviation $s(\langle \mu \rangle)$ of the resulting sample sequences, the electronic noise is $s_{\text{elec}} := s(0)$, and the pileup noise is $s_{\text{pu}}(\langle \mu \rangle) := \sqrt{s^2(\langle \mu \rangle) - s^2(0)}$.

Chapter 8

Conclusions and Outlook

The LHC physics program for the next two decades will require extensive upgrades to both the accelerator and the experiments. To accelerate data taking, the accelerator’s instantaneous luminosity will be increased by at least a factor of five from its design value. This is necessary to achieve two of the goals of the ATLAS physics program: reducing the uncertainties on measurements of Standard Model processes, particular those involving the Higgs boson; and the continued search for particles that have been predicted by extensions of the Standard Model. At the same time, trigger efficiencies of the ATLAS detector must be improved in order to keep up with the increased event rate without discarding interesting data.

Two keys to this are an accurate energy reconstruction at the calorimeter cell level and an efficient suppression of out-of-time pileup effects. To this end, it does not suffice to minimize the noise in the analog readout electronics of the calorimeter; the algorithms that reconstruct the deposited energy must be optimized as well. This opens up an enormous parameter space that must be explored if one wants to find the optimal energy reconstruction.

In this thesis, a variety of energy reconstruction algorithms have been reviewed in an effort to map out this space. A strong focus has been placed on Wiener filters and matched filters such as the Optimal Filter (OF); they have been analyzed both in the time and the frequency domain. A filter dubbed *matched Wiener filter* has been studied theoretically and compared to the original Wiener filter. To the author’s knowledge, this has been the first time in the context of ATLAS calorimetry.

In addition to linear filters, more advanced algorithms have been presented. Examples include Matching Pursuit, as an example of sparse-signal restoration algorithms, and convolutional neural networks. This lays a foundation for future studies of these algorithms in the context of energy reconstruction.

The AREUS simulation, which has been presented and described in detail in this thesis, has been maintained and improved continuously. The simulation of thermal noise, which originally only comprised white noise, has been extended with pink-noise and spectral-noise models. Randomness of the simulation of pileup hits has been improved by implementing the *random-skip* feature. The analog-electronics simulation of the HEC has been refined and that of the FCal has been made possible. This is of great importance for current and future research into energy-reconstruction algorithms for the Phase-II Upgrade, as virtually all studies require a

faithful simulation of the analog electronics.

Additionally, the implementation of the Wiener filter with forward correction (WFFC) has been extended to make its selection criterion configurable. (Originally, the selection algorithm was based exclusively on the post-peak sample.) Depending on the simulated cell and the pileup scenario, this has improved the energy resolution of the WFFC by 15 to 100 %.

The AREUS simulation has been verified using Run-2 calibration data. This effort has surfaced several bugs in the simulation and data handling, which have subsequently been corrected. At the same time, the SPICE simulation of pulse shapes for the Phase-I Upgrade has been verified by comparison to official documentation and documented for future use. For Phase-II Upgrade studies, a database has been aggregated that contains simulated pulse shapes and noise power spectral densities (PSD) for six representative cells of the electromagnetic calorimeters.

Using AREUS, three parameters of the analog readout electronics have been optimized for the Phase-II Upgrade: the time constant of the shaper, the number of low-pass filters in the analog shaper, and the sampling rate of the ADC. These studies answered vital design questions in the lead-up to the technical design report for the Phase-II Upgrade of the LAr calorimeters. The decision between a sampling rate of 40 MHz and 80 MHz was particularly pressing, as it had a strong influence on the subsequent hardware design process.

After this, an in-depth comparison between three of the reviewed filters has been made: the OF as a well-understood baseline algorithm; the Wiener filter as another linear filter with higher emphasis on good signal identification; and the WFFC as a non-linear extension that uses this property to improve its energy resolution. Each filter has been optimized for best energy resolution. In particular, it has been shown that increasing the OF filter depth to $P = 9$ can improve the energy resolution by 2 to 3 %. This improvement is best observable at high luminosities in the EMEC, a domain that has not been investigated before.

It has been shown that the WFFC consistently outperforms the OF whenever electronic noise is negligible in comparison to pileup effects. This improvement is strongest (10 %) in the presampler and at high $|\eta|$. The degradation of the WFFC performance has been traced back to noise spikes falsely classified as signal peaks. This effect is strongest in the central region of the middle layer, where cell hits are rare. As luminosity increases, the probability of a cell not to be hit in any given bunch crossing goes to zero and the effect vanishes.

One of the problems that have been encountered is that the WFFC is susceptible to the parameterization of the pulse undershoot. Though it is possible to tune these parameters to each cell's pulse shape, it would be more scalable to modify the WFFC such that these parameters become superfluous. One possibility is to simplify the current forward correction such that it uses a single FIR filter deep enough to cover the entire undershoot.

For future studies, there are multiple avenues to consider. The technique of forward correction has shown promising results; methods to make it robust against electronic noise have been suggested. Similarly, algorithms for sparse signal restoration, while being considerably more complex, seem well suited as well. Finally, artificial neural networks have shown considerable success on a multitude of tasks and should be considered for energy reconstruction as well.

Finally, this thesis has demonstrated that after the Phase-II Upgrade, pileup effects will

increase, but not homogeneously so. Any energy reconstruction algorithm to be used at ATLAS will have to perform well both with negligible and non-negligible electronic noise. If it has to be tuned for a specific noise mixture, this process has to be both automatic and predictable.

In conclusion, this work provides largely improved simulation studies that have been vital in further developing the energy reconstruction in the ATLAS LAr calorimeters. New algorithms have been presented that make it possible to surpass the theoretical limits of linear filters. This effort is necessary for precise and efficient measurements of electrons and photons at the LHC in the high-luminosity era.

Appendix A

Derivation of the Liquid-Argon Electronics Pulse Shape

In order to derive the shape of current pulses induced into the readout electronics by ionization events, the following approximations are useful:

- The problem is one-dimensional. Ignoring the dimensions transverse to the axis between electrodes simplifies the math and is allowable as long as readout cells are smaller than a typical electromagnetic shower.
- The electric field between the electrodes is constant. This is reasonable as long as the high-voltage supply can react quickly enough to the (rather small) drop in voltage that is caused by electron absorption at the anode, ion recombination at the cathode, and the current protection resistors between the high-voltage supply and the electrodes.
- Ion charge does not build up. This approximation facilitates the derivation, but does not actually hold up in the ATLAS LAr calorimeter. Especially in the FCal, ionization of the liquid argon occurs at such a high rate that new events occur before argon ions from previous events have recombined. These residual ions screen the electric field between electrodes and distort the signal pulse. This effect has already been investigated in detail [257] and is ignored here.

In the simplest geometry, a LAr calorimeter consists of two planar electrodes: The cathode at $z = 0$ is connected to the ground, the anode at $z = a$ is connected to a high voltage V_0 . The distance between them is a . The electrodes thus form a capacitor with a homogeneous electric field $\mathbf{E} = \frac{V_0}{a} \mathbf{e}_z$ between them.

If an ionization event occurs, positive and negative charges (argon ions and free electrons resp.) are produced in the space between the electrodes. The electric field $\mathbf{E}$ causes the negative charges to drift towards the anode and the positive charges to drift to the electrode. Moving




Figure A.1: An ionizing particle crosses a LAr ionization chamber at a small angle. When reduced to one dimension, this corresponds to a localized line charge. The drift time t_{d-x} is variable and depends on the distance between the particle track and the electrodes.

the electric charge drains energy from the capacitor at a rate:

$$\frac{dE}{dt} = \frac{dz}{dt} E q(t) \quad (\text{A.1})$$

$$= \frac{dz}{dt} \frac{V_0}{a} q(t), \quad (\text{A.2})$$

where $\frac{dz}{dt}$ is the velocity with which the charge is moved and $q(t)$ is the amount of charge between the electrodes at a given time t . The fact that horizontal movement is being neglected has been used, as well as the definition of E .

The energy stored in a capacitor is:

$$E = \frac{1}{2} \frac{Q^2}{C}, \quad (\text{A.3})$$

where C is its (constant) capacitance and Q (not to be mixed up with q) is the charge stored on its electrodes. Because $C := \frac{Q}{V}$, the charge is always proportional to the voltage across the electrodes. Hence, without external influence, the drifting charges in the liquid argon would reduce the electrode voltage over time.

However, the calorimeter's high-voltage supply keeps the voltage constant via a recharging current. The current is proportional to the change in energy as follows:

$$\frac{dQ}{dt} = \frac{d}{dt} \sqrt{2CE} \quad (\text{A.4})$$

$$= \frac{dE}{dt} \frac{C}{\sqrt{2CE}} \quad (\text{A.5})$$

$$= \frac{dE}{dt} \frac{1}{V}, \quad (\text{A.6})$$




Figure A.2: An ionizing particle crosses a LAr ionization chamber at a large angle. When reduced to one dimension, this corresponds to a uniform ionization between the electrodes. The drift time t_d is nearly constant; it depends only on controllable factors like the voltage V_0 or the argon purity.

where the voltage $V = V_0$ is kept constant. Since the recharging current compensates the energy that is put into the drifting ionization charges, it is:

$$\frac{dQ}{dt} = \frac{dz}{dt} \frac{1}{a} q(t), \quad (\text{A.7})$$

From experiments [258, 259], it is known that at sufficiently high voltages, electrons drift through liquid argon at a constant speed $\frac{dz}{dt} = \frac{a}{t_d}$ (where t_d is the *drift time*) and that argon ions move about two orders of magnitude slower. Due to this, both the movement and the screening effect of the argon ions are neglected here.

The exact pulse shape then depends on $q(t)$. For example, if a track is only at a slightly odd angle to the electrodes, $q(t)$ would be constant Ne for a time t_{d-x} , where N is the number of produced ionization pairs, e is the elementary charge, d is the distance between electrodes and x is the minimum distance between the track and the anode. The pulse would then be approximately rectangular (see fig. A.1).

However, due to the accordion shape of the Electromagnetic Calorimeters (EMB and EMECs), almost all ionization tracks cross the entire distance between the electrodes (see fig. A.2). Due to the electrons' constant drift speed, the charge between electrodes can be assumed as:

$$q(t) = \begin{cases} Ne \left(1 - \frac{t}{t_d}\right) & \text{if } t < t_d, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.8})$$

Putting everything together, the typical current pulse produced by an ionization track in the LAr calorimeters is triangular in shape:

$$I(t) := \frac{dQ}{dt} = \begin{cases} \frac{Ne}{t_d} \left(1 - \frac{t}{t_d}\right) & \text{if } t < t_d, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.9})$$

List of Figures

2.1	The Standard Model after electroweak symmetry breaking	4
2.2	The Standard Model before electroweak symmetry breaking	6
3.1	The CERN accelerator complex	13
3.2	The standard filling scheme at the LHC	14
3.3	Cut-away view of the ATLAS detector	16
3.4	Cut-away view of the ATLAS inner detector	18
3.5	Cut-away view of the ATLAS calorimeter system	19
3.6	Cut-away view of the ATLAS Muon Spectrometer	21
3.7	Sketch of an EMB module depicting all four layers	25
3.8	Cell misalignment of the EMB front layer due to sagging	27
3.9	Example of an FCal trigger tower	28
3.10	The calorimeter response of the Electromagnetic Calorimeters	28
3.11	Block diagram of the LAr readout electronics	30
3.12	Block diagram of the LAr FEB	31
3.13	Thermal and pileup noise over the pulse peaking time	33
4.1	LHC plan for current and future upgrades	36
4.2	Exploded representation of an EMB trigger tower in η - ϕ space	37
4.3	Phase-I Upgrade of the LAr readout electronics	38
4.4	Draft of the LATOME firmware architecture	39
4.5	Evolution of the global readout architecture	41
4.6	The L1 calorimeter trigger before and after the Phase-I Upgrade	42
4.7	Phase-II Upgrade of the TDAQ system	43
4.8	Connectivity diagram between the calorimeters and the FEXs	44
4.9	Phase-II Upgrade of the LAr readout electronics	45
5.1	Typical response of the readout electronics to a hit in a calorimeter cell	50
5.2	Typical one-sided PSD of electronic noise in the analog LAr readout	55
5.3	The effect of peak broadening on the filter response	59
5.4	The peak-broadening transfer function for various pre- and post-peak values	60
5.5	The Maximum Finder fails to detect a hit	64
5.6	The Maximum Finder underestimates the energy of overlapping pulses	65

5.7	The forward-correction algorithm	66
5.8	An alternative forward-correction algorithm	66
5.9	A pathological sparse-signal restoration problem	69
5.10	Atoms from a typical LAr pulse and their mutual coherence with each other	69
5.11	A feed-forward neural network	72
5.12	Typical LSTM and GRU layers	73
5.13	A convolutional neural network	74
5.14	A dilated convolutional neural network	74
6.1	Example of a processing chain defined in AREUS	77
6.2	Example of several filter algorithms working cooperatively	79
6.3	Bucket similarity when using the random-skip feature	82
6.4	The ideal pulse shapes used in this thesis	85
6.5	Linearity of the simulated analog electronics	85
6.6	Typical thermal noise in the LAr readout electronics	87
6.7	Comparison of simulated and ideal noise	89
6.8	Difference between two methods of spectral-noise simulation	90
6.9	Visualization of ENOB simulation	92
6.10	Typical distributions of the reconstruction error	97
7.1	Comparison of sample mean and median for Δ_{sig} and Δ_{tot}	100
7.2	Comparison of the sample standard deviation and the one-sigma half range	101
7.3	Coefficients of variation for each filter and pileup scenario	102
7.4	Noise depending on pileup for different training sizes	103
7.5	Electronic noise in six reference calorimeter cells	104
7.6	Peaking time as a function of shaping time	106
7.7	Thermal and pileup noise as functions of peaking time	107
7.8	Noise depending on pileup for different shaping functions	109
7.9	Noise depending on pileup for different shaper time constants	109
7.10	Noise depending on pileup for different sampling rates	110
7.11	Noise depending on pileup for different OF depths	111
7.12	Noise depending on pileup for fixed and tuned OF coefficients	112
7.13	Noise depending on pileup for different desired output of the Wiener filter	113
7.14	Noise depending on pileup for different Wiener filter depths	114
7.15	Noise depending on pileup for different desired output of the WFFC	115
7.16	Histograms of the false positives of the WFFC	115
7.17	Noise depending on pileup for different WFFC selection algorithms	116
7.18	Noise depending on pileup for different WFFC depths	117
7.19	Noise depending on pileup for different WFFC margins	118
7.20	Noise depending on pileup for different filters	119
7.21	Fraction of BCs with hits	119
7.22	Bias depending on pileup for different filters	120
7.23	OF and WFFC noise depending on pileup for all simulated cells	121
7.24	Pileup over electronic noise before filtering	122

A.1	An ionizing particle crosses a LAr chamber at a small angle	128
A.2	An ionizing particle crosses a LAr chamber at a large angle	129

List of Tables

3.1	Segmentation of the LAr calorimeters into cells	26
6.1	Samples used to simulate background hits	83
6.2	Least significant bits in the electromagnetic calorimeter	93

Glossary

ADC	analog-to-digital converter
ALICE	A Large Ion Collider Experiment
AMI	ATLAS Metadata Infrastructure
ANN	artificial neural network
AREUS	ATLAS Readout Electronics Upgrade Simulation
ASIC	application-specific integrated circuit
ATCA	Advanced Telecom Computing Architecture
Athena	ATLAS offline software
ATLAS	proper name, formerly “A Toroidal LHC Apparatus”
BC	bunch crossing
BCID	bunch crossing identifier
BCR	Bunch Counter Reset
BEH	Brout–Englert–Higgs mechanism, part of the Standard Model that explains vector boson masses
BFGS	Broyden–Fletcher–Goldfarb–Shanno algorithm
BIBO	bounded-input-bounded-output
BPD	Basis Pursuit Denoising
BSM	beyond the Standard Model
CBNT	Combined N-Tuple, the predecessor of the n-tuple format called D3PD, which in turn has been superceded by the hierarchic format xAOD
CERN	European Organisation for Nuclear Research
CKM	Cabbibo–Kobayashi–Maskawa matrix, describes quark flavor mixing
CMS	Compact Muon Solenoid
CNN	convolutional neural network
COTS	commercial off-the-shelf
CSC	Cathode Strip Chambers

- CTP** Central Trigger Processor
- DAQ** data acquisition
- DCS** Detector Control System
- DFT** discrete Fourier transform
- DRE** dynamic range enhancement
- DTFT** discrete-time Fourier transform
- DØ** detector at the Tevatron
- eFEX** Electron Feature Extractor
- EMB** Electromagnetic Barrel Calorimeter
- EMEC** Electromagnetic End-Cap Calorimeters
- ENOB** effective number of bits
- EWSB** electroweak symmetry breaking
- EYETS** extended year-end technical stop
- FBP** Forward–Backward Pursuit
- FCal** Forward Calorimeters
- FEB** Front-End Board
- FEB2** Front-End Board 2
- FEC** front-end crate
- FELIX** Front End Link Interface Exchange
- FEX** Feature Extractor
- fFEX** Forward Feature Extractor
- FIR** finite impulse response
- FPGA** field-programmable gate array
- FPR** false-positive rate
- GERDA** Germanium Detector Array
- gFEX** Global Feature Extractor
- GIM** Glashow–Iliopoulos–Maiani mechanism, an explanation for the suppression of flavor-changing neutral currents; requires the charm quark to exist
- GRU** Gated Recurrent Unit
- GWS** Glashow–Weinberg–Salam model, describes the electroweak interaction
- H1** detector at HERA
- HEC** Hadronic End-Cap Calorimeters
- HERA** Hadron-Electron Ring Accelerator

- HLC1** preamp/shaper chip for the FEB2 boards of the Phase-II Upgrade, based on 65 nm CMOS technology
- HL-LHC** High-Luminosity LHC
- HLT** High-Level Trigger
- IAXO** International Axion Observatory
- IBL** Insertable B-Layer
- INL** integrated non-linearity
- IP** intellectual property
- IQR** interquartile range
- ISR** Intersecting Storage Rings
- jFEX** Jet Feature Extractor
- L0** level 0 trigger stage
- L1** level 1 trigger stage
- L2** level 2 trigger stage
- LAr** liquid argon
- LASP** Liquid Argon Signal Processor
- Lasso** least absolute shrinkage and selection operator
- LATOME** LAr Trigger Processing Mezzanine
- LAUROC** preamp/shaper chip for the FEB2 boards of the Phase-II Upgrade, based on 130 nm CMOS technology
- LDPB** LAr Digital Processing Blade
- LDPS** LAr Digital Processing System
- LEP** Large Electron–Positron Collider
- LHC** Large Hadron Collider
- LHCb** Large Hadron Collider Beauty
- LHCf** Large Hadron Collider Forward
- LOCx2** transmitter ASIC based on a low-overhead code
- LS1** Long Shutdown 1
- LS2** Long Shutdown 2
- LS3** Long Shutdown 3
- LSB** least significant bit
- LSB** Layer Sum Board
- LSP** lightest stable supersymmetric particle

- LSTM** Long Short-Term Memory
- LTDB** LAr Trigger Digitizer Board
- MC** Monte-Carlo simulation
- MDAC** multiplicative digital-to-analog converter
- MDT** Monitored Drift Tubes
- MIM** metal–insulator–metal, a capacitor technology
- MLP** multi-layer perceptron
- MoEDAL** Monopole and Exotics Detector at the LHC
- MP** Matching Pursuit
- MSE** mean squared error
- MSSM** minimal supersymmetric Standard Model
- muxer** Multiplexer
- n-tuple** a number of conventions on saving data in a ROOT file: n-tuple files have no internal directory structure, but only contain one or more trees; the branches of each tree should contain only primitive types and collections thereof
- OF** Optimal Filter
- OLS** Orthogonal Least Squares
- OMP** Orthogonal Matching Pursuit
- one-sigma half range** half the width of the interval that contains 68.27 % of the data
- one-sigma interval** the interval that contains the most central 68.27 % of the data
- PMNS** Pontecorvo–Maki–Nakagawa–Sakata matrix, describes neutrino flavor mixing
- POPOP** a secondary (i.e. wavelength-shifting) scintillator
- pp** proton–proton
- PRNG** pseudo-random number generator
- PS** Proton Synchrotron
- PSB** Proton Synchrotron Booster
- PSD** power spectral density
- PTP** para-Terphenyl, a primary scintillator
- QCD** quantum chromodynamics
- R807-808** an otherwise unnamed particle detector at the ISR [260]
- ReLU** rectified linear unit
- RNN** recurrent neural network
- ROB** Readout Buffer

- ROD** Readout Driver
- RoI** region of interest
- ROOT** C++ library and file format widely used in high-energy physics
- ROS** Readout System
- RPC** Resistive-Plate Chambers
- SBR** Single Best Replacement
- SCA** Switched-Capacitor Array
- SCT** Semiconductor Tracker
- SM** Standard Model
- SMLR** Single Most Likely Replacement
- SNO+** Successor to the Sudbury Neutrino Observatory
- SP** Subspace Pursuit
- SPICE** library and program for electronics simulation
- SPS** Super Proton Synchrotron
- TBB** Tower-Builder Board
- TDAQ** trigger and data acquisition
- TDB** Tower-Driver Board
- TDNN** time-delayed neural network
- TDR** technical design report
- TGC** Thin-Gap Chambers
- TLFN** time-lagged feed-forward network
- TOTEM** Total Elastic and Diffractive Cross Section Measurement
- TPR** true-positive rate
- TRT** Transition Radiation Tracker
- TTC** trigger, timing, and control system
- USA15** main ATLAS underground services area
- UX15** Underground Experimental Hall 15
- VEV** vacuum expectation value
- WF** Wiener filter
- WFFC** Wiener filter with forward correction
- μ the average number of pp interactions per bunch crossing
- τ the time constant of the shaper

Bibliography

References for Chapter 2

- [1] A. Pais and S. B. Treiman, *How Many Charm Quantum Numbers Are There?*, Phys. Rev. Lett. **35** (1975) 1556.
- [2] H. Fritzsch, M. Gell-Mann, and H. Leutwyler, *Advantages of the color octet gluon picture*, Physics Letters B **47** (1973) 365, ISSN: 0370-2693, URL: <http://www.sciencedirect.com/science/article/pii/0370269373906254>.
- [3] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30** (26 1973) 1346, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.30.1346>.
- [4] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, Phys. Rev. Lett. **30** (26 1973) 1343, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.30.1343>.
- [5] S. Weinberg, *Non-Abelian Gauge Theories of the Strong Interactions*, Phys. Rev. Lett. **31** (7 1973) 494, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.31.494>.
- [6] S. L. Glashow, *Partial-symmetries of weak interactions*, Nuclear Physics **22** (1961) 579, ISSN: 0029-5582, URL: <http://www.sciencedirect.com/science/article/pii/0029558261904692>.
- [7] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264.
- [8] A. Salam, “Weak and Electromagnetic Interactions”, *8th Nobel Symposium* (Lerum, Sweden, May 19–25, 1968), ed. by N. Svartholm, Almqvist and Wiksell and Wiley, 1968 367, URL: https://doi.org/10.1142/9789812795915_0034.
- [9] P. W. Anderson, *Plasmons, Gauge Invariance, and Mass*, Phys. Rev. **130** (1 1963) 439, URL: <https://link.aps.org/doi/10.1103/PhysRev.130.439>.
- [10] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (9 1964) 321, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.321>.

- [11] P. Higgs, *Broken symmetries, massless particles and gauge fields*, Physics Letters **12** (1964) 132, ISSN: 0031-9163, URL: <http://www.sciencedirect.com/science/article/pii/0031916364911369>.
- [12] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (16 1964) 508, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [13] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (20 1964) 585, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.585>.
- [14] P. W. Higgs, *Spontaneous Symmetry Breakdown without Massless Bosons*, Phys. Rev. **145** (4 1966) 1156, URL: <https://link.aps.org/doi/10.1103/PhysRev.145.1156>.
- [15] B. Björken and S. Glashow, *Elementary particles and SU(4)*, Physics Letters **11** (1964) 255, ISSN: 0031-9163, URL: <http://www.sciencedirect.com/science/article/pii/0031916364904330>.
- [16] M. L. Perl et al., *Evidence for Anomalous Lepton Production in e^+e^- Annihilation*, Phys. Rev. Lett. **35** (22 1975) 1489.
- [17] S. Carlip, *Quantum gravity: a progress report*, Reports on Progress in Physics **64** (2001) 885, URL: <https://doi.org/10.1088%2F0034-4885%2F64%2F8%2F301>.
- [18] E. Noether, *Invariante Variationsprobleme*, German, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse **1918** (1918) 235, URL: <http://eudml.org/doc/59024>.
- [19] L. Boyle, *Standard Model Of Particle Physics–Most Complete Diagram.png*, Adapted work. Published under CC BY-SA 4.0, 2018, URL: <https://commons.wikimedia.org/w/index.php?oldid=285432513>.
- [20] M. Tanabashi et al., *Review of Particle Physics*, Phys. Rev. D **98** (3 2018) 030001, URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>.
- [21] P. Tarjanne and V. L. Teplitz, *SU(4) Assignments for the Vector Resonances*, Phys. Rev. Lett. **11** (1963) 447.
- [22] Y. Hara, *Unitary triplets and the eightfold way*, Phys. Rev. **134** (1964) B701.
- [23] M. Gell-Mann, *A schematic model of baryons and mesons*, Physics Letters **8** (1964) 214, ISSN: 0031-9163, URL: <http://www.sciencedirect.com/science/article/pii/S0031916364920013>.
- [24] G. Zweig, *An SU₃ model for strong interaction symmetry and its breaking; Version 1*, tech. rep. CERN-TH-401, CERN, 1964, URL: <http://cds.cern.ch/record/352337>.
- [25] S. L. Glashow, J. Iliopoulos, and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, Phys. Rev. D **2** (7 1970) 1285, URL: <https://link.aps.org/doi/10.1103/PhysRevD.2.1285>.

- [26] J. J. Aubert et al., *Experimental Observation of a Heavy Particle J*, Phys. Rev. Lett. **33** (23 1974) 1404, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.33.1404>.
- [27] J.-E. Augustin et al., *Discovery of a Narrow Resonance in e^+e^- Annihilation*, Phys. Rev. Lett. **33** (23 1974) 1406, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.33.1406>.
- [28] *CERN Courier Volume 14, Number 12, December 1974*, (1974), URL: <https://cds.cern.ch/record/1729724>.
- [29] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, Progress of Theoretical Physics **49** (1973) 652, ISSN: 0033-068X, eprint: <http://oup.prod.sis.lan/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf>, URL: <https://doi.org/10.1143/PTP.49.652>.
- [30] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (4 1964) 138, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.138>.
- [31] S. W. Herb et al., *Observation of a Dimuon Resonance at 9.5 GeV in 400-GeV Proton-Nucleus Collisions*, Phys. Rev. Lett. **39** (5 1977) 252, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.39.252>.
- [32] F. Abe et al., *Observation of Top Quark Production in $\bar{p}p$ Collisions with the Collider Detector at Fermilab*, Phys. Rev. Lett. **74** (14 1995) 2626, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.74.2626>.
- [33] S. Abachi et al., *Search for High Mass Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV*, Phys. Rev. Lett. **74** (13 1995) 2422, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.74.2422>.
- [34] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *A Historical Profile of the Higgs Boson*, (2012), arXiv: 1201.6045 [hep-ph].
- [35] G. 't Hooft, *Renormalizable Lagrangians for Massive Yang-Mills Fields*, Nucl. Phys. **B35** (1971) 167, [,201(1971)].
- [36] G. 't Hooft and M. J. G. Veltman, *Regularization and Renormalization of Gauge Fields*, Nucl. Phys. **B44** (1972) 189.
- [37] F. Hasert et al., *Search for elastic muon-neutrino electron scattering*, Physics Letters B **46** (1973) 121, ISSN: 0370-2693, URL: <http://www.sciencedirect.com/science/article/pii/0370269373904942>.
- [38] F. Hasert et al., *Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment*, Physics Letters B **46** (1973) 138, ISSN: 0370-2693, URL: <http://www.sciencedirect.com/science/article/pii/0370269373904991>.

- [39] G. Arnison et al., *Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s}=540$ GeV*, Physics Letters B **122** (1983) 103, ISSN: 0370-2693, URL: <http://www.sciencedirect.com/science/article/pii/0370269383911772>.
- [40] G. Arnison et al., *Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider*, Physics Letters B **126** (1983) 398, ISSN: 0370-2693, URL: <http://www.sciencedirect.com/science/article/pii/0370269383901880>.
- [41] T. N. P. H. W. Group, “Combined CDF and Do Upper Limits on Standard Model Higgs Boson Production with up to 8.6 fb⁻¹ of Data”, 2011, arXiv: 1107.5518 [hep-ex], URL: http://lss.fnal.gov/cgi-bin/find_paper.pl?conf-11-354.
- [42] T. N. P. H. W. Group, “Updated Combination of CDF and Do Searches for Standard Model Higgs Boson Production with up to 10.0 fb⁻¹ of Data”, 2012, arXiv: 1207.0449 [hep-ex], URL: http://lss.fnal.gov/cgi-bin/find_paper.pl?conf-12-318.
- [43] G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. **B716** (2012) 1, arXiv: 1207.7214 [hep-ex].
- [44] S. Chatrchyan et al., *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. **B716** (2012) 30, arXiv: 1207.7235 [hep-ex].
- [45] ATLAS Collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, (2019).
- [46] A. M. Sirunyan et al., *Combination of searches for Higgs boson pair production in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. Lett. **122** (2019) 121803, arXiv: 1811.09689 [hep-ex].
- [47] G. W. Bennett et al., *Final report of the E821 muon anomalous magnetic moment measurement at BNL*, Phys. Rev. D **73** (7 2006) 072003, URL: <https://link.aps.org/doi/10.1103/PhysRevD.73.072003>.
- [48] J. Grange et al., *Muon (g-2) Technical Design Report*, tech. rep. arXiv:1501.06858. FERMILAB-FN-0992-E, Comments: 666 pages, 2015, URL: <http://cds.cern.ch/record/1983396>.
- [49] N. Raha, *The Current Status of the Fermilab Muon g-2 Experiment*, Universe **5** (2019) 43.
- [50] E. Armengaud et al., *Conceptual design of the International Axion Observatory (IAXO)*, Journal of Instrumentation **9** (2014) T05002, arXiv: 1401.3233 [physics.ins-det], URL: <https://doi.org/10.1088%2F1748-0221%2F9%2F05%2FT05002>.

- [51] K. Abe et al., *The T2K experiment*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **659** (2011) 106, ISSN: 0168-9002, URL: <http://www.sciencedirect.com/science/article/pii/S0168900211011910>.
- [52] D. Ayres et al., *The NOvA Technical Design Report*, (2007).
- [53] R. Acciarri et al., *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE)*, (2016), arXiv: 1601.05471 [physics.ins-det].
- [54] K. Abe et al., *Hyper-Kamiokande Design Report*, (2018), arXiv: 1805.04163 [physics.ins-det].
- [55] K.-H. Ackermann et al., *The Gerda experiment for the search of $0\nu\beta\beta$ decay in ^{76}Ge* , The European Physical Journal C **73** (2013) 2330, ISSN: 1434-6052, URL: <https://doi.org/10.1140/epjc/s10052-013-2330-0>.
- [56] S. Andringa et al., *Current Status and Future Prospects of the SNO+ Experiment*, Adv. High Energy Phys. **2016** (2016) 6194250, arXiv: 1508.05759 [physics.ins-det].
- [57] M. Agostini et al., *Results on Neutrinoless Double- β Decay of ^{76}Ge from Phase I of the GERDA Experiment*, Phys. Rev. Lett. **111** (12 2013) 122503, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.122503>.

References for Chapter 3

- [20] M. Tanabashi et al., *Review of Particle Physics*, Phys. Rev. D **98** (3 2018) 030001, URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>.
- [58] L. Evans and P. Bryant, *LHC Machine*, Journal of Instrumentation **3** (2008) S08001, URL: <https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08001>.
- [59] *CERN Courier Volume 48, Number 8, October 2008*, (2008), URL: <https://cds.cern.ch/record/1734296>.
- [60] *CERN Courier Volume 48, Number 9, November 2008*, (2008), URL: <https://cds.cern.ch/record/1734312>.
- [61] *CERN Courier Volume 49, Number 1, January-February 2009*, (2009), URL: <https://cds.cern.ch/record/1734323>.
- [62] *CERN Courier Volume 50, Number 4, May 2010*, (2010), URL: <https://cds.cern.ch/record/1734520>.
- [63] *CERN Courier Volume 53, Number 1, January-February 2013*, (2013), URL: <https://cds.cern.ch/record/1734891>.
- [64] *CERN Courier Volume 52, Number 4, May 2012*, (2012), URL: <https://cds.cern.ch/record/1734802>.

- [65] *CERN Courier Volume 55, Number 6, July / August 2015*, (2015),
URL: <https://cds.cern.ch/record/2036709>.
- [66] *CERN Courier Volume 58, Number 10, December 2018*, (2018),
URL: <https://cds.cern.ch/record/2649360>.
- [67] *CERN Courier Volume 59, Number 4, July/August 2019*, (2019),
URL: <https://cds.cern.ch/record/2681906>.
- [68] G. Aad et al., *The ATLAS Experiment at the CERN Large Hadron Collider*,
Journal of Instrumentation **3** (2008) So8003.
- [69] S. Chatrchyan et al., *The CMS experiment at the CERN LHC*,
Journal of Instrumentation **3** (2008) So8004.
- [70] K. Aamodt et al., *The ALICE experiment at the CERN LHC*,
Journal of Instrumentation **3** (2008) So8002.
- [71] A. A. Alves Jr. et al., *The LHCb Detector at the LHC*,
Journal of Instrumentation **3** (2008) So8005.
- [72] O. Adriani et al., *The LHCf detector at the CERN Large Hadron Collider*,
Journal of Instrumentation **3** (2008) So8006.
- [73] G. Anelli et al., *The TOTEM Experiment at the CERN Large Hadron Collider*,
Journal of Instrumentation **3** (2008) So8007.
- [74] J. Pinfold et al., *Technical Design Report of the MoEDAL Experiment*,
tech. rep. CERN-LHCC-2009-006, MoEDAL-TDR-001, CERN, 2009,
URL: <http://cds.cern.ch/record/1181486>.
- [75] E. Mobs, *The CERN accelerator complex - August 2018. Complexe des accélérateurs du CERN - Août 2018*, (2018), General Photo,
URL: <https://cds.cern.ch/record/2636343>.
- [76] O. S. Brüning et al., *LHC Design Report*, CERN Yellow Reports: Monographs,
CERN, 2004, URL: <https://cds.cern.ch/record/782076>.
- [77] R. Bailey and P. Collier, *Standard Filling Schemes for Various LHC Operation Modes*,
tech. rep. LHC-PROJECT-NOTE-323, second revision: CERN, 2003,
URL: <https://cds.cern.ch/record/691782>.
- [78] ATLAS Collaboration, *ATLAS: technical proposal for a general-purpose pp experiment at the Large Hadron Collider at CERN*, LHC Tech. Proposal, CERN, 1994,
URL: <https://cds.cern.ch/record/290968>.
- [79] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *A Historical Profile of the Higgs Boson*,
(2012), arXiv: 1201.6045 [hep-ph].
- [80] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, The European Physical Journal C **76** (2016) 653,
ISSN: 1434-6052, URL: <https://doi.org/10.1140/epjc/s10052-016-4466-1>.
- [81] J. Pequenaio, “Computer generated image of the whole ATLAS detector”, 2008,
URL: <https://cds.cern.ch/record/1095924>.

- [82] R. K. Ellis, J. Stirling, and B. R. Webber, *QCD and collider physics*, 1. paperback ed., Cambridge Univ. Press, 2003, ISBN: 9780521581899, URL: http://slubdd.de/katalog?TN_libero_mab213615594.
- [83] J. Pequeno, “Computer generated image of the ATLAS inner detector”, 2008, URL: <https://cds.cern.ch/record/1095926>.
- [84] ATLAS Collaboration, *ATLAS Insertable B-Layer Technical Design Report*, tech. rep. CERN-LHCC-2010-013. ATLAS-TDR-19, 2010, URL: <https://cds.cern.ch/record/1291633>.
- [85] M. Backhaus, *The upgraded Pixel Detector of the ATLAS Experiment for Run 2 at the Large Hadron Collider*, Nucl. Instrum. Meth. **A831** (2016) 65.
- [86] ATLAS Collaboration, *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties*, JINST **12** (2017) P05002, arXiv: 1702.06473 [hep-ex].
- [87] J. Pequeno, “Computer Generated image of the ATLAS calorimeter”, 2008, URL: <https://cds.cern.ch/record/1095927>.
- [88] J. Pequeno, “Computer generated image of the ATLAS Muons subsystem”, 2008, URL: <https://cds.cern.ch/record/1095929>.
- [89] ATLAS Collaboration, *ATLAS muon spectrometer: Technical Design Report*, Technical Design Report ATLAS, CERN, 1997, URL: <http://cds.cern.ch/record/331068>.
- [90] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, tech. rep. arXiv:2007.12539. 10, 2020 60, URL: <https://cds.cern.ch/record/2725146>.
- [91] ATLAS Collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, tech. rep. CERN-LHCC-2013-018. ATLAS-TDR-023, 2013, URL: <https://cds.cern.ch/record/1602235>.
- [92] J. H. Marshall, *On Liquid Argon Ionization Chambers of Centimeter Size*, Review of Scientific Instruments **25** (1954) 232, eprint: <https://doi.org/10.1063/1.1771029>, URL: <https://doi.org/10.1063/1.1771029>.
- [93] S. E. Derenzo, R. E. Muller, R. G. Smits, and L. W. Alvarez, *The Prospect of High Spatial Resolution for Counter Experiments at NAL: A New Particle Detector Using Electron Multiplication in Liquid Argon*, eConf **C690609V3** (1969) 79.
- [94] S. E. Derenzo et al., *Recent developments in high-resolution noble liquid counters*, tech. rep., 1970, URL: <http://cds.cern.ch/record/362737>.
- [95] G. Knies and D. Neuffer, *Development and test of a liquid-argon shower detector*, Nuclear Instruments and Methods **120** (1974) 1, ISSN: 0029-554X, URL: <http://www.sciencedirect.com/science/article/pii/0029554X74900974>.

- [96] W. Willis and V. Radeka, *Liquid-argon ionization chambers as total-absorption detectors*, Nuclear Instruments and Methods **120** (1974) 221, ISSN: 0029-554X, URL: <http://www.sciencedirect.com/science/article/pii/0029554X74900391>.
- [97] J. Engler et al., *A liquid-argon ionization calorimeter for detection of electromagnetic and hadronic showers*, Nuclear Instruments and Methods **120** (1974) 157, ISSN: 0029-554X, URL: <http://www.sciencedirect.com/science/article/pii/0029554X74901190>.
- [98] C. W. Fabjan and F. Gianotti, *Calorimetry for Particle Physics*, Rev. Mod. Phys. **75** (4 2003) 1243, URL: <https://link.aps.org/doi/10.1103/RevModPhys.75.1243>.
- [99] A. Artamonov et al., *The ATLAS Forward Calorimeter*, Journal of Instrumentation **3** (2008) P02010, URL: <https://doi.org/10.1088%2F1748-0221%2F3%2F02%2FP02010>.
- [100] B. Aubert et al., *Liquid argon calorimetry with LHC-performance specifications*, (1990), URL: <https://cds.cern.ch/record/220650>.
- [101] ATLAS Collaboration, *ATLAS liquid-argon calorimeter: Technical Design Report*, Technical Design Report ATLAS, CERN, 1996, URL: <https://cds.cern.ch/record/331061>.
- [102] N. J. Buchanan et al., *Design and implementation of the Front End Board for the readout of the ATLAS liquid argon calorimeters*, Journal of Instrumentation **3** (2008) P03004, URL: <https://doi.org/10.1088%2F1748-0221%2F3%2F03%2FP03004>.
- [103] J. Jongmanns et al., *Performance of the upgraded PreProcessor of the ATLAS Level-1 Calorimeter Trigger*, tech. rep. ATL-COM-DAQ-2019-035, Access restricted to CERN users: CERN, 2019, URL: <https://cds.cern.ch/record/2661780>.
- [104] J. Barry, E. Lee, and D. Messerschmitt, *Digital Communication*, 3rd ed., Vol. 1, Springer US, 2004, ISBN: 9780792375487, URL: <https://books.google.de/books?id=hPx70ozDJlwC>.
- [105] D. Bertolini, P. Harris, M. Low, and N. Tran, *Pileup Per Particle Identification*, JHEP **10** (2014) 059, arXiv: 1407.6013 [hep-ph].

References for Chapter 4

- [91] ATLAS Collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, tech. rep. CERN-LHCC-2013-018. ATLAS-TDR-023, 2013, URL: <https://cds.cern.ch/record/1602235>.

- [106] A. G. et al.,
High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. o.1,
CERN Yellow Reports: Monographs, CERN, 2017,
URL: <https://cds.cern.ch/record/2284929>.
- [107] P. Strizenec, *ATLAS LAr Calorimeter Commissioning for LHC Run-3*, (2020),
URL: <https://cds.cern.ch/record/2725147>.
- [108] G. Apollinari, I. Béjar Alonso, O. Brüning, M. Lamont, and L. Rossi,
The HL-LHC project, Website, 2020,
URL: <https://hlumilhc.web.cern.ch/content/hl-lhc-project>.
- [109] ATLAS Collaboration,
Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment,
tech. rep. CERN-LHCC-2012-022. LHCC-I-023, draft version for comments:
CERN, 2012, URL: <https://cds.cern.ch/record/1502664>.
- [110] ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*,
tech. rep. CERN-LHCC-2015-020. LHCC-G-166, CERN, 2015,
URL: <http://cds.cern.ch/record/2055248>.
- [111] ATLAS Collaboration,
Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter,
tech. rep. CERN-LHCC-2017-018. ATLAS-TDR-027, CERN, 2017,
URL: <http://cds.cern.ch/record/2285582>.
- [112] L. Xiao et al., *LOCx2, a low-latency, low-overhead, 2×5.12 -Gbps transmitter ASIC for the ATLAS Liquid Argon Calorimeter trigger upgrade*,
Journal of Instrumentation **11** (2016) C02013,
URL: <https://doi.org/10.1088%2F1748-0221%2F11%2F02%2Fc02013>.
- [113] M. Aleksa et al.,
ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report,
tech. rep. CERN-LHCC-2013-017. ATLAS-TDR-022,
final version presented to December 2013 LHCC, 2013,
URL: <https://cds.cern.ch/record/1602230>.
- [114] G. Aad et al., *ATLAS LAr Calorimeter trigger electronics phase I upgrade: LATOME Firmware Specifications*, tech. rep. ATL-COM-LARG-2017-052, internal note:
CERN, 2017, URL: <https://cds.cern.ch/record/2294187>.
- [115] ATLAS Collaboration,
Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System,
tech. rep. CERN-LHCC-2017-020. ATLAS-TDR-029, CERN, 2017,
URL: <http://cds.cern.ch/record/2285584>.
- [116] N. J. Buchanan et al., *Radiation qualification of the front-end electronics for the readout of the ATLAS liquid argon calorimeters*, Journal of Instrumentation **3** (2008) P10005,
URL: <https://doi.org/10.1088%2F1748-0221%2F3%2F10%2Fp10005>.

- [117] N. Morange, *LAr Preamplifier/Shaper: Specifications Document for HL-LHC*, tech. rep. ATL-COM-LARG-2019-029, draft, restricted to CERN users: CERN, 2019, URL: <https://cds.cern.ch/record/2701460>.
- [118] F. Dulucq et al., “LAUROC: ’A new electronically cooled line-terminating preamplifier for the ATLAS liquid argon calorimeter upgrade’”, *Proceedings, 2016 IEEE Nuclear Science Symposium and Medical Imaging Conference: NSS/MIC 2016: Strasbourg, France*, 2016 8069900.
- [119] G. De Geronimo, H. Chen, H. Liu, and M. Newcomer, “HLCI: A Front-End ASIC for Liquid Argon Calorimeters”, *Proceedings, 2017 IEEE Nuclear Science Symposium and Medical Imaging Conference and 24th international Symposium on Room-Temperature Semiconductor X-Ray & Gamma-Ray Detectors (NSS/MIC 2017): Atlanta, Georgia, USA, October 21-28, 2017*, 2017.
- [120] N. Morange et al., *Characterization of Preamplifier and Shaper Prototypes for the Phase-II Upgrade of the LAr Calorimeter Readout*, tech. rep. ATL-COM-LARG-2019-018, internal note: CERN, 2019, URL: <https://cds.cern.ch/record/2683319>.
- [121] R. Xu and T. Andeen, *COLUTAV3 Update*, Access restricted to CERN users, 2019, URL: <https://indico.cern.ch/event/846334/contributions/3561187/>.

References for Chapter 5

- [111] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter*, tech. rep. CERN-LHCC-2017-018. ATLAS-TDR-027, CERN, 2017, URL: <http://cds.cern.ch/record/2285582>.
- [114] G. Aad et al., *ATLAS LAr Calorimeter trigger electronics phase I upgrade: LATOME Firmware Specifications*, tech. rep. ATL-COM-LARG-2017-052, internal note: CERN, 2017, URL: <https://cds.cern.ch/record/2294187>.
- [122] G. Heinzel, A. Rüdiger, and R. Schilling, *Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new at-top windows*, (2002), URL: <http://hdl.handle.net/11858/00-001M-0000-0013-557A-5>.
- [123] J. O. Smith, *Mathematics of the Discrete Fourier Transform (DFT)*, online book, W3K Publishing, 2007, ISBN: 978-0-9745607-4-8, URL: <https://ccrma.stanford.edu/~jos/mdft/>.
- [124] N. Wiener, *Generalized harmonic analysis*, *Acta Math.* **55** (1930) 117, URL: <https://doi.org/10.1007/BF02546511>.
- [125] A. Khinchin, *Korrelationstheorie der stationären stochastischen Prozesse*, *Mathematische Annalen* **109** (1934) 604, ISSN: 1432-1807, URL: <https://doi.org/10.1007/BF01449156>.

- [126] M. Parseval des Chênes, *Mémoires présentés à l'Institut des Sciences, Lettres et Arts, par divers savans, et lus dans ses assemblées*, Sciences, mathématiques et physiques (Savans étrangers) **1** (1806) 638.
- [127] F. Flandoli and M. Barsanti, *Course on Probability and Stochastic Processes*, lecture notes, 2011, URL: <http://users.dma.unipi.it/~flandoli/Automazione.html>.
- [128] E. R. Kanasewich, *Time sequence analysis in geophysics*, University of Alberta, 1981, ISBN: 978-0-88864-074-1, URL: <https://books.google.com/books?id=k8SSLy-FYagC&pg=PA185>.
- [129] R. E. Kálmán and R. S. Bucy, *New Results in Linear Filtering and Prediction Theory*, Journal of Basic Engineering **83** (1 1961) 95, ISSN: 0098-2202, URL: <https://dx.doi.org/10.1115/1.3658902>.
- [130] N. Wiener, *Extrapolation, interpolation, and smoothing of stationary time series with engineering applications*, Technology Press of Massachusetts Institute of Technology, 1949.
- [131] A. N. Kolmogorov, *Stationary sequences in Hilbert space*, Bull. Math. Univ. Moscow **2** (1941) 1.
- [132] R. E. Kálmán, *A New Approach to Linear Filtering and Prediction Problems*, Journal of Basic Engineering **82** (1 1960) 35, ISSN: 0098-2202, URL: <https://dx.doi.org/10.1115/1.3662552>.
- [133] S. L. Lauritzen, *Time Series Analysis in 1880: A Discussion of Contributions Made by T.N. Thiele*, International Statistical Review / Revue Internationale de Statistique **49** (1981) 319, ISSN: 03067734, 17515823, URL: <http://www.jstor.org/stable/1402616>.
- [134] W. Wirtinger, *Zur formalen Theorie der Funktionen von mehr komplexen Veränderlichen*, Mathematische Annalen **97** (1927) 357, ISSN: 1432-1807, URL: <https://doi.org/10.1007/BF01447872>.
- [135] R. Burckel and R. Remmert, *Theory of Complex Functions*, 4th ed., Graduate Texts in Mathematics, Springer New York, 2012, ISBN: 9781461209393, URL: <https://books.google.de/books?id=kIHlBwAAQBAJ>.
- [136] M. Leyton, “Minimum Bias and Underlying Event Measurements with ATLAS”, *Proceedings, 3rd International Workshop on Multiple Partonic Interactions at the LHC (MPI@LHC 2011): Hamburg, Germany, 21-25 Nov 2011*, 2012 11, arXiv: 1202.2090 [hep-ex].
- [137] O. Toeplitz, *Zur Theorie der quadratischen und bilinearen Formen von unendlichvielen Veränderlichen*, Mathematische Annalen **70** (1911) 351, ISSN: 1432-1807, URL: <https://doi.org/10.1007/BF01564502>.

- [138] G. I. Gonçalves, B. S. Peralva, L. M. A. Filho, A. S. Cerqueira, and J. M. Seixas, "Energy Estimation Based on Wiener-Hopf Filtering for the ATLAS Tile Calorimeter", *XXI ENMC - Encontro Nacional de Modelagem Computacional*, 2018, URL: <https://github.com/ingoncalves/enmc-2018/blob/master/enmc2018.pdf>.
- [139] D. O. Gonçalves, *Energy reconstruction of the ATLAS Tile Calorimeter under high pile-up conditions using the Wiener Filter*, tech. rep. ATL-TILECAL-PROC-2019-002, CERN, 2019, URL: <http://cds.cern.ch/record/2674807>.
- [140] J. Duarte, L. de Andrade Filho, E. de Simas Filho, P. Farias, and J. de Seixas, *Online energy reconstruction for calorimeters under high pile-up conditions using deconvolutional techniques*, *Journal of Instrumentation* **14** (2019) P12017, URL: <https://doi.org/10.1088/2F1748-0221/2F14%2F12%2Fp12017>.
- [141] W. E. Cleland and E. G. Stern, *Signal processing considerations for liquid ionization calorimeters in a high rate environment*, *Nucl. Instrum. Meth.* **A338** (1994) 467.
- [142] S. Stärz, "Energy Reconstruction and high-speed Data Transmission with FPGAs for the Upgrade of the ATLAS Liquid Argon Calorimeter at LHC", Presented 19 May 2015, 2015, URL: <https://cds.cern.ch/record/2030122>.
- [143] D. G. Bailey, "Space Efficient Division on FPGAs", *Electronics New Zealand Conference (EnzCon'06)*, 2006 206, URL: <https://github.com/ingoncalves/enmc-2018/blob/master/enmc2018.pdf>.
- [144] J. P. Grohs, "Simulation of the upgraded Phase-1 Trigger Readout Electronics of the Liquid-Argon Calorimeter of the ATLAS Detector at the LHC", Presented 29 Feb 2016, 2015, URL: <https://cds.cern.ch/record/2135931>.
- [145] M. S. O'Brien, A. N. Sinclair, and S. M. Kramer, *Recovery of a sparse spike time series by L_1 norm deconvolution*, *IEEE Transactions on Signal Processing* **42** (1994) 3353, ISSN: 1053-587X.
- [146] L. M. Li and T. P. Speed, *Deconvolution of Sparse Positive Spikes*, *Journal of Computational and Graphical Statistics* **13** (2004) 853, ISSN: 10618600, URL: <http://www.jstor.org/stable/27594081>.
- [147] S. Mallat and Z. Zhang, *Matching Pursuits with Time-frequency Dictionaries*, *Trans. Sig. Proc.* **41** (1993) 3397, ISSN: 1053-587X, URL: <http://dx.doi.org/10.1109/78.258082>.
- [148] S. van der Walt, S. C. Colbert, and G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, *Computing in Science Engineering* **13** (2011) 22, ISSN: 1521-9615.
- [149] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition", *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, vol. 1, 1993 40.

- [150] Z. Z. Geoffrey M. Davis Stephane G. Mallat, *Adaptive time-frequency decompositions*, Optical Engineering **33** (1994) 2183 - 2191 ,
URL: <https://doi.org/10.1117/12.173207>.
- [151] S. Chen, S. Billings, and W. Luo, *Orthogonal least squares methods and their application to non-linear system identification*, International Journal of Control **50** (1989) 1873,
eprint: <https://doi.org/10.1080/00207178908953472>,
URL: <https://doi.org/10.1080/00207178908953472>.
- [152] T. Blumensath and M. E. Davies,
On the Difference Between Orthogonal Matching Pursuit and Orthogonal Least Squares,
tech. rep., University of Southampton, 2007,
URL: <https://eprints.soton.ac.uk/142469/>.
- [153] C. Morales-Pérez, J. Rangel-Magdaleno, I. Cruz-Vega, J. Ramirez-Cortes, and
H. Peregrina-Barreto,
“FPGA implementation of Orthogonal Matching Pursuit algorithm”,
2016 13th International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE), 2016 1.
- [154] W. Dai and O. Milenkovic,
Subspace Pursuit for Compressive Sensing Signal Reconstruction,
IEEE Transactions on Information Theory **55** (2009) 2230, ISSN: 0018-9448.
- [155] N. B. Karahanoglu and H. Erdogan,
Compressed sensing signal recovery via forward-backward pursuit,
Digital Signal Processing **23** (2013) 1539, ISSN: 1051-2004, URL: <http://www.sciencedirect.com/science/article/pii/S1051200413001218>.
- [156] F. Wang, G. Sun, Z. Li, and J. He,
Fusion Forward-Backward Pursuit Algorithm for Compressed Sensing,
International Journal of Wireless Information Networks **24** (2017) 436,
ISSN: 1572-8129, URL: <https://doi.org/10.1007/s10776-017-0331-x>.
- [157] J. Kormylo and J. Mendel,
Maximum likelihood detection and estimation of Bernoulli - Gaussian processes,
IEEE Transactions on Information Theory **28** (1982) 482, ISSN: 0018-9448.
- [158] C.-Y. Chi and J. Mendel,
Improved maximum-likelihood detection and estimation of Bernoulli-Gaussian processes,
IEEE Transactions on Information Theory **30** (1984) 429, ISSN: 0018-9448.
- [159] C. Soussen, J. Idier, D. Brie, and J. Duan,
From Bernoulli-Gaussian Deconvolution to Sparse Signal Restoration,
IEEE Transactions on Signal Processing **59** (2011) 4572, ISSN: 1053-587X.
- [160] J. J. Fuchs, *Recovery of exact sparse representations in the presence of bounded noise*,
IEEE Transactions on Information Theory **51** (2005) 3601, ISSN: 0018-9448.

- [161] J. A. Tropp, *Just relax: convex programming methods for identifying sparse signals in noise*, IEEE Transactions on Information Theory **52** (2006) 1030, ISSN: 0018-9448.
- [162] D. L. Donoho, M. Elad, and V. N. Temlyakov, *Stable recovery of sparse overcomplete representations in the presence of noise*, IEEE Transactions on Information Theory **52** (2006) 6, ISSN: 0018-9448.
- [163] S. Chen, D. Donoho, and M. Saunders, *Atomic Decomposition by Basis Pursuit*, SIAM Journal on Scientific Computing **20** (1998) 33, eprint: <https://doi.org/10.1137/S1064827596304010>, URL: <https://doi.org/10.1137/S1064827596304010>.
- [164] F. Santosa and W. Symes, *Linear Inversion of Band-Limited Reflection Seismograms*, SIAM Journal on Scientific and Statistical Computing **7** (1986) 1307, eprint: <https://doi.org/10.1137/0907087>, URL: <https://doi.org/10.1137/0907087>.
- [165] R. Tibshirani, *Regression Shrinkage and Selection via the Lasso*, Journal of the Royal Statistical Society. Series B (Methodological) **58** (1996) 267, ISSN: 00359246, URL: <http://www.jstor.org/stable/2346178>.
- [166] M. Wright, *The interior-point revolution in optimization: history, recent developments, and lasting consequences*, Bulletin of the American mathematical society **42** (2005) 39.
- [167] A. Nemirovski and M. Todd, *Interior-point methods for optimization*, Acta Numerica **17** (2008) 191.
- [168] S. Mehrotra, *On the Implementation of a Primal-Dual Interior Point Method*, SIAM Journal on Optimization **2** (1992) 575, eprint: <https://doi.org/10.1137/0802028>, URL: <https://doi.org/10.1137/0802028>.
- [169] M. Zibulevsky and M. Elad, *L1-L2 Optimization in Signal and Image Processing*, IEEE Signal Processing Magazine **27** (2010) 76, ISSN: 1053-5888.
- [170] C. G. Broyden, *The Convergence of a Class of Double-rank Minimization Algorithms*, IMA Journal of Applied Mathematics **6** (1970) 76, ISSN: 0272-4960, eprint: <http://oup.prod.sis.lan/imamat/article-pdf/6/1/76/2233756/6-1-76.pdf>, URL: <https://doi.org/10.1093/imamat/6.1.76>.
- [171] R. Fletcher, *A new approach to variable metric algorithms*, The Computer Journal **13** (1970) 317, ISSN: 0010-4620, eprint: <http://oup.prod.sis.lan/comjnl/article-pdf/13/3/317/988678/130317.pdf>, URL: <https://doi.org/10.1093/comjnl/13.3.317>.
- [172] D. Goldfarb, *A family of variable-metric methods derived by variational means*, Math. Comp. **24** (1970) 23, ISSN: 0025-5718, URL: <https://doi.org/10.2307/2004873>.

- [173] D. F. Shanno, *Conditioning of quasi-Newton methods for function minimization*, Math. Comp. **24** (1970) 647, ISSN: 0025-5718, URL: <https://doi.org/10.2307/2004840>.
- [174] Q. Liu et al., *Fast Neural Network Training on FPGA Using Quasi-Newton Optimization Method*, IEEE Transactions on Very Large Scale Integration (VLSI) Systems **26** (2018) 1575, ISSN: 1063-8210.
- [175] F. Rosenblatt, *The Perceptron: A perceiving and recognizing automaton Project PARA*, tech. rep. 85-460-1, Cornell Aeronautical Laboratory, 1957.
- [176] N. Madysa, “Identification of Hadronic Tau Lepton Decays at the ATLAS Detector Using Artificial Neural Networks”, Presented 12 Oct 2015, 2015, URL: <https://cds.cern.ch/record/2127017>.
- [177] M. Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, Software available from tensorflow.org, 2015, URL: <https://www.tensorflow.org/>.
- [178] F. Chollet et al., *Keras*, <https://keras.io/>, 2015, URL: <https://keras.io/>.
- [179] P. Goldsborough, *A Tour of TensorFlow*, CoRR **abs/1610.01178** (2016), arXiv: 1610.01178, URL: <http://arxiv.org/abs/1610.01178>.
- [180] S. Schramm, *Machine learning at CERN: ATLAS, LHCb, and more*, tech. rep. ATL-PHYS-PROC-2018-119, CERN, 2018, URL: <https://cds.cern.ch/record/2639991>.
- [181] ATLAS Collaboration, *A neural network clustering algorithm for the ATLAS silicon pixel detector. A neural network clustering algorithm for the ATLAS silicon pixel detector*, JINST **9** (2014) P09009. 38 p, URL: <https://cds.cern.ch/record/1712337>.
- [182] M. Wolter, *Tau identification using multivariate techniques in ATLAS*, tech. rep. ATL-PHYS-PROC-2009-016. ATL-COM-PHYS-2008-286, CERN, 2008, URL: <https://cds.cern.ch/record/1152704>.
- [183] W. S. Freund, *An Ensemble of Neural Networks for Online Electron Filtering at the ATLAS Experiment*, tech. rep. ATL-DAQ-PROC-2019-009, CERN, 2019, URL: <https://cds.cern.ch/record/2675025>.
- [184] R. Livneh, “A Neural Network Approach to Muon Triggering in ATLAS”, Presented on 19 Mar 2007, 2007, URL: <https://cds.cern.ch/record/1024579>.
- [185] B. Schlag, “Jet Reconstruction in the ATLAS Level-1 Calorimeter Trigger with Deep Artificial Neural Networks”, Presented 20 Aug 2018, 2018, URL: <https://cds.cern.ch/record/2670301>.
- [186] ATLAS Collaboration, *Generalized Numerical Inversion: A Neural Network Approach to Jet Calibration*, tech. rep. ATL-PHYS-PUB-2018-013, CERN, 2018, URL: <https://cds.cern.ch/record/2630972>.

- [187] G. P. Niedermayer,
“Investigations of Calorimeter Clustering in ATLAS using Machine Learning”,
Presented 19 Dec 2017, 2018, URL: <https://cds.cern.ch/record/2299903>.
- [188] ATLAS Collaboration, *Convolutional Neural Networks with Event Images for Pileup Mitigation with the ATLAS Detector*, tech. rep. ATL-PHYS-PUB-2019-028, CERN, 2019,
URL: <https://cds.cern.ch/record/2684070>.
- [189] B. S. M. Peralva, A. S. Cerqueira, L. M. A. Filho, and J. M. Seixas,
Optimal Signal Selection for a Highly Segmented Calorimeter,
tech. rep. ATL-TILECAL-PROC-2010-002, CERN, 2010,
URL: <https://cds.cern.ch/record/1259887>.
- [190] P. C. Grubitz, “Investigating Recurrent Neural Networks for Applications in Energy Reconstruction at the ATLAS Liquid Argon Calorimeter”,
Master’s thesis: TU Dresden, 2019, URL:
https://iktp.tu-dresden.de/IKTP/pub/19/Grubitz_Masterarbeit.pdf.
- [191] H. D. Block, *The Perceptron: A Model for Brain Functioning*,
Rev. Mod. Phys. **34** (1 1962) 123,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.34.123>.
- [192] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller, “Efficient BackProp”, English,
Neural Networks: Tricks of the Trade, ed. by G. Montavon, G. B. Orr, and K.-R. Müller,
vol. 7700, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012 9,
ISBN: 978-3-642-35288-1.
- [193] C. Igel and M. Hüsken, “Improving the Rprop learning algorithm”,
Second International Symposium on Neural Computation (NC 2000), vol. 2000,
ICSC, 2000 115.
- [194] N. K. Treadgold and T. D. Gedeon,
Simulated annealing and weight decay in adaptive learning: the SARPROP algorithm,
IEEE Transactions on Neural Networks **9** (1998) 662, ISSN: 1045-9227.
- [195] A. E. Bryson and Y.-C. Ho,
Applied Optimal Control: optimization, estimation, and control, revised,
Blaisdell Publishing Company, 1969.
- [196] D. E. Rumelhart, G. E. Hinton, and R. J. Williams,
Learning representations by back-propagating errors, Nature **323** (6088 1986) 533,
ISSN: 1476-4687, URL: <https://doi.org/10.1038/323533a0>.
- [197] J. L. Elman, *Finding Structure in Time*, Cognitive Science **14** (1990) 179, eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1207/s15516709cog1402_1,
URL: https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402_1.

- [198] S. Hochreiter and J. Schmidhuber, *Long Short-Term Memory*, *Neural Computation* **9** (1997) 1735, eprint: <https://doi.org/10.1162/neco.1997.9.8.1735>, URL: <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [199] K. Cho et al., *Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation*, arXiv e-prints, arXiv:1406.1078 (2014) arXiv:1406.1078, arXiv: 1406.1078 [cs.CL].
- [200] G. Weiss, Y. Goldberg, and E. Yahav, *On the Practical Computational Power of Finite Precision RNNs for Language Recognition*, arXiv e-prints, arXiv:1805.04908 (2018) arXiv:1805.04908, arXiv: 1805.04908 [cs.LG].
- [201] J. Chung, Ç. Gülçehre, K. Cho, and Y. Bengio, *Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling*, CoRR **abs/1412.3555** (2014), arXiv: 1412.3555, URL: <http://arxiv.org/abs/1412.3555>.
- [202] S. Bai, J. Z. Kolter, and V. Koltun, *An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling*, arXiv e-prints, arXiv:1803.01271 (2018) arXiv:1803.01271, arXiv: 1803.01271 [cs.LG].
- [203] G. Chevalier, *The LSTM cell.png*, Adapted work. Published under CC BY 4.0, 2018, URL: https://commons.wikimedia.org/w/index.php?title=File:The_LSTM_cell.png&oldid=344169785.
- [204] J. E. Blad, *Gated Recurrent Unit, base type.svg*, Adapted work. Published under CC BY-SA 4.0, 2018, URL: https://commons.wikimedia.org/w/index.php?title=File:Gated_Recurrent_Unit,_base_type.svg&oldid=297295513.
- [205] Y. LeCun et al., *Backpropagation Applied to Handwritten Zip Code Recognition*, *Neural Computation* **1** (1989) 541, eprint: <https://doi.org/10.1162/neco.1989.1.4.541>, URL: <https://doi.org/10.1162/neco.1989.1.4.541>.
- [206] K. J. Lang, A. H. Waibel, and G. E. Hinton, *A time-delay neural network architecture for isolated word recognition*, *Neural Networks* **3** (1990) 23, ISSN: 0893-6080, URL: <http://www.sciencedirect.com/science/article/pii/089360809090044L>.
- [207] E. A. Wan, "Time series prediction by using a connectionist network with internal delay lines", *Santa Fe Institute Studies in the Sciences of Complexity – Proceedings*, ed. by A. Weigend and N. Gershenfeld, vol. 17, Addison-Wesley, 1994 195.
- [208] F. Yu and V. Koltun, *Multi-Scale Context Aggregation by Dilated Convolutions*, arXiv e-prints, arXiv:1511.07122 (2015) arXiv:1511.07122, arXiv: 1511.07122 [cs.CV].

- [209] A. van den Oord et al., *WaveNet: A Generative Model for Raw Audio*, arXiv e-prints, arXiv:1609.03499 (2016) arXiv:1609.03499, arXiv: 1609.03499 [cs.LG].

References for Chapter 6

- [111] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter*, tech. rep. CERN-LHCC-2017-018. ATLAS-TDR-027, CERN, 2017, URL: <http://cds.cern.ch/record/2285582>.
- [136] M. Leyton, “Minimum Bias and Underlying Event Measurements with ATLAS”, *Proceedings, 3rd International Workshop on Multiple Partonic Interactions at the LHC (MPI@LHC 2011): Hamburg, Germany, 21-25 Nov 2011*, 2012 11, arXiv: 1202.2090 [hep-ex].
- [141] W. E. Cleland and E. G. Stern, *Signal processing considerations for liquid ionization calorimeters in a high rate environment*, Nucl. Instrum. Meth. **A338** (1994) 467.
- [142] S. Stärz, “Energy Reconstruction and high-speed Data Transmission with FPGAs for the Upgrade of the ATLAS Liquid Argon Calorimeter at LHC”, Presented 19 May 2015, 2015, URL: <https://cds.cern.ch/record/2030122>.
- [144] J. P. Grohs, “Simulation of the upgraded Phase-1 Trigger Readout Electronics of the Liquid-Argon Calorimeter of the ATLAS Detector at the LHC”, Presented 29 Feb 2016, 2015, URL: <https://cds.cern.ch/record/2135931>.
- [148] S. van der Walt, S. C. Colbert, and G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, Computing in Science Engineering **13** (2011) 22, ISSN: 1521-9615.
- [210] N. Madysa, *Phase-II Studies with AREUS*, version 3.0.0, Access restricted to CERN users, 2019, URL: <https://gitlab.cern.ch/AREUS/AREUS>.
- [211] J. P. Grohs and S. Stärz, *AREUS*, tech. rep. ATL-LARG-INT-2014-003, CERN, 2014, URL: <https://cds.cern.ch/record/1745502>.
- [212] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino, and D. Quarrie, “The Athena Control Framework in Production, New Developments and Lessons Learned”, *Computing in High Energy Physics and Nuclear Physics*, 2005 456, URL: <https://cds.cern.ch/record/865624>.
- [213] J. Vlissides, R. Helm, R. Johnson, and E. Gamma, *Design patterns: Elements of reusable object-oriented software*, Addison-Wesley, 1994, ISBN: 0-201-63361-2.
- [214] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852, arXiv: 0710.3820 [hep-ph].

- [215] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, Comput. Phys. Commun. **134** (2001) 41.
- [216] ATLAS Collaboration, *Physics Analysis Workbook, Release 21*, Access restricted to CERN users, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/PhysicsAnalysisWorkBookRel21>.
- [217] S. Agostinelli et al., *Geant4 – a simulation toolkit*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506** (2003) 250, ISSN: 0168-9002, URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [218] J. Allison et al., *Recent developments in Geant4*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **835** (2016) 186, ISSN: 0168-9002, URL: <http://www.sciencedirect.com/science/article/pii/S0168900216306957>.
- [219] J. P. Grohs, S. Stärz, N. Madysa, and P. Horn, *AreusCaloTools, An ATHENA package for AREUS*, version 1.0.2, Access restricted to CERN users, 2019, URL: <https://gitlab.cern.ch/AREUS/AreusCaloTools>.
- [220] ATLAS Collaboration, *CBNTAA: Athena-Aware combined ntuple*, Access restricted to CERN users, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ComBinedNTupleAthena>.
- [221] ATLAS Collaboration, *CaloD3PD*, Access restricted to CERN users, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/CaloD3PD>.
- [222] H. Ma, *CaloCellDetailsFillerTool.cxx, Excerpt from ATHENA*, URL: <https://gitlab.cern.ch/atlas/athena/blob/release/22.0.6/PhysicsAnalysis/D3PDMaker/CaloD3PDMaker/src/CaloCellDetailsFillerTool.cxx>.
- [223] J. P. Grohs, *CaloCellIdentifier.h, Excerpt from AREUS*, Access restricted to CERN users, URL: <https://gitlab.cern.ch/AREUS/AREUS/blob/3.0.0/Common/include/CaloCellIdentifier.h>.
- [224] ATLAS Collaboration, *Calorimeter Claibration*, Access restricted to CERN users, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/CalorimeterCalibration>.
- [225] L. Courneyea et al., *Computation and validation of the electronic calibration constants for the ATLAS Liquid Argon Calorimeters*, tech. rep. ATL-LARG-INT-2010-007, CERN, 2010, URL: <https://cds.cern.ch/record/1278462>.
- [226] J. Collot, B. Seligman, and R. Schaffer, *LArHit.h, Excerpt from ATHENA*, URL: <https://gitlab.cern.ch/atlas/athena/blob/release/22.0.6/LArCalorimeter/LArSimEvent/LArSimEvent/LArHit.h>.

- [227] G. Parrou, S. Negroni, and G. Unal, *LArBarrelCalculator.cxx, Excerpt from ATHENA*, URL: <https://gitlab.cern.ch/atlas/athena/blob/release/22.0.6/LArCalorimeter/LArG4/LArG4Barrel/src/LArBarrelCalculator.cxx>.
- [228] ATLAS Collaboration, *Further ATLAS tunes of PYTHIA6 and Pythia 8*, tech. rep. ATL-PHYS-PUB-2011-014, CERN, 2011, URL: <https://cds.cern.ch/record/1400677>.
- [229] ATLAS Collaboration, *AMI Query: Tags e1133, s3142, s3144*, Access restricted to CERN users, URL: https://ami.in2p3.fr/app?subapp=tagsShow&userdata=e1133_s3142_s3144.
- [230] ATLAS Collaboration, *AMI Query: mc15_14TeV.119995.Pythia8_A2MSTW2008LO_minbias_inelastic_low.merge.HITS.e1133_s3142_s3144*, Access restricted to CERN users, URL: https://ami.in2p3.fr/AMI/servlet/net.hep.atlas.Database.Bookkeeping.AMI.Servlet.Command?Command=FormGetElementInfo+-elementID=%22810173%22+-entityName=%22dataset%22+-processingStep=%22production%22+-project=%22mc15_001%22&Converter=/xsl/AMIProdFrameHTML.xsl.
- [231] ATLAS Collaboration, *AMI Query: mc15_14TeV.119996.Pythia8_A2MSTW2008LO_minbias_inelastic_high.merge.HITS.e1133_s3142_s3144*, Access restricted to CERN users, URL: https://ami.in2p3.fr/AMI/servlet/net.hep.atlas.Database.Bookkeeping.AMI.Servlet.Command?Command=FormGetElementInfo+-elementID=%22810156%22+-entityName=%22dataset%22+-processingStep=%22production%22+-project=%22mc15_001%22&Converter=/xsl/AMIProdFrameHTML.xsl.
- [232] P. Borgeaud et al., *The LArg Tower Builder Board: calculation, simulation, measurements.*, (2002), URL: <http://cds.cern.ch/record/685385>.
- [233] L. Kurchaninov, *Modeling of the HEC Electronics Chain*, tech. rep. ATLAS-HEC 2019, MPP, 2001, URL: https://wwwatlas.mpp.mpg.de/HEC_notes/hec109/index.html.
- [234] L. W. Nagel and D. Pederson, *SPICE (Simulation Program with Integrated Circuit Emphasis)*, tech. rep. UCB/ERL M382, EECS Department, University of California, Berkeley, 1973, URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/1973/22871.html>.
- [235] N. Seguin-Moreau, *New Preamplifier design*, Access restricted to CERN users, 2015, URL: <https://indico.cern.ch/event/446613/contributions/1939610/>.
- [236] N. Seguin-Moreau, *Update on LAUROC PA/Shaper*, Access restricted to CERN users, 2018, URL: <https://indico.cern.ch/event/759147/contributions/3154593/>.

- [237] H. Abreu et al.,
Performance of the electronic readout of the ATLAS liquid argon calorimeters,
JINST **5** (2010) P09003.
- [238] M. Gardner, *White and brown music, fractal curves and one-over-f fluctuations*,
Scientific american **238** (1978) 16.
- [239] R. Whittle, J. McCartney, and A. Herriman, *DSP generation of Pink (1/f) Noise*, 1999,
URL: <http://www.firstpr.com.au/dsp/pink-noise/>.
- [240] N. Brodu, *Real-time update of multi-fractal analysis on dynamic time series using incremental discrete wavelet transforms*, arXiv preprint (2005),
arXiv: nlin/0511041 [nlin.CD].
- [241] R. F. Voss, "Fractals in nature: From characterization to simulation",
The Science of Fractal Images, ed. by H.-O. Peitgen and D. Saupe,
Springer New York, 1988 21, ISBN: 978-1-4612-3784-6,
URL: https://doi.org/10.1007/978-1-4612-3784-6_1.
- [242] N. C. Beaulieu and C. C. Tan,
Letter FFT Based Generation of Bandlimited Gaussian Noise Variates,
European Transactions on Telecommunications **10** (1999) 545, eprint:
<https://onlinelibrary.wiley.com/doi/pdf/10.1002/ett.4460100510>,
URL:
<https://onlinelibrary.wiley.com/doi/abs/10.1002/ett.4460100510>.
- [243] H. T. Yura and S. G. Hanson, *Digital simulation of an arbitrary stationary stochastic process by spectral representation*, J. Opt. Soc. Am. A **28** (2011) 675,
URL: <http://josaa.osa.org/abstract.cfm?URI=josaa-28-4-675>.
- [244] M. Shinozuka and C.-M. Jan,
Digital simulation of random processes and its applications,
Journal of Sound and Vibration **25** (1972) 111, ISSN: 0022-460X, URL:
<http://www.sciencedirect.com/science/article/pii/0022460X72906001>.
- [245] N. Seguin-Moreau, *PA/Shaper Schedule and LAUROC2 Update*,
Access restricted to CERN users, 2019,
URL: https://indico.cern.ch/event/801784/sessions/304217/attachments/1822744/2985702/LAUROC1_3April2019.pdf.
- [246] F. C. Leone, L. S. Nelson, and R. B. Nottingham, *The Folded Normal Distribution*,
Technometrics **3** (1961) 543, ISSN: 00401706,
URL: <http://www.jstor.org/stable/1266560>.
- [247] N. Madysa,
Neural Networks for Trigger-Level Signal Classification and Energy Reconstruction,
Access restricted to CERN users, 2018,
URL: <https://indico.cern.ch/event/734395/contributions/3043440/>.

- [248] N. Madysa, *Sparse-Representation-Based Filter for the EM CaloTrigger Path*, Access restricted to CERN users, 2018,
URL: <https://indico.cern.ch/event/752186/contributions/3117769/>.

References for Chapter 7

- [111] ATLAS Collaboration,
Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter,
tech. rep. CERN-LHCC-2017-018. ATLAS-TDR-027, CERN, 2017,
URL: <http://cds.cern.ch/record/2285582>.
- [114] G. Aad et al., *ATLAS LAr Calorimeter trigger electronics phase I upgrade: LATOME Firmware Specifications*, tech. rep. ATL-COM-LARG-2017-052, internal note: CERN, 2017, URL: <https://cds.cern.ch/record/2294187>.
- [249] ATLAS Collaboration, *Calibration Runs 322719, 32271 and 3227199*, file directory, Restricted to CERN users, 2018,
URL: <file:///eos/atlas/atlascerngroupdisk/det-larg/ElecCalib/Auto/>.
- [250] L. Serin, *EM Characteristics Values*, Access restricted to CERN users, 2015,
URL: <https://indico.cern.ch/event/464389/contributions/1970407/>.
- [251] J. W. Gibbs, *Fourier's Series*, Nature **59** (1539 1899) 606, ISSN: 1476-4687,
URL: <https://doi.org/10.1038/059200b0>.
- [252] J. W. Gibbs, *Fourier's Series*, Nature **59** (1522 1898) 200, ISSN: 1476-4687,
URL: <https://doi.org/10.1038/059606a0>.
- [253] N. J. Buchanan et al., *ATLAS liquid argon calorimeter front end electronics*, Journal of Instrumentation **3** (2008) P09003,
URL: <https://doi.org/10.1088%2F1748-0221%2F3%2F09%2Fp09003>.
- [254] W. Lampl et al., *Digitization of LAr calorimeter for CSC simulations*, tech. rep. ATL-LARG-PUB-2007-011. ATL-COM-LARG-2007-008, CERN, 2007,
URL: <http://cds.cern.ch/record/1057879>.
- [255] V. Grassi, *Status of LAr HV*, Access restricted to CERN users, 2012,
URL: <https://indico.cern.ch/event/209421/contributions/1502209/>.
- [256] B. Trocme et al., *Status of EMBPS at 1200 V*, Access restricted to CERN users, 2012,
URL: <https://indico.cern.ch/event/208547/contributions/409792/>.

References for Chapter A

- [257] J. P. Rutherford, *Signal Degradation due to Charge Buildup in Noble Liquid Ionization Calorimeters*, (2001), Accepted as Scientific Note, URL: <https://cds.cern.ch/record/493996>.
- [258] L. S. Miller, S. Howe, and W. E. Spear, *Charge Transport in Solid and Liquid Ar, Kr, and Xe*, Phys. Rev. **166** (1968) 871.

-
- [259] K. Yoshino, U. Sowada, and W. F. Schmidt,
Effect of molecular solutes on the electron drift velocity in liquid Ar, Kr, and Xe,
Phys. Rev. **A14** (1976) 438.

Acknowledgments

The work presented in this thesis would not have been possible without the help of many people. I would like to thank them personally and acknowledge their contributions.

Special thanks go to my doctoral advisor, Arno Straessner, who introduced me to the highly interesting topic of liquid-argon calorimetry. He was always available for discussions and always knew the right people to ask about minutiae of the detector and its simulation. Without his guidance since my graduate studies, I would not be the scientist I am today.

I also thank Peter Krieger for his assessment of this thesis and for ensuring that it holds up to a strict scientific standard. I thank Brigitte Vachon and Nicolas Morange for organizing the Phase-II Upgrade simulation and performance effort and for coordinating the multitude of contributions to the technical design report of the Phase-II Upgrade. Their questions, suggestions, and other input has proven invaluable for this work. I'd like to thank Nicolas Morange further for his supervision of my qualification as an ATLAS author.

In a similar vein, I thank Alessandro Ambler and Sana Ketabchi for their contributions to the AREUS project. Their research and implementation of the HEC and FCal simulation resp. has improved a crucial tool for future upgrade studies. I also extend my thanks to Luciano Manhães de Andrade Filho and Bernardo Sotto-Maior Peralva for the intensive discussions on the Wiener filter. These discussions had a major impact on the structure and focal points of this thesis.

Furthermore, I'd like to thank in no particular order: Nathalie Seguin-Moreau for the simulation of the LAUROC chip, without which this work would not have been possible; Laurent Serin for his collection of representative calorimeter cells and their noise information; Sven Menke for providing baseline electronic noise to compare the AREUS simulation with; Pavol Strizenec for providing Run-2 calibration data; Christopher Anelli for his help and input regarding minimum-bias hit samples for upgrade simulation studies; Martin Aleksa for his indefatigable work for the liquid-argon calorimeters project in general and his valuable input for this work in particular; Hucheng Chen and Hervé Deschamps for the fruitful discussions on simulations of the analog electronics; Steffen Stärz for his invaluable knowledge of and unwavering interest in the development of AREUS. I also thank everyone else in the liquid-argon community who has raised important questions or made interesting suggestions for this work whenever intermediary results were presented.

I am deeply grateful to everyone I had the pleasure to work with in Dresden; without Maximilian Hils, Philipp Horn, Anne-Sophie Berthold, Nick Fritzsche, Olga Novgorodova, Rainer Hentges, Martin Serflin, Robert Roß, Andreas Glatte, or Clemens Grubitz, the past years would

not have been the same and work would not have been as fun and productive. It has been a great pleasure and boon to study at the TU Dresden.

Finally, I thank Lisa, Thomas, and my family for their never-ending supply of support, even and especially when things got tough; my friends for keeping an eye on my work–life balance; and Morgan for being my source of happiness for six years and counting.

Versicherung

Hiermit versichere ich, dass ich die vorliegende Arbeit ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe; die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sind als solche kenntlich gemacht.

Die Arbeit wurde bisher weder im Inland noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Die vorliegende Dissertation wurde am Institut für Kern- und Teilchenphysik der Technischen Universität Dresden unter wissenschaftlicher Betreuung von Prof. Dr. Arno Straessner angefertigt.

Es haben keine früheren erfolglosen Promotionsverfahren stattgefunden.

Ich erkenne die Promotionsordnung des Bereichs Mathematik und Naturwissenschaften der der Technischen Universität Dresden vom 23.02.2011 in der Fassung vom 23.05.2018 an.

Dresden, den _____

M. Sc. Nico Madysa