University of Groningen

Online cluster-finding algorithms for the PANDA electromagnetic calorimeter

Tiemens, Marcel

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Tiemens, M. (2017). *Online cluster-finding algorithms for the PANDA electromagnetic calorimeter*. [Thesis fully internal (DIV), University of Groningen]. Rijksuniversiteit Groningen.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

rijksuniversiteit
 groningen

Online Cluster-Finding Algorithms for the $\bar{P}ANDA$ Electromagnetic Calorimeter

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus prof. dr. E. Sterken
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op
maandag 20 november 2017 om 14:30 uur

door

Marcel Tiemens

geboren op 18 januari 1988
te Groningen

Promotor:

Prof. dr. N. Kalantar-Nayestanaki

Copromotores:

Dr. M. Kavatsyuk

Dr. J.G. Messchendorp

Beoordelingscommissie:

Prof. dr. W. Kühn

Prof. dr. H. Löhner

Prof. dr. A. Pellegrino

Contents

1	Introduction	7
1.1	Physics at the Femtometer Scale	8
1.2	Mysteries of QCD: Exotic States of Matter	10
1.3	A Needle in a Haystack	14
1.4	Current Efforts to Find the Needles	16
1.5	Outline of this Thesis	18
2	Theoretical Motivation	19
2.1	Atomic Structure – Quantisation of Energy Levels	20
2.2	Atomic Structure – Splitting of Energy Levels	20
2.3	Positronium	22
2.4	From Electrodynamics to Chromodynamics – The Quarkonium System	24
2.5	Struggles with the Strong Interaction	25
2.5.1	The Strong Coupling Constant α_s	25
2.5.2	Energy Scales in QCD	28
2.5.3	Calculations in QCD – Effective Field Theories	28
2.5.4	Calculations in QCD – Lattice QCD	30
2.5.5	Calculations in QCD – Quantum Numbers	32
2.6	Decay Modes in Heavier Quarkonium Systems	33
2.7	The Hunt for Exotic States of Matter	34
2.8	Experimental Observation of Resonances	35
3	Experimental Facilities	39
3.1	FAIR	40
3.2	HESR	41
3.3	$\bar{P}ANDA$	43
3.3.1	Physics Programme	43
3.3.2	Detector Setup	44

3.3.2.1	Target System	44
3.3.2.2	Micro Vertex Detector	46
3.3.2.3	Straw Tube Tracker	46
3.3.2.4	Cherenkov Light Detector	47
3.3.2.5	Electromagnetic Calorimeter	47
3.3.2.6	Time-of-Flight System	49
3.3.2.7	Solenoid Magnet	49
3.3.2.8	Muon Counters	49
3.3.2.9	Forward Spectrometer	50
4	The EMC: From Particle to Byte	51
4.1	Scintillators	52
4.1.1	General Properties	52
4.1.2	Lead Tungstate	56
4.2	Readout Electronics	58
4.3	Readout of the Detector	59
4.3.1	Digitisation	60
4.3.2	Data Concentration	62
4.3.3	Event Building	63
4.3.4	Synchronisation	67
5	Simulation of the $\bar{P}ANDA$ EMC	69
5.1	Software Package	70
5.2	Event-based Simulation	70
5.2.1	General Description	70
5.2.2	Technical Details	71
5.3	Time-based Simulation	72
5.3.1	Digitisation and Signal Pre-processing	73
5.4	Event Reconstruction	76
6	Cluster Finding	81
6.1	Final-State EM Particle Reconstruction	82
6.2	Cluster Finding - Input	82
6.3	Distributed Cluster-Finding	83
6.3.1	Motivation	83
6.3.2	Precluster Merging	83
6.4	Cluster Finding - Algorithms	85

6.4.1	Description and Parameters	86
6.4.2	Algorithm Implementations	87
6.4.3	Bump Splitting	88
6.5	Performance Assessment by PandaROOT Simulation	89
6.5.1	Parameter Optimisation	89
6.5.2	Performance of the Cluster-Finding Algorithms	97
6.6	Performance Under More Challenging Circumstances	105
6.7	Performance Assessment by Hardware Simulation	108
6.7.1	Performance of the Algorithm on a Hardware Prototype . . .	108
6.7.2	Simulation of the Burst Building Network	110
7	Outlook and Conclusion	113
7.1	Outlook	114
7.2	Conclusion	115
	Summary	123
	Samenvatting	129
	Acknowledgements/Dankwoord	133
	List of Acronyms	137

Chapter 1

Introduction

This chapter will start with a discussion on the subatomic world, and the discoveries that can and have been made. Recent discoveries pave the way to continue this line of research, to unveil the way Mother Nature works at the femtometer scale. To find the proverbial needle in the haystack, several enormous experiments are being and have been set up. The enormous amount of data they collect provides significant challenges, and the experiments use different approaches to tackle these challenges. Three of them, the Compact Muon Solenoid (CMS) and LHCb experiments at the Large Hadron Collider (LHC) at CERN, and the future antiProton ANnihilations at Darmstadt ($\bar{P}ANDA$) experiment at the Facility for Antiproton and Ion Research (FAIR), will be discussed in some more detail.

1.1 Physics at the Femtometer Scale

In 1673, the Dutch scientist Anthonie van Leeuwenhoek made the world's first observation of microcellular organisms, whose existence was not known, or even suspected, before that time. The observation inspired a search to find ever smaller building blocks of the world around us. However, optical microscopes soon reached their limit, and in order to go to smaller scales, new techniques had to be devised to increase the *resolution*: the smallest distance at which two points can be seen as two separate entities. Physically, to extract information about these points, the distance between them should be larger than the wavelength of the light that is used to study them. This means that optical microscopes cannot resolve objects that are smaller than ~ 500 nm.

In 1924, Louis de Broglie proposed that all matter can exhibit wave-like behaviour, following the development of Quantum Mechanics. The famous formula that came with his hypothesis links the particle wavelength λ to its momentum p : $\lambda = h/p$, with h the Planck constant. This finding opened up a new window to investigate nature at the smallest scales. One simply has to increase the particle momentum to decrease the wavelength, thereby increasing the resolution. In the 1950s, the development of synchrotron particle accelerators dawned a new era, hurling subatomic particles at ever larger momenta. By colliding these particles (not necessarily of the same type), physicists were able to probe their internal structure. First, the structure of atomic nuclei could be investigated, and in the 1960s, advances in accelerator science enabled to probe the inner structure of the protons and neutrons. In addition, they discovered a vast variety of particles and found that they can be grouped together by the type of their constituents, the *quarks* — the smallest building blocks known today. Figure 1.1 presents an overview of these and other fundamental building blocks. In the *Standard Model of Particle Physics* (or Standard Model, for short), the theory that was set up to describe the subatomic world, the quarks get their mass by interacting with a particle known as the Higgs boson (Figure 1.1). The Higgs boson was recently discovered at the LHC at CERN, Geneva, Switzerland [1], confirming the theory's prediction. However, no direct predictions for the value of the masses of composite particles are given by the model, and the values found in several measurements give rise to one of the mysteries of the Standard Model. The proton, for example, is made up of three quarks: two *up* quarks and a *down* quark, having masses of $2.2 \text{ MeV}/c^2$ and $4.7 \text{ MeV}/c^2$, respectively [2]. This adds up to $10.1 \text{ MeV}/c^2$, but the proton mass is $938 \text{ MeV}/c^2$, about a factor 100 higher!

Figure 1.1. Schematic overview of the fundamental building blocks of matter. Colour-shaded areas indicate which forces act on which set of particles by the ‘shadow’ casted from the base of the participating blocks. Differently shaded blocks show the family they belong to. Particles of the first family were the first to be discovered and are the only stable particles. The other families can be seen as heavier versions of the first family and are increasingly unstable [2]. The red blocks show the carriers of the different forces: the photon for the electromagnetic force, the W and Z boson for the weak nuclear force, and the gluon for the strong nuclear force. See also section 2.3.

The gained mass is thought to be generated by the interaction between the quarks and gluons, but how this works is poorly understood. It is possible to gain some insight into this mechanism by precisely measuring the energy levels and transitions therein of a specific quark combination. The charmonium meson, composed of a *charm* quark and its antiquark, is a promising candidate for this investigation. The relatively heavy mass of the *charm* quark ($\sim 1.3 \text{ GeV}/c^2$ [2]) allows an essentially non-relativistic approach, while its energy level spectrum (Figure 1.2) can still be calculated perturbatively. See section 2.5 for more details.

Another interesting feature of QCD is known as ‘confinement’. Section 2.5.1 describes this phenomenon in more detail. In essence, it states that quarks cannot be found alone; they group together to form the observable subatomic particles, the *hadrons*. Only specific combinations are allowed, but there is no bound on the number of quarks that can group together. However, the vast majority of observed hadrons consists of two or three quarks. Composite subatomic particles which do not fit this scheme of ‘conventional’ hadrons are referred to as exotic states of matter. Other, even more exotic states are known as glueballs, in which the force carriers

Figure 1.2. The invariant mass spectrum of charmonium, showing results from both potential-model calculations and measured values, indexed by their quantum number (horizontal axis). Notation: $[name](n^{2s+1}L)$ or $[name](mass [MeV/c^2])$, with n the principal quantum number, s the spin of the state, and L the orbital angular momentum in spectroscopic notation. J is omitted in the notation, because it is already shown on the horizontal axis. Unknown charged states are indicated by the letter Z, neutral unknown states with negative parity by Y, and other unknown neutral states by X. Only those states are shown for which the spin-parity has been determined. For a complete list, see the latest PDG tables [2].

of the strong interaction combine together to form a bound state, something that is not possible in e.g. electrodynamics. Observation of and subsequent measurements on these exotic states of matter would provide valuable insight into the theory by e.g. confirming its predictions.

1.2 Mysteries of QCD: Exotic States of Matter

The charmonium meson is being studied extensively by various experiments, and its energy level spectrum is shown in Figure 1.2. The precise features will be addressed in chapter 2. For this discussion, note only the coloured boxes. Below the open-charm threshold¹ at $3.73 \text{ GeV}/c^2$, indicated by a dashed line in Figure 1.2, the yellow boxes, which indicate states that have been both predicted and measured, show near-perfect agreement between theory and experiment. Above the threshold, the story changes. The gray boxes show that there are many states predicted to exist, but those have

¹Below this energy threshold, only ‘pure’ charmonium states (i.e. $c\bar{c}$) can be formed, while above it, in addition, also D meson pairs can be created, which consist of a charm quark and a lighter quark. Consequently, these have nonzero charm — an additional quantum number that was introduced to indicate the quark content of particles. c quarks contribute +1 to this number; $\bar{c}$ quarks add -1.

Figure 1.3. (left) Decay mode of the Y(4260) state to a J/ψ via the $Z_c(3900)$. (right) Observation of the $Z_c(3900)$ signal in the combination of the J/ψ with a charged pion in the BESIII experiment [9].

not (yet) been observed. Even more interesting, the red boxes show states that were observed, but not predicted. It is important to note that the predictions do not take the existence of exotic states into account.

The observation of the ‘red’ states in Figure 1.2 prompted a search for the nature of these particles — could they be the exotic states of matter? A plethora of theoretical interpretations have been proposed. The interpretations need experimentally measured quantities to verify them. As an example, the Y(4260) state will be discussed. To find out more about this particle, its decay modes were measured, for example by tuning the BEPCII e^+e^- collider [55] to a centre-of-mass energy of 4,260 MeV/c², and by measuring the decay products of the Y(4260) with the BESIII detector [3]. In particular, it was checked how strongly it couples to the known charmonium states, i.e. if it can decay directly to these states, or if there are intermediate states involved. Figure 1.3 (left) shows one interesting decay mode via a particle with a mass around 3,900 MeV/c², the recently observed $Z_c(3900)$, showing up with a statistical significance of more than 8σ in Figure 1.3 (right). This decay was seen in several experiments [4–7]. The most interesting feature of this decay mode is that the $Z_c(3900)$ decays via a *charged* pion to the J/ψ state, which, because it is a $c\bar{c}$ state, has no charge. This means that the $Z_c(3900)$ must also be charged, which in turn means that it cannot be a ‘pure’ charmonium state. The $Z_c(3900)$ must therefore consist of at least four quarks. Assuming it has four quarks, there are several interpretations (see Figure 2.3), two of which will be explored below:

1. A tetraquark, which is a ‘true’ bound state of four quarks;
2. A $D^0\bar{D}^*$ molecule, which has two D mesons that are bound together by pion

exchange (like in a nucleus).

These two interpretations can be distinguished experimentally by looking at the decay products of the $Z_c(3900)$. For the molecule, the most likely decay mode would be the breaking of the bond between the D mesons [8], i.e.

$$Z_c(3900) \rightarrow D^0 \bar{D}^*.$$

The tetraquark state, on the other hand, would favour decays to the J/ψ meson by emitting a charged pion. The width² of the $Z_c(3900)$ in this channel was found to be $\Gamma = (46 \pm 10 \pm 20) \text{ MeV}/c^2$ [9]. When looking at the mass spectrum of a reconstructed pair of D^0 and $\bar{D}^*$ mesons, the BESIII collaboration found a peak at $3,885 \text{ MeV}/c^2$ [10], with a width of $\Gamma = (24.8 \pm 3.3 \pm 11) \text{ MeV}$. Because both the mass and the width are close to the $Z_c(3900)$, it is assumed to be the same particle. The ratio of the *partial* widths of these two modes provides a means of identifying the most probable interpretation. The width, Γ , like the ones quoted before, is the sum of the partial widths, $\partial\Gamma$, which reflect some particular decay of the state, e.g. $Z_c(3900) \rightarrow D\bar{D}^*$. The partial width is calculated from the total width and the branching fraction (section 2.8). From these measurements, the ratio of the partial widths comes out at [10]:

$$\frac{\partial\Gamma(Z_c(3900) \rightarrow D^0 \bar{D}^*)}{\partial\Gamma(Z_c(3900) \rightarrow \pi J/\psi)} = 6.2 \pm 1.1 \pm 2.8,$$

which seems to favour the molecule description. To make a more definitive statement, more research is needed, as this quantity by itself cannot rule out other interpretations. However, this example demonstrates how experimentally measured quantities can help guide the interpretations of (possibly) exotic states. As the $Y(4260)$ state features a decay mode to the $Z_c(3900)$, which was just argued to be an exotic hadron, it is likely that the $Y(4260)$ itself is also an exotic state.

Another question of interest regards the existence of isospin partners of the $Z_c(3900)$, i.e. particles with the same mass, but a different charge. Historically, isospin was conceived by thinking of such a set of particles as being different manifestations of the same particle. This proved especially useful in constructing the theory of the strong interaction (section 2.4), where the strong force would not discriminate between isospin partners. Hence, the existence of isospin partners

²The *width* of a state is inversely proportional to its mean life time, indicating the stability of a particle state.

provides valuable insight in the coupling of the strong force to the particle. The $Z_c(3900)$ was found to be part of an isospin triplet [7]. For its exotic status, and the research that was done to identify its properties, the $Z_c(3900)$ was quoted as one of the physics highlights of 2013 by the American Physical Society (APS) [11]. The existence of the unknown states shows that there is still much to be learned in the field of hadron physics and the strong interaction. In particular, the $X(3872)$, which was the first unexpected state to be observed [12], is still not understood. It is known to have a very narrow width ($\sim 1.2 \text{ MeV}/c^2$) and, from the fact that the branching fractions B (see section 2.8), $B(X(3872) \rightarrow \rho J/\psi) \approx B(X(3872) \rightarrow \omega J/\psi)$, it has a large probability for isospin breaking; the ρ and ω mesons have different isospin numbers, so the branching fractions, which indicate how strongly the $X(3872)$ couples to the ρ and ω , are not expected to be the same. Furthermore, a recent investigation by the BESIII collaboration showed that radiative transitions from the $Y(4260)$ to the $X(3872)$ exist [13], implying that the structure of the two particles might have a common nature.

Several experiments are running or are being developed to explore these exciting features, as to ultimately determine how matter is formed from quarks and gluons in QCD. To help shed light on the unknown states, the $\bar{P}ANDA$ experiment is currently being developed by a collaboration of about 500 scientists from 17 countries, which will employ proton-antiproton annihilations at energies in the charmonium mass region to create particles of interest. The use of antiprotons relaxes the requirements on the magnitude of the beam momentum, as the annihilation energy is available for particle production. With the (anti)proton being a composite particle, nearly all charmonium states can be directly and copiously populated. In contrast, most of the currently running experiments can either only directly populate states with the same quantum numbers as the e^+e^- system, like BESIII [3], CLEO-c [7], and BELLE [14] (which does produce a very ‘clean’ signal, i.e. little background), or rely on decays of other particles, and cannot be tuned to the energy of the exotic states in the charmonium energy range, e.g. CMS [16].

The $\bar{P}ANDA$ collaboration aims to research key items in the strong and weak interaction, including the topics described above. The $\bar{P}ANDA$ experimental apparatus will be installed at the Facility for Antiproton and Ion Research (FAIR) in Darmstadt, Germany. Replacing the proton target with a target of heavier elements enables the study of exotic states in a dense environment, providing information to understand its

nature. The $\bar{P}ANDA$ detector will be able to perform high-precision measurements on the produced states. Chapter 3 provides a more comprehensive description of the experiment and the facility that houses it. The present project is centered around the $\bar{P}ANDA$ experiment and aims to further its development.

1.3 A Needle in a Haystack

To accurately measure the properties of the currently discovered unexpected states, a large number of them needs to be produced. States that have not yet been observed will likely have a small production cross section, so a large number of events needs to be generated in order to create a detectable signal above the background. The $\bar{P}ANDA$ experiment is able to produce such states. However, the fact that the $p\bar{p}$ collisions are able to produce a plethora of particles in the charmonium mass region is both advantageous and disadvantageous. The disadvantage is that there will be a very large hadronic background, as only a small portion of the produced particles will be interesting to study. To ensure a sufficient production of the states of interest, the experiment will run at a high interaction rate: about $2 \cdot 10^7$ annihilations per second. This is considerably higher than e.g. the $\sim 4 \cdot 10^3$ interactions/s in the BESIII experiment, or even the $\sim 2 \cdot 10^5$ annihilations/s in the BELLE experiment [14]. About 2,400 Gbps (gigabit per second) of data is expected to be generated — too much to be stored.

Historically, data reduction has been achieved by using a hardware trigger, which would allow the storage of data if a certain condition is fulfilled. This is needed, because:

- The amount of data generated is simply too much to be able to store.
- The majority of produced data does not contain the type of events that are being looked for.

If e.g. a process like $X \rightarrow 4\gamma$ would warrant studying for some particle X , only events that have at least four photons would be tagged by the trigger for storage. In experiments like CMS, even though the interaction rate is far higher — reaching up to $\sim 10^9$ /s — the hardware trigger offers a sufficient reduction factor (see section 1.4). For the $\bar{P}ANDA$ experiment, the story is different. The $Y(4260)$ example will be used to illustrate the challenge. A prominent decay mode for this state is

$$p\bar{p} \rightarrow Y(4260) \rightarrow J/\psi \pi^0 \pi^0 \rightarrow e^+ e^- + 4\gamma.$$

The hardware trigger would then be set to require the detection of four photons and an e^+e^- pair. However, conventional hardware triggers were based on primitive conditions, such as detector multiplicities. With such conditions, it would be difficult to distinguish electrons from photons by their energy depositions in the calorimeter. Hence, to make the distinction, information of other detector systems is already required at this stage. For example, of the minimum of six energy depositions in the calorimeter, two of which would need to lie in the same solid-angle region as a reconstructed track in the tracking detector. Even this extended triggering scheme would not be able to remove background channels like

$$\begin{aligned} p\bar{p} &\rightarrow J/\psi \pi^0 \eta \rightarrow e^+e^- + 4\gamma, \\ p\bar{p} &\rightarrow J/\psi \eta \eta \rightarrow e^+e^- + 4\gamma, \\ p\bar{p} &\rightarrow J/\psi \pi^+ \pi^- \pi^0 \pi^0 \rightarrow e^+e^- + 4\gamma, \end{aligned}$$

and, indeed, any process which produces an e^+e^- pair and 4 photons. In the third case, the $\pi^+\pi^-$ pair was of low energy and eluded detection. As these background channels do not have to proceed via the production of the $Y(4260)$, the combined cross section³ is much higher for these processes: 282 μb versus 77 pb [15]. Hence, the hardware trigger will not be able to achieve a satisfactory data reduction factor for the study of this state. This is a general issue in the $\bar{P}ANDA$ experiment and, therefore, a new data readout scheme has been proposed. Looking back at the example, requiring the $e^+e^- + 4\gamma$ combination to peak around 4,260 MeV/c² would eliminate the channel with the charged pions, and requiring two-photon combinations to originate from a π^0 takes care of the other background channels mentioned above. Hence, by using advanced physics-inspired properties, a sufficient data reduction factor should be obtainable. This requires that almost complete event information is reconstructed in real time, putting constraints on the readout hardware of the detector. In section 4.3, specifics will be addressed on the implementation of this concept, called *triggerless readout*. In this concept, all hits are stored until the event selection takes place, and the detector subsystems prepare the data for this event selection. For example, an algorithm to identify groups of hits in the electromagnetic calorimeter corresponding to e.g. a photon will run online in the early stages of the readout system, providing vital

³Obtained using the DPM generator (see section 5.2.2 for a description), by considering the number of events out of 10^7 events with exactly two neutral and two charged pions, which was taken to be the most realistic alternative, as no charmonium states or particle decays are simulated by the generator. This can be justified by considering that, without particle identification techniques, the charged pions could easily be mistaken for electrons and positrons. The total inelastic $p\bar{p}$ cross section at 4.260 GeV/c² was extracted from the PDG plots [2].

Figure 1.4. Schematic perspective view of the CMS detector, located at the Large Hadron Collider at CERN, Geneva, Switzerland [16].

input for the event selection process. Incidentally, options for this algorithm will be discussed in chapter 6 of this work.

1.4 Current Efforts to Find the Needles

There are several experiments currently looking for their favourite needles, and they need to sift through vast amounts of data to find them. The CMS experiment mentioned above was one of the experiments at CERN that first observed the Higgs boson. The detector fully and symmetrically encloses the interaction point for maximal coverage, see Figure 1.4. To reduce the amount of data, CMS employs a trigger system that discards data with a reduction factor of about 10^6 [16]. The trigger operates in two steps: the Level-1 (L1) trigger, which uses information from the calorimeters and the muon system to discard events, and the High-Level Trigger (HLT), which has access to the complete data and uses a computing ‘farm’ to perform complex calculations similar to those made in the offline analysis, if required for specific interesting events. The L1 trigger works in three hierarchies. The first is based on energy depositions in the calorimeters and identified muon tracks or track segments. The next stage combines the available information and builds a ranking of detected particles, like electron and muon candidates, using a pattern logic. The final stage determines

Figure 1.5. Schematic side view of the LHCb detector, located at the Large Hadron Collider at CERN, Geneva, Switzerland [17].

the highest ranked objects across the experiment and uses them to determine if the data should be stored. After the L1 trigger, the amount of data has been reduced by a factor $\sim 10^4$. The HLT optionally sets to work after this, and performs a similar function as the triggerless readout discussed in the previous section. The difference is that the HLT takes the output of the L1 trigger, and builds events from that data, i.e. at a reduced rate, while the $\bar{P}ANDA$ detector has no L1 trigger, and has to build events at the full interaction rate. In the triggerless readout concept, as explained in more detail in section 4.3, the different stages of the readout system pre-build event information to ease the load on the final event-building system. This kind of processing, done at this level, is unique to the $\bar{P}ANDA$ experiment.

Another experiment, which needs to reduce the amount of data generated, is the LHCb experiment at CERN [17]. The design is radically different from the CMS detector, as the spectrometer encompasses only a single jet of particles (Figure 1.5). This is sufficient for the LHCb physics program, which focusses on B mesons (b -quark systems), as these are produced in a tight cone (because of the high energies and head-on collisions at the LHC), extending symmetrically in both ways along the beam direction. This means that the detector will only observe half of the produced events, but to extend it is unattainable from both a financial and a construction point of view. To achieve a sufficient data reduction factor, the LHCb collaboration uses

a two-step approach, similar to what is used in the CMS experiment. The first, the Level-0 (L0) trigger, locates the hadron, electron, and photon candidates with the largest transverse⁴ momentum, p_T , in the calorimeters. This is because B mesons are expected to produce particles with a large transverse momentum, due to their large mass. From the muon chambers, the muons with the largest p_T are also selected by the trigger. This allows to reduce the data rate by a factor ~ 40 . The second trigger level, also a HLT, makes use of the complete data. It verifies that the data it received from the L0 trigger contain only high- p_T particle candidates, and performs again a similar processing as is foreseen in the triggerless readout concept.

1.5 Outline of this Thesis

After the theoretical motivation has been provided in chapter 2, and the experimental setup is described in chapters 3 and 4, the details of my specific research project will be described in the remaining part. The primary objective is to develop an algorithm that can search for groups of hits in real-time in one of the detector subsystems, the electromagnetic calorimeter (EMC). The EMC's data processing scheme will serve as an example for that of the other subdetectors and it can be used to reconstruct certain decay chains independent of the rest of the detector (see, for example, section 5.4). The data processing scheme is discussed in section 4.3. Searching for groups of hits may sound trivial, but the conditions under which the algorithm must perform this task, as a result of the high interaction rate, complicate matters significantly. After testing the algorithm using simulations in section 6.5, a 'real' test using a detector readout prototype will be evaluated in section 6.7.1. In addition, two possible topologies of the data acquisition system will be discussed in sections 4.3.3 and 6.7.2. Using the aforementioned algorithm and a simulated data set allows to determine the requirements of this network. The hardware and firmware for this network is currently being developed, and will be tested in the near future. In chapter 7, a summary of the findings, recommendations, and points of improvement will be presented, alongside an outlook with suggestions on how to proceed in the future.

⁴Transverse, i.e. perpendicular, to the beam momentum.

Chapter 2

Theoretical Motivation

To set the stage for some of the main motivations behind this line of research, this chapter will start with a basic review on charmonium, starting from Atomic Physics and Quantum Electrodynamics, on which the theory of the strong interaction is built. A brief overview of several theoretical approaches will be presented, along with some results from theoretical calculations, that are compared to experimental measurements. The theoretical models described in this chapter usually require experimental measurements to fix some parameters, before predictions can be made. These predictions can, then, be tested against experimental observations. In the field of particle physics, theory and experiment depend on each other for guidance and interpretations. More comprehensive reviews on this subject are, for example, given in [18, 19].

The general notation that will be used throughout this chapter is that vectors will be indicated by **bold print**. Furthermore, the use of natural units, i.e. the speed of light c , Planck's reduced constant $\hbar = h/2\pi$, and the electric permittivity of vacuum ϵ_0 all set to 1, will be implied, unless stated otherwise.

Figure 2.1. The existence of specific and complementary lines in the absorption and emission spectra of the elements (in this case, hydrogen) strongly hinted at the quantisation of electron orbits [20].

2.1 Atomic Structure – Quantisation of Energy Levels

In 1909, Hans Geiger and Ernest Marsden, under the theoretical guidance of Ernest Rutherford, unintentionally demonstrated that atoms have a compact, positively charged nucleus, by scattering alpha particles off a metal foil. To make the atoms electrically neutral, they are surrounded by negatively charged electrons orbiting the nucleus. However, it was known that charges in curved paths emit electromagnetic radiation. To explain the stability of atoms, Niels Bohr proposed in 1913 that electrons can only occupy certain discrete orbits around the nucleus, and could only jump between them by absorbing or emitting a quantum of energy in the form of electromagnetic radiation. This also explained the observation of spectral absorption and emission lines in the visible light spectrum of e.g. hydrogen (Figure 2.1).

2.2 Atomic Structure – Splitting of Energy Levels

The electrons orbit the nucleus at a speed v that is close to the speed of light, c . They do so in the presence of the electric field $\mathbf{E}$ of the nucleus. Using the Lorentz transformation from special relativity, this electric field is seen by the electron as a magnetic field $\mathbf{B}$ (in natural units):

$$\mathbf{B} = -\mathbf{v} \times \mathbf{E}. \quad (2.1)$$

Using the electric potential, V , associated with the electric field, this can be rewritten to

$$\mathbf{B} = \frac{1}{m_e e} \left(\frac{1}{r} \frac{\partial V}{\partial r} \right) \mathbf{r} \times m_e \mathbf{v} = \frac{1}{m_e e} \left(\frac{1}{r} \frac{\partial V}{\partial r} \right) \mathbf{L}, \quad (2.2)$$

with m_e and e the electron mass and charge, r the radial distance, and $\mathbf{L} = \mathbf{r} \times m_e \mathbf{v}$ the orbital angular momentum. The dynamics of the electron follows from the interaction of the electron's intrinsic magnetic moment $\boldsymbol{\mu} = -g_s \mu_B \mathbf{S}$ with this magnetic field.

The factor $g_s \simeq 2$ is called the anomalous magnetic moment, $\mathbf{S}$ is the electron's spin, and $\mu_B = e/2m_e$ is the Bohr magneton. The interaction is given by the Hamiltonian:

$$H = -\boldsymbol{\mu} \cdot \mathbf{B} = a\mathbf{S} \cdot \mathbf{L}, \quad (2.3)$$

where all constants and the contribution of the potential have been absorbed into a . The dot product between $\mathbf{S}$ and $\mathbf{L}$ shows that there is an interaction between the electron spin and the angular momentum, called *spin-orbit coupling* or $L - S$ coupling. This interaction gives rise to a splitting of the energy levels, which can be shown to be:

$$\Delta E_{s-o} = \frac{\alpha^2}{nl(l+1)} E(n), \quad (2.4)$$

with α the fine-structure constant (see the last paragraph in Section 2.3 for its significance), and l the angular momentum quantum number. Because v is close to c , relativistic corrections to the electron's momentum should also be taken into account. The combination of these effects leads to a fine structure in the energy levels.

There is another contributor to the energy levels. Reviewing the discussion above, the electron and the nucleus both have a magnetic moment. According to classical electrodynamics, a magnetic dipole moment $\boldsymbol{\mu}$ gives rise to a magnetic field (in natural units)

$$\mathbf{B}_{dip}(\mathbf{r}) = \frac{1}{4\pi r^3} (3(\boldsymbol{\mu} \cdot \hat{\mathbf{r}})\hat{\mathbf{r}} - \boldsymbol{\mu}), \quad (2.5)$$

at some position $\mathbf{r}$, with $\hat{\mathbf{r}}$ the unit vector pointing radially outwards from the source. Hence, like before, the magnetic moment of one particle, let's call it $\boldsymbol{\mu}_1$, will respond to the magnetic dipole field, $\mathbf{B}_2$, created by the other. This interaction is, again, given by the Hamiltonian

$$\begin{aligned} H &= \boldsymbol{\mu}_1 \cdot \mathbf{B}_2 \\ &= -g_s \mu_B \mathbf{S}_1 \left[\frac{3(-g_s \mu_B \mathbf{S}_2 \cdot \hat{\mathbf{r}})\hat{\mathbf{r}} + g_s \mu_B \mathbf{S}_2}{4\pi r^3} \right] \\ &= \frac{(g_s \mu_B)^2}{4\pi r^3} \left[3(\mathbf{S}_2 \cdot \hat{\mathbf{r}})(\mathbf{S}_1 \cdot \hat{\mathbf{r}}) - \mathbf{S}_1 \cdot \mathbf{S}_2 \right], \end{aligned} \quad (2.6)$$

in which there appear terms of the form $\mathbf{S}_1 \cdot \mathbf{S}_2$ – the *spin-spin interaction*, leading to the hyperfine structure in the energy levels of the system. Historically, a notation to indicate the angular momentum level employed the use of letters. This notation, called spectroscopic notation, is still used today to indicate excitations in angular momentum. In this scheme, $l = 0$ is designated by the letter S , $l = 1$ by P , and $l = 2$

by D , and so on. This leads to the notation used in Figure 1.2.

2.3 Positronium

Although the preceding discussion revolved around atoms, the formalism can be readily adapted to a system where two particles of the same (absolute) charge orbit each other. An example of such a system is positronium (Ps), the bound state of an electron and a positron. This section will discuss how the energy levels in such a system are characterised. The electron and positron are so-called fermions, and are in addition elementary particles¹, meaning that their spin quantum number can assume exactly two values. Because of this, spin can be thought of as ‘pointing’ up or down for such particles. This analogy will be used in the remaining text. For a combination of two particles, their spins can both be pointing in the same direction (‘aligned’) or opposite to each other (‘anti-aligned’). Because of interactions between the spins, the lowest-energy possibility is to have them anti-aligned.

The energy levels, or particle ‘states’ in the spectrum can be calculated using Quantum Electrodynamics (QED), the theory of the electromagnetic interaction. The basic structure can be obtained from the non-relativistic Hamiltonian H , a construction that describes the motion of particles. The general form is $H = T + V$, with T the kinetic energy of the particle(s) and V the potential in which they move, which gives rise to a force on the particle(s). Since the Ps system consists of two particles of equal mass rotating around each other, the Hamiltonian is most conveniently expressed in the center-of-mass system using spherical coordinates. V is given by the Coulomb potential $-e^2/r$, with e the electric charge and r the distance from the centre-of-mass. To calculate the energy levels, it is necessary to find the negative eigenvalues of the energy E from the equation of motion, the Schrödinger equation:

$$\hat{H}\phi = \left(\frac{\hat{\mathbf{p}}_r^2}{m_e} + \frac{\hat{\mathbf{L}}^2}{m_e r^2} - \frac{e^2}{r} \right) \phi = E\phi. \quad (2.7)$$

ϕ is the wavefunction of the state, which describes its position and momenta, $\hat{\mathbf{p}}_r^2$ is the radial-pointing-momentum operator, $\hat{\mathbf{L}}^2$ is the angular momentum operator, and m_e is the mass of the electron. After a bit of algebra, the solution reads: $E_n = R/n^2$,

¹Particles are called elementary when they possess no internal structure (to the best of current knowledge), i.e. they are themselves not made up of other particles.

Figure 2.2. A few possible Feynman diagrams of the electromagnetic interaction between the electron and the positron. Solid straight lines represent the path of a fermion, solid wiggly lines the path of a force carrier of the electroweak interaction (W and Z bosons, photons). (a) Single photon exchange. (b) Situation (a), but including an e^+e^- loop. In QED, such pairs are allowed to come into existence for a time Δt dictated by the Heisenberg uncertainty relation: $\Delta E \Delta t \leq \hbar$, with ΔE the energy of the ‘virtual’ particles. This process is also called ‘vacuum polarisation’. (c) Situation (a), but including a loop where the electron emits a photon and catches it again. This is called ‘self energy’ in QED.

with R the Rydberg constant. To get more accurate and realistic values for the energy levels, relativistic effects need to be taken into account, as well as $L - S$ coupling, and so-called ‘loop corrections’ (see Figure 2.2b and c). In the Standard Model, and hence in QED, particles interact by the exchange of other particles, the mediators of the force (Figure 1.1). For QED, these are the mediators of the electromagnetic interaction, the photons. The process is depicted in Figure 2.2a. A more extensive description of the calculated corrections can be found in [23], and references therein.

The strength of the electromagnetic interaction is given by the dimensionless fine-structure constant α , which equals $\frac{e^2}{4\pi} \approx \frac{1}{137}$ in natural units. For an illustrative example of this, consider two electric charges of charge e , a distance d apart. Following the above, they will interact by exchanging a photon. The electrostatic repulsion is given by $\frac{e^2}{4\pi\epsilon_0 d}$. For the photon to be exchangeable, its wavelength must fit exactly between the charges: $\lambda = \frac{2\pi d}{i}$, with i some integer. To set the normalisation scale, take the photon with the lowest energy, so $i = 1$. The ratio between the energy needed to overcome the electrostatic repulsion and the energy of the photon is $\frac{e^2}{4\pi\epsilon_0 d} / \frac{hc}{\lambda} = \frac{e^2}{4\pi\epsilon_0 d} \times \frac{2\pi d}{hc} = \frac{e^2}{4\pi\epsilon_0 \hbar c} = \alpha$. From this, it follows that α , also called the coupling constant, appears in a calculation everytime when two electromagnetic particles interact. In the diagrams in Figure 2.2, α will hence appear at every vertex, and as $\alpha < 1$ it follows that diagrams (b) and (c) will be suppressed with respect to (a). This is the basis behind perturbative calculations in QED: although it is possible (and even necessary) to extend the diagrams in Figure 2.2 with more and more loops, each added loop will contribute less and less to the process. This means that with a small number of corrections, a decent estimate of the quantity of interest can be computed. To get

a more precise number, more loops should be added to the calculation. It should be noted, however, that this is not necessarily a trivial matter. This also forms the basis of what is called perturbation theory, where some quantity that is difficult or impossible to calculate exactly, can be approximated from a Taylor-series-like expansion of the Lagrangian, $\mathcal{L} = T - V$, around a relevant parameter that is much smaller than 1 (in this case, α).

2.4 From Electrodynamics to Chromodynamics – The Quarkonium System

Analogously to positronium, it is also possible to create bound states of quarks, called *quarkonium*. Like positronium, such a bound state is realised by combining a quark with an antiquark. There are, however, some notable differences. Quarks are found to not only carry an electric charge, but also a colour charge. The new quantum number colour, unique to particles affected by the strong interaction, was introduced to explain the structure of hadrons that were observed experimentally: no free quarks have ever been observed, and the Ω^- baryon, with three s quarks, would have two of them with identical quantum numbers, which is not allowed, since they are fermions. A most striking feature of the colour quantum number, that follows from the former observation, is that it must equal zero for all observable particles. Since quarks themselves have a non-zero colour, this implies that they may only be observed as hadrons, where the total colour charge of the hadron is zero, also called ‘white’, or, more specifically, according to the underlying symmetry group, $SU(3)$, colour singlets. As there are 6 values of the colour charge (red, green, and blue, and their corresponding anticolours), white hadrons can be constructed from a colour-anticolour combination or by combining three colours. The first possibility is called a *meson* – it follows that all quarkonium systems are mesons. The latter combination yields a *baryon*. Combinations of more quarks are allowed, but no experimental evidence was found for their existence until recently [4, 5] (see also section 1.2).

Another important difference with QED is that the force carriers of the strong interaction, the gluons, have non-zero colour themselves. This is in sharp contrast to the force carriers of the electromagnetic interaction, the photons, which do have zero electric charge, and therefore cannot directly couple to each other or themselves. Since gluons have non-zero colour charge, they can in principle combine with quark

Figure 2.3. Possible interpretations of the newly-observed unexpected states. $q(\bar{q})$ is an *up* or a *down* (anti)quark, lines with spiral loops represent gluons. (a) A hybrid meson, which has an on-shell gluon component. (b) A hadronic molecule, a loosely bound states of two mesons, e.g. $D - \bar{D}^*$, bound by pion exchange. (c) A tetraquark (diquark-antidiquark) state. (d) Hadrocharmonium, a charmonium state folded inside a system of light quarks in an atom-like configuration.

systems to form a bound quark-antiquark-gluon system, called a hybrid meson (see Figure 2.3a), or even link together with other gluons to form a bound state consisting purely of gluons, known as a glueball. Such states are called *exotic states* of matter, and although they are allowed to exist in QCD, they have yet to be unambiguously observed in experiments.

The phenomenon of requiring zero total colour is known as colour confinement, see section 2.5.1. An important consequence of colour confinement is that it is impossible to make e.g. a quark beam, like an electron beam can be created. The shy nature of quarks also has very profound consequences for the interactions, decays, and energy levels in (bound) quark systems, although many similarities to their lepton sisters still remain, revealing a more general underlying structure from bound fermion-antifermion systems, see e.g. Figure 2.4. Some of the properties of the quarkonium systems will be explored in the next sections.

2.5 Struggles with the Strong Interaction

2.5.1 The Strong Coupling Constant α_s

Analogously to the QED case, the strong coupling constant α_s indicates the strength of the strong interaction. It is not unreasonable to suspect that the non-zero colour

Figure 2.4. The experimentally-determined energy levels of positronium, charmonium, and bottomonium [2]. Horizontal dashed lines indicate the open charm/bottom thresholds. The general structure of the level diagrams is very similar for the different systems, supporting the claims about their origin made in this chapter.

charge of the gluons might affect the behaviour of α_s . To investigate the strong coupling, the so-called β function is defined:

$$\beta(g) = \mu \frac{\partial}{\partial \mu} g(\mu), \quad (2.8)$$

where g is the colour charge, which determines the coupling constant in the same way as in QED: $\alpha_s = \frac{g^2}{4\pi}$, and μ is the energy scale at which the coupling is evaluated. As β is proportional to μ , it is indicative of the energy associated with the coupling at a certain energy scale, and it is reasonable to let $\beta = 0$ at $g = 0$. There are three possibilities for the evolution of β : always 0, increasing to positive values with g , or increasing to negative values with g . The second option implies that g increases with energy, meaning that it gets weaker at low energy. This leads to ‘screening’². The third option facilitates the opposite, leading to ‘anti-screening’. A calculation with corrections with a single loop (like in Figure 2.2) will already be able to discriminate between these options, by expanding β around its origin using perturbation theory. The result from this calculation yields $\beta = \frac{g^3}{16\pi^2}(-\frac{11}{3}N_c + \frac{2}{3}N_f)$, where the third power in g comes from the fact that the loop diagrams have three vertices, and N_c and N_f are the number of colours and fermions subject to the strong force, respectively.

²In QED, a charge in space can be surrounded by virtual e^+e^- pairs that will align themselves in the electric field of the charge as to shield it. This reduces the effective charge, and leads to the electric charge e , and with it, the EM coupling constant α , to become smaller at lower energy (i.e. larger distances). This effect is called ‘screening’.

Figure 2.5. (left) Experimental extractions of the strong coupling constant α_s at different momentum-transfer (Q) values [24]. (right) Combined results from theoretical calculations (yellow band) and some measurements (coloured shapes) for α_s at different Q values [25].

With three colours and maximally six quark flavours to take into consideration, $\beta < 0$, implying that there is an anti-screening effect, and the coupling constant α_s depends on the energy scale. This is called the ‘running’ of the coupling constant. It turns out that the coupling constant satisfies

$$\alpha_s(q^2) = \frac{4\pi}{(11 - \frac{2}{3}N_f) \ln(q^2/\Lambda^2)}, \quad (2.9)$$

where q is the momentum transfer between the two quarks, and Λ , defined by $\ln \Lambda^2 = \ln \mu^2 - (\frac{g^2}{16\pi^2}(11 - \frac{2}{3}N_f))^{-1}$, sets the relative energy scale at which α_s is to be evaluated. q is larger when the spatial distance between quarks is small, and in this limit, Eq. 2.9 shows that α_s tends to zero, implying that quarks can move around as almost free particles when they are close.

In the limit where q is small (and the spatial distance is large), α_s increases rapidly, creating a force that pulls the quarks ever stronger together. A mechanical analogy would be a system of masses connected by a spring. As the distance grows, the potential energy in the system increases. At some point, the energy exceeds the quark pair-production threshold $E > 2m_q c^2$, and a new ‘real’³ quark-antiquark pair is created. For this reason, it is impossible to separate quarks to make free-particle states, leading to the confinement feature of QCD. Some measurements from which the coupling constant is obtained at different energy scales are shown in Figure 2.5.

³Also called ‘on-shell’. In contrast, the virtual $q\bar{q}$ pairs that are created in vacuum fluctuations are called ‘off-shell’, because they can never be directly measured.

2.5.2 Energy Scales in QCD

The preceding discussion revealed that there are different energy scales in QCD, at which the different features emerge. It makes sense to identify an energy scale, below which α_s becomes definitely too large for perturbation theory to apply, and thus confinement becomes predominant. This occurs at the characteristic QCD confinement scale $\Lambda_{QCD} \approx 200$ MeV. To perform calculations on e.g. quarkonium systems, three energy scales are important [18]: the quark mass m_q at which quarkonium annihilation and production takes place, the relative momentum $p_{rel} \approx m_q v \approx 1/r$, with v the relative velocity of and r the typical distance between the quarks, at which quarkonium binding takes place, and the binding energy $E_b \approx m_q v^2$. The typical hierarchy between these scales is $m_q \gg p_{rel} \gg E_b$. For heavier quarks, $m_q \gg \Lambda_{QCD}$ and $v \ll 1$, so such systems can be treated perturbatively.

2.5.3 Calculations in QCD – Effective Field Theories

The running of the coupling constant is problematic for computations, as the use of perturbation theory relies on the assumption that the expansion parameter is small, which for small q values no longer holds. The existence of multiple energy scales, as described above, further complicates things: the different scales tend to get entangled in loop-correction calculations. However, for heavier quarkonium systems, the well-separated energy scales allow to probe the different regimes at which the coupling constant operates. Heavy quarkonia with different radii experience different strengths of the Coulomb-like and confinement potentials.

One solution to circumvent the problems created by the not-so-constant coupling constant is to resort to an effective description of the processes of interest. This relies on the assumption that small fluctuations with respect to the length or energy scale of interest will average out. For example, in the neutron decay, described in section 2.8, it is sufficient to describe the process on the hadron level. It is not needed to take the dynamics of the constituent quarks into consideration. Because there are multiple energy scales in QCD, it is possible in some scenarios to integrate (average) out contributions from lower energy scales, leading to an Effective Field Theory (EFT). For the theory to be reliable, it must be matched to ‘real’ QCD at the scale at which the EFT becomes valid, which poses constraints on the form of the effective Lagrangian. The ruling theories at the different scales will be discussed below.

- If m_q is large, relativistic corrections will be small. Below this scale, energies

are no longer sufficient to create heavy quark-antiquark pairs, so contributions to the Lagrangian at the scale m_q or larger can be integrated out – to describe the motion of the Moon around the Earth, the motion of the Earth around the Sun need not be taken into account. This leads to the EFT Non-relativistic QCD (NRQCD), where fluctuations up to the cut-off Λ_{cut} are integrated out of the QCD Lagrangian, with $E_b, p_{rel}, \Lambda_{QCD} \gg \Lambda_{cut} \gg m_q$ [26].

- At smaller energies, contributions up to p_{rel} can also be integrated out, which is done in potential NRQCD (pNRQCD). When $p_{rel} \gg \Lambda_{QCD}$, the matching of pNRQCD to NRQCD can be done perturbatively. This is called the *weakly-coupled* pNRQCD regime. If the quarkonium system is small, so p_{rel} is large, the quark potentials can be calculated using perturbation theory. The energy levels of the system follow from the expectation value of the potentials, and non-perturbative terms that are not part of the potential. The latter show up in the energy level and decay calculations in the form of local or non-local electric and magnetic condensates, which interact purely by gluon exchange [18]. Within pNRQCD, these objects can be correlated to what is called the ‘gluelump’ mass [26], an on-shell gluon component, opening the door for hybrid mesons.

Using a technique called the ‘low-momentum sum rules’ in pNRQCD, precise determinations of the masses of the c and b quark have been obtained. The results [27] are $m_c = 1.279 \pm 0.013$ GeV and $m_b = 4.162 \pm 0.016$ GeV for the c and the b quark mass, respectively. These calculations use measurements on the charm production cross section by the BABAR experiment and the bottom production cross section by the CLEO experiment as input parameters. Another determination using a non-relativistic version of the aforementioned sum rules yielded $m_c = 1.25 \pm 0.04$ GeV [28] and $m_b = 4.19 \pm 0.06$ GeV [29]. The charm (bottom) quark mass calculation uses experimental data on charmonium (bottomonium) 1^{--} states. With the quark masses obtained, the energy levels of the first resonances can be calculated. However, because the quarkonium radius r cannot be determined directly, it is not clear up to what point $p_{rel} \approx 1/r \gg \Lambda_{QCD}$ holds. In general, it is assumed that the $\Upsilon(1S)$, η_b , B_c and possibly the J/ψ and the η_c mesons fall into the weakly-coupled regime. By taking the effects into account of charm quarks in the loop diagrams, the mass of the B_c was determined to the value 6307(17) MeV [30], which agrees within 2σ with the experimental measurement $6285.7 \pm 5.3(stat) \pm 1.2(syst)$ MeV obtained by the CDF collaboration [31].

Another interesting quantity that can be calculated is the hyperfine splitting, which gives information on the spin structure of the quarkonium state, and even α_s through

$$\Delta m_{hf}[\eta_b(1S)]_{LO} = \frac{C_F^4 \alpha_s^4 m_b}{3}, \quad (2.10)$$

with $\Delta m_{hf}[\eta_b(1S)]_{LO}$ the hyperfine splitting of the η_b at leading order (LO, i.e. without loop corrections), and $C_F = (N_c^2 - 1)/2N_c$. The hyperfine splitting was calculated to be $\Delta m_{hf}[\eta_b(1S)] = 41 \pm 8$ MeV [32] for bottomonium, and $\Delta m_{hf}[\eta_c(1S)] = 112$ MeV for charmonium, which seems to agree with the experimental value $\Delta m_{hf}[\eta_c(1S)] = 117.7 \pm 1.3$ MeV [19]. Another value for the charmonium hyperfine splitting was obtained by considering an interference term in the η_c line shape, and comes out at $\Delta m_{hf}[\eta_c(1S)] = 111.9 \pm 0.8$ MeV for a constructive scenario, and $\Delta m_{hf}[\eta_c(1S)] = 108.8 \pm 0.4$ MeV for a destructive scenario [33]. Also, a value for the B_c meson was obtained: $\Delta m_{hf}[B_c(1S)] = 50 \pm 17_{-12}^{+15}$ MeV [34], where the second error comes from the uncertainty in α_s . Using Eq. 2.10, the value of $\Delta m_{hf}[\eta_b(1S)]$ results in a value for α_s , evaluated at an energy corresponding to the mass of the Z boson: $\alpha_s(m_Z) = 0.124 \pm 0.001 \pm 0.001 \pm 0.001$ [35], where the first error comes from experimental input, the second from the gluon condensate, and the third from the uncertainty in the b mass. This is more than 5σ away from the PDG value, $\alpha_s(m_Z) = 0.1181 \pm 0.0011$ [2].

- When $p_{rel} \approx \Lambda_{QCD}$, the approximations above no longer hold. This is called the *strongly coupled* regime. Away from the $q\bar{q}$ production threshold, the only contribution to the Lagrangian is the quarkonium singlet field (i.e. the colour-neutral system as a whole). The Lagrangian then contains a potential that can be expanded in powers of $1/m_q$ [36]. All masses can be calculated by solving the Schrödinger equation using the resulting potential.

2.5.4 Calculations in QCD – Lattice QCD

In general, but particularly in the strong-coupling regime, analytic or perturbative calculations in QCD are challenging. In order to obtain a reliable result, computations are done numerically on a finite-sized grid with a finite spacing. This is called lattice-regularized QCD, or *Lattice QCD* (LQCD). This allows calculations from first principles. Moving from a continuous description of space-time to a discrete one sets a ‘natural’ cut on the momentum transfer, and with it p_{rel} , of the order $1/a$,

where a is the lattice spacing. a needs to be much smaller than the length scale of the physical process under investigation not to affect the results, and to ensure that they do not depend on the lattice spacing a , usually several simulations are performed with different values of a . By extrapolating the results to the limit $a \rightarrow 0$, approximations to the ‘real’ physical situation, in this context referred to as ‘continuum QCD’, can be obtained. These simulations help to determine key values of the physical properties affected by the strong interaction, such as the mass, decay width, and scattering lengths of light hadrons, and to guide experimental efforts to obtain values for the Cabibbo-Kobayashi-Maskawa (CKM [58, 59]) matrix elements, which determine the mixing of quark flavours, i.e. the probability that one flavour⁴ transits to another, but generally require one experimentally-determined quantity to gauge the results.

In LQCD, the fermions reside on the lattice sites, and the gauge fields (gluons, photons, W and Z bosons) form the connections between them. To overcome the need for vast (or even unattainable) computational resources, the so-called *quenched approximation* has been used, where only valence quarks are simulated, and sea quarks are omitted⁵. With the steady increase in computing power over the years, the available resources have presently become sufficient to support unquenched, or *dynamical*, calculations. Currently, $N_f = 2$ or even $N_f = 2 + 1$ dynamical (sea) quark flavours, N_f , can be taken into account, the first option containing the *up* and *down* quarks, and the latter also the *strange* quark. There exist different models to describe the evolution of the system; see e.g. [37] for a recent overview.

Taking realistic low values for the light quark masses is still computationally expensive, and to overcome this, these masses are enlarged. The unrealistically high masses have profound consequences to the masses of the lightest hadrons, e.g. the pions and kaons. This is especially problematic in the strongly-coupled regime, where contributions from pion exchange become important. To later ‘reconnect’ the results from such calculations to the physical world, Chiral Perturbation Theory is commonly used, where the masses of the lightest quarks are put to zero. This can be justified by considering that $m_q \ll \Lambda_{QCD}$ for these quarks. Afterwards, the result from the calculation needs to be extrapolated to correct for the non-zero-ness of the light-quark masses, moving away from what is called the ‘chiral limit’. A recent study used the

⁴The different types of quarks, shown in Figure 1.1, are called ‘flavours’.

⁵In this context, ‘valence’ quarks are those that are physical, and show up in the final state, and ‘sea’ quarks are the virtual quark-antiquark pairs that are spontaneously created and annihilated in the vacuum.

difference in the squared masses of the $u\bar{u}$ and $d\bar{d}$ mesons in this framework to obtain a value for the u and d absolute quark masses: $m_u = 2.27(6)(5)(4)$ MeV, and $m_d = 4.67(6)(5)(4)$ MeV [38], where the quoted errors are statistical, systematic, and electromagnetic (from the QED quenching error) contributions. This comes close to the PDG values $m_u = 2.2^{+0.6}_{-0.4}$ MeV, and $m_d = 4.7^{+0.5}_{-0.4}$ MeV [2], which used a weighted average of several LQCD calculations. Other results include determinations of the masses of the D meson (which has a *charm* and a *down* quark) and the D_s meson (composed of a *charm* and a *strange* quark): $m_D = 1.868(7)$ GeV and $m_{D_s} = 1.962(6)$ GeV [39], which agree with their experimental counterparts $m_D = 1.869$ GeV and $m_{D_s} = 1.968$ GeV [2]. The mass of the B_s meson has also been calculated: $m_{B_s} = 6304 \pm 12^{+18}_{-0}$ MeV, which agrees with the values quoted in Section 2.5.3. This recent progress, combined with the fact that results can be calculated from first principles, makes LQCD the most promising theoretical tool to investigate QCD, and guide experimental efforts.

2.5.5 Calculations in QCD – Quantum Numbers

The simplest (to calculate, at least) quarkonium state is the ground state, where the quarks have their spins anti-aligned and they have no orbital angular momentum between them. The total spin of the system is $S = s_{\bar{q}} + s_q = -\frac{1}{2} + \frac{1}{2} = 0$, and so the total angular momentum is $J = L + S = 0 + 0 = 0$. It is common to categorise these states by specifying the total angular momentum J , parity P , and the charge conjugation C in the following manner: J^{PC} . A parity transformation results in a spatial inversion with respect to the origin of the wavefunction of the object under consideration, $\mathbf{r} \rightarrow -\mathbf{r}$. The parity of a state is $+1$ if the wavefunction is not changed by this transformation, and -1 if it is. Fermions and antifermions have opposite parity. Hence, the intrinsic parity of any quark-antiquark system (of the same flavour) is given by $P_{intrinsic} = P_q \cdot P_{\bar{q}} = 1 \cdot -1 = -1$. The total parity P can be calculated by multiplying (as parity is a multiplicative quantum number) the intrinsic parity by the angular momentum contribution, $(-1)^L$. Therefore, if there is no angular momentum involved (i.e. $L = 0$), the intrinsic parity is equal to the total parity. Charge conjugation transforms particles into their anti-selves, and vice versa. It can be shown that C can be calculated from $C = (-1)^{L+S}$. The quarkonium $S = 0$ ground state hence has negative parity, making it a pseudo-scalar state, with $J^{PC} = 0^{-+}$.

The quarkonium states can be excited in the same way as the positronium states.

Figure 2.6. Quark diagrams with connected (left) and disconnected (right) charm quark lines of the charmonium meson (see text). $q(\bar{q})$ is an *up* or a *down* (anti)quark.

Examples include the spin excitation to the $J^{PC} = 1^{--}$ vector ground state, and radial excitations to higher n states. In Figure 2.4, the energy level diagrams of positronium and two quarkonium systems, charmonium and bottomonium, are compared.

2.6 Decay Modes in Heavier Quarkonium Systems

Heavier quarkonium systems have the possibility to decay to lighter systems, something that has no analogy in positronium, as it is already made up of the lightest leptons. As an example, charmonium, the bound system of a charm quark and its antiquark, will be considered. The relatively high mass of the charm quark ($\sim 1.28 \text{ GeV}/c^2$ [2]) allows an essentially non-relativistic treatment of the charmonium meson, making theoretical calculations from perturbation theory feasible. This feature has made charmonium quite popular, as it opens new opportunities to help understand QCD quantitatively.

For any heavy quarkonium system, the preferred decay mode is a ‘break-up’ to two lighter systems, each containing a quark of the parent system. Exemplary for charmonium would be a reaction like

$$c\bar{c} \rightarrow c\bar{u} + \bar{c}u \quad (2.11)$$

(see Figure 2.6 (left)). This type of decay is favoured, because it features connected lines in the quark diagrams, as can be seen in Figure 2.6 (left). However, for the reaction to be kinematically allowed, the mass of the $c\bar{c}$ system would need to be at least that of the daughter systems combined, which is not the case for the $n = 1$ charmonium states with $L = 0$ or 1 (see Figure 1.2). Consequently, these states must

decay via modes such as those shown in Figure 2.6 (right). These modes feature disconnected quark lines, and are known to be heavily suppressed⁶. More gluons are involved following the quark-antiquark annihilation, and their energies are much higher than that of the gluons that are responsible for the creation of the light quark-antiquark pair in reaction 2.11. As, in addition, at higher energies, the coupling constant is smaller, such contributions are suppressed. Because the decay possibilities are suppressed, states below the kinematic threshold for reaction 2.11, the open-charm threshold, are relatively long-lived ($\sim 10^{-23}$ s), resulting in very narrow resonances in the mass spectrum (see, for instance, Figure 2.8).

2.7 The Hunt for Exotic States of Matter

As explored in section 1.2, several particle states have been observed that were not predicted to exist by the models described above. An interesting feature is that all of them appear above the open-charm threshold and, indeed, they are nearly all close to the threshold for the production of meson pairs, see e.g. Figure 1.2. Close to these thresholds, the models outlined in the previous sections are not always valid, because the binding energy becomes of the order of Λ_{QCD} . As long as $E_b \ll m_q$, NRQCD retains its validity.

Several options are explored by various collaborations to identify the nature of the unexpected states. These include interpretations [18] (pages 39-51) as hybrids, tetraquark states, hadronic molecules, and hadrocharmonia (see Figure 2.3). As of yet, no consensus has been reached, as no model can unambiguously exclude the others based on the available data.

Figure 2.7 shows the calculated energy level spectrum of charmonium from LQCD, including some predictions for hybrid states using leading-order Chiral Perturbation Theory [41]. The masses at which these appear coincide with some of the X, Y, and Z states (Figure 1.2), but further measurements are needed to see if they can be assigned to them.

⁶Known as the *OZI rule*, after S. Okubo, G. Zweig and J. Iizuka [40].

Figure 2.7. Summary of LQCD calculations for the charmonium spectrum up to masses of $4.5 \text{ GeV}/c^2$, including predictions for hybrid states [41]. All states are labeled by their J^{PC} quantum numbers, and the mass of the η_c meson is subtracted from all results to reduce the systematic error. The red boxes show the lightest hybrid candidates, and the blue ones show the first excited hybrid states. The green boxes display the charmonium states, and the black line-segments represent the experimental values [2]. The height of each box indicates a 1σ deviation from the mean, which is the vertical center of the box. The lowest non-interacting $D\bar{D}$ and $D_s\bar{D}_s$ levels are indicated by the gray and green dashed lines, for the experimental and calculated values of the $D(D_s)$ meson, respectively.

2.8 Experimental Observation of Resonances

In this section, the final part of this chapter, the experimental discovery and measurement of states, like the ones in the previous sections, will be addressed.

Bottom-up Approach | Experimentally, particle states are generally observed by constructing the invariant mass, $M = \sqrt{E^2 - p^2}$, from observed final-state particles⁷, or by subsequently reconstructing the intermediate states. The former can be used when the particles decays directly to the final-state particles, such as in the decay $\text{Ps} \rightarrow \gamma\gamma$. Plotting $M = \sqrt{E_{tot}^2 - p_{tot}^2}$, where E_{tot} and p_{tot} are the total energy and momentum of the two detected photons, will give a peak at the (rest) mass of the Ps state. This enhancement of a signal in the spectrum (see, for example, Figure 2.8) is also often called a *resonance*. The intermediate-state approach is preferred

⁷These final-state particles would need to be long-lived, so they can reach the detector. Assuming the decay products to move approximately at the speed of light, their life-time τ will need to be $\tau \geq d/c$, with d the distance from the creation point to the detector.

Figure 2.8. The invariant mass spectrum of an e^+e^- combination, showcasing the original discovery of the J/ψ meson, with an invariant mass of $3.1 \text{ GeV}/(c^2)$. The shown spectrum was measured by Ting *et al.* at Brookhaven National Lab in 1974 by colliding a proton beam on a Beryllium target at high energy [21]. They called the new particle ‘ J ’. Simultaneous, the group of Richter *et al.* at Stanford observed the same particle [22] in an electron/positron collider experiment, which they called ψ . A compromise was reached, agreeing to call the particle ‘ J/ψ ’. For this discovery, Ting and Richter shared the 1976 Nobel Prize in Physics.

when the decay from the particle of interest occurs via multiple intermediate steps. An example would be $X \rightarrow \pi^0 \pi^0 \rightarrow 4\gamma$. The particle X can be directly found from the four photons in the final state, but the signal-to-noise ratio can be greatly improved if the intermediate steps are first constructed, assuming the direct decay mode $X \rightarrow 4\gamma$ does not exist. This imposes an additional constraint that cleans up the signal by forcing the decay reconstruction to proceed via the intermediate steps. In the example, once the two π^0 s have been found, the particle X can be identified from their invariant mass spectrum as before. This technique is useful in investigations if intermediate resonances exist, like the observation of the $Z_c(3900)$ in section 1.2. It is, hence, used when a *specific* decay mode of the particle of interest, like $Y(4260) \rightarrow \pi^- Z_c(3900)^+$, $Z_c(3900)^+ \rightarrow \pi^+ J/\psi$, $J/\psi \rightarrow e^+e^-$, also referred to as an *exclusive mode*, warrants studying.

It is clear that in order to construct the invariant mass, the energies and momenta of the final-state particles (in the preceding example, the photons) need to be known to a good accuracy. The more accurate those quantities can be obtained, the more accurate the invariant mass of the parent particle can be determined. This is particularly important when there are two parent particles that are very close in mass, and both have the same final-state configuration. In that case, high-precision measurements are mandatory to disentangle the states. Another example is when the

particle of interest is at the top of a lengthy decay chain, such that many intermediate particles need to be reconstructed first. If the resolution of the measured properties of the final-state particles is low, one can never expect to reach a satisfactory resolution for the properties of the particle at the beginning of the chain, as there are inherent efficiency losses at each reconstruction step. The measurement of some of these quantities is performed using an electromagnetic calorimeter, which is described in section 3.3.2.5 and chapter 4. The device measures the energies of photons, electrons, and partially that of charged pions, and also provides information on their momenta.

Top-down Approach | Another technique is to use the so-called ‘missing mass’, or ‘recoil mass’, which exploits the kinematics of a two-body decay, i.e. reactions like $A \rightarrow B + C$, for some particles A, B, C . If one of the particles on the r.h.s. of the reaction, say, B , can be directly observed in the detector, which is the case if it is a stable final-state particle like a photon or an electron, its four-momentum can be determined from that measurement. If the four-momentum of the initial state A is known in addition (e.g. from tuning the beam momenta in a collider experiment), the state of interest, C , can be searched for by plotting $M = \sqrt{(E_{init}^2 - p_{init}^2) - (E_B^2 - p_B^2)}$, where E_{init} and p_{init} , and E_B and p_B are the energy and momentum of the initial state and the B particle, respectively. This provides a more efficient way to identify the state of interest, as much less intermediate steps are involved, but only works for very specific decays. This approach is advantageous if the decay products of the particle of interest (C in this example) are not important, a so-called *inclusive mode*, i.e. $A \rightarrow B + C, C \rightarrow \text{anything}$. This way, it is possible to determine the total number of initial-state particles (here, C) that are produced. By combining this result with a study of an exclusive mode of the same particle, correcting for efficiencies, allows to extract the branching fraction (see below). Additionally, the recoil mass technique can also be used to identify particles that cannot be (easily) measured otherwise. For this to work well, the four-momenta of the parent and other daughter particles need to be accurately known. An example is the decay of the neutron: $n \rightarrow p + e^- + \bar{\nu}_e$, where the mass of the electron anti-neutrino $\bar{\nu}_e$ may be determined if the four-momenta of the neutron n , proton p , and electron e^- are known.

The previous discussion explained how states are observed, and how their masses are found. However, to identify which state has likely been observed, its angular momentum and spin quantum numbers must also be determined. This is achieved by investigating the angular distribution of the decay products. As angular momentum

is conserved, the angles under which the decay products are emitted are restricted to follow certain distributions that depend on the angular momentum of the parent.

Branching fraction | In general, unstable particles have more than one decay option. By measuring how many decays of a certain type are observed and the initial number of particles that were created, the ratio of these quantities indicates how likely this decay mode is to occur. If this *branching fraction* is non-zero, the decay is allowed, and the initial particle is said to couple to that particular final state. How strong this coupling is, provides information on the particle's structure and other properties.

Chapter 3

Experimental Facilities

To research the kind of physics described in the previous chapter, several experimental facilities exist, and several more are planned or already under construction at the time of writing. As discussed in Chapter 1, this work will mainly focus on the $\bar{P}ANDA$ experiment, that is currently being developed. Before discussing the experiment itself, the facilities housing it will be addressed.

Figure 3.1. Schematic overview of the future Facility for Antiproton and Ion Research (FAIR), with the beamlines shown in red, and the existing GSI facility, with its beamlines shown in blue. Shown in the figure are the linear accelerators UNILAC and p-Linac, the Experimental Storage Ring (ESR), Fragment Separator (FRS), High Energy Storage Ring (HESR), Collector Ring (CR), Super Ion Synchrotrons (SIS18/100), Super Fragment-Separator (SFRS), and the locations of the experiments. The numbers 1 and 2 indicate the target stations for heavy-ion production and for antiproton production, respectively.

3.1 FAIR

In Figure 3.1, the future Facility for Antiproton and Ion Research (FAIR), that is currently under construction near Darmstadt, Germany, is depicted. FAIR [42] will be an extension of the existing GSI¹ facility, using the existing accelerators as a pre-accelerator for the new double-ring (heavy) ion synchrotron SIS100/300, with magnetic rigidities of 100 and 300 Tm, for SIS100 and SIS300 (planned for the next phase), respectively, having a circumference of 1100 m. The facility comprises of a complex system of cooler and storage rings and experimental setups, one of which is *PANDA*. The synchrotrons will deliver ion beams at average intensities of $3 \cdot 10^{11}$ ions per second and momenta of 2-35 GeV/c per nucleon. The primary beam can be directed to impinge target stations (indicated by 1 and 2 in Figure 3.1) to also produce intensive secondary beams, providing antiprotons and exotic nuclei for ground-breaking experiments. These experiments are divided into four pillars:

- Atomic and Plasma Physics, which will be researched by the APPA Collabora-

¹Gesellschaft für Schwerionenforschung, a research facility for heavy ion research, nowadays called GSI Helmholtzzentrum für Schwerionenforschung.

Figure 3.2. Schematic overview of the High Energy Storage Ring [47].

tion [43],

- Compressed Baryonic Matter, covered by the CBM experiment [44],
- Nuclear Physics and Astrophysics, which will be investigated by the NUSTAR Collaboration [45], and
- Hadron and Hypernuclear Physics, which will be addressed by the $\bar{P}ANDA$ experiment [46].

For the $\bar{P}ANDA$ experiment, the secondary beam of interest is the antiproton beam, produced at target station 2 in Figure 3.1. After production, the antiprotons are sent through the Collector Ring (CR), where they are stochastically cooled to a momentum resolution of $\sigma_p/p \approx 3 \cdot 10^{-2}$ [42], with σ_p the standard deviation in the momentum p , and then accumulated in the Accumulator Ring (CESR, planned for the next phase). After a sufficient antiproton density is reached, they are guided to the High Energy Storage Ring (HESR) for use in experiments.

3.2 HESR

Designed specifically for $\bar{P}ANDA$, the racetrack-shaped HESR stores the antiproton beam and cools² it further to improve the momentum resolution. To achieve this, an electron cooler will be installed in the lower of the two straight sections of the

²If the average momentum of an antiproton bunch would be subtracted from each particle in the bunch, then the antiprotons would essentially move around randomly. This can be interpreted as the temperature of the bunch; hence, reducing the internal momentum spread can be viewed as ‘cooling’ the bunch.

ring (Figure 3.2), together with a stochastic cooling system comprised of kickers just before the electron cooler and high-sensitivity pick-ups at the end of the opposing straight section. Electron cooling, developed in 1966 at INP, Novosibirsk, Russia, employs a beam of electrons with the same average velocity as the antiprotons, that is injected parallel to the beam. The momentum spread of the electrons is much lower, because they can be accelerated using an electrostatic potential. The antiprotons undergo Coulomb scattering on the electrons, and thereby transfer momentum. At ‘thermal’ equilibrium, all particles have the same momentum, but as the electrons have a much lower mass, their velocity will be much higher. The electrons are extracted, and the resulting momentum spread in the antiproton beam is reduced. Stochastic cooling is another technique to reduce the momentum spread in a particle beam, developed at CERN, Geneva, Switzerland, in 1968. In this technique, momentum information on small groups of beam particles, or even individual particles, depending on the sensitivity, is read out non-destructively and electromagnetically by the *pick-ups*. The low intensity of the signals induced by these groups of particles make this information difficult, but not impossible, to obtain. The signal is amplified and sent to the *kickers*, which use an electromagnet to exert a correcting force on the particle group. Because of the smallness of this effect, and because most particles cannot be addressed individually, this operation needs to be repeated several times to obtain a reasonable amount of cooling. The HESR features two modes of operation: High Luminosity (HL), where more antiprotons are stored in the ring to achieve a higher luminosity, at the cost of momentum resolution, and High Resolution (HR) with complementary features to the HL mode. Stochastic cooling is foreseen to reduce the momentum spread to $1.5 \cdot 10^{-4}$, which is sufficient for the HL mode. To improve the resolution to $4 \cdot 10^{-5}$ (Root Mean Square (RMS), a common statistical measure for the width of a presumed Gaussian distribution) for the HR mode, the electron cooler will be used, with electron energies up to 5 MeV. The antiproton beam enters the HESR from the top right in Figure 3.2, with momenta between 1.5 and 15 GeV/c. The expected mean lifetime of the beam is 0.43 h when its momentum is 1.5 GeV/c, and 1.97 h for 15 GeV/c [47]. The operational specifications are listed in Table 3.1 for the two respective modes. To be able to probe the type of physics described in Chapter 2, the $\bar{P}ANDA$ experiment requires that the HESR delivers at least 10^{10} antiprotons per run³, with the beam momentum in the range between 1.5 and 15 GeV/c and an RMS momentum spread $\sigma_p/p \leq 10^{-4}$. From Table 3.1, it is clear that these requirements are met by the HESR.

³A *run* is the operational period until a new injection of antiprotons is needed.

HESR Mode	p (GeV/c)	RMS momentum spread (σ_p/p)	Peak luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	Number of antiprotons
HL	1.5 – 15	$1 \cdot 10^{-4}$	$2 \cdot 10^{32}$	10^{11}
HR	1.5 – 8.9	$4 \cdot 10^{-5}$	$2 \cdot 10^{31}$	10^{10}

Table 3.1. HESR specifications for the two modes of operation (High Luminosity and High Resolution), when the target density is $4 \cdot 10^{15}$ atoms/cm².

3.3 $\bar{P}ANDA$

Using the antiproton beam from the HESR to illuminate an internal proton target, combined with a magnetic spectrometer, the $\bar{P}ANDA$ experiment will be able to probe a rich hadron physics programme. $\bar{P}ANDA$ will be a fixed-target type experiment, where, in contrast to colliding-beam experiments, particle production will be boosted in the forward direction. This is reflected in the asymmetric design of the detector (Figure 3.3).

3.3.1 Physics Programme

The goals of the experiment include the physical processes outlined in Chapter 2, and will therefore improve our current understanding of the strong interaction and the structure of hadrons. In particular, the programme entails the investigation of QCD bound states, as precision measurements are needed to help guide the different theoretical models (section 2.5.3 and 2.5.4), and to look for exotic states of matter (section 2.7). In the hyperon physics programme, hyperon pair production proceeds via the creation of a strange-antistrange quark pair, giving rise to non-perturbative QCD dynamics, which can be studied from the proton-antiproton annihilations. By replacing the proton target by heavier nuclei, is it possible to extend the programme to include hypernuclear physics and the study of hadrons in a nuclear medium. The latter is expected to help to understand the origin of hadron masses, in the context of spontaneous chiral symmetry breaking (section 2.5.4). Hypernuclear physics involves the formation of so-called hypernuclei by introducing hyperons in nuclei, either by binding hyperons to existing nuclei, or by replacing an up or down quark by a strange quark. This allows to probe nuclear structure and to investigate the forces between hyperons and ‘regular’ nucleons, which will shed light on the strong interaction at larger distance scales.

Next to these spectroscopic goals, it will also be possible to measure electromagnetic form factors of the proton via Deeply Virtual Compton Scattering and the process $p\bar{p} \rightarrow e^+e^-$. The large amounts of D mesons that will be produced will allow additional studies on rare weak decays, and $D_{(s)}$ -meson spectroscopy [48].

3.3.2 Detector Setup

From the discussion in section 2.8, it is clear that in order to detect particles of interest, accurate measurements of the energy, momenta and angular distributions are mandatory. Because no single detector is able to simultaneously measure these properties with sufficient accuracy, the detector setup of $\bar{P}ANDA$ combines several systems with different functionality, like most particle physics experiments. These systems are nested into each other in an onion-like construction to maximally encase the interaction region. This section describes the different detector subsystems that are currently being developed, starting from the interaction point outwards. Each subsection strives to provide a brief description of the subsystem, along with some general properties and features. More details can be found in the reference of the corresponding section. An overview of the detector is shown in Figure 3.3.

3.3.2.1 Target System

The design of the solenoid magnet allows different types of target systems [49] to be implemented. These include both gaseous and non-gaseous targets. Currently, two possible implementations are being developed, both of which are capable of providing a target density of $4 \cdot 10^{15}$ hydrogen atoms/cm², as is required to fulfil the design goals of the high luminosity mode.

Cluster Jet | A cluster jet target gives a homogeneous and adjustable target density. To produce the jet, a cooled, pressurized gas is expanded into vacuum through a nozzle, causing the gas to condensate into clusters, which can then serve as targets. As the jet comes in at a direction transverse to the antiproton beam axis, the interaction point is well defined in the plane perpendicular to the beam axis (by the overlap with the beam), but needs to be reconstructed at a later stage along the beam direction. Density fluctuations are $\leq 3.5\%$ (RMS), and the position resolution is $\leq 50 \mu\text{m}$.

Hydrogen Pellet | Pellet targets consists of droplets of frozen molecules dripping from a nozzle into vacuum at a fixed frequency, giving rise to a high target density. The dripping leads to a non-uniform time distribution. In particular, several pellets

Figure 3.3. Overview of the $\bar{P}ANDA$ detector setup, which has a height of about 5 m and a length of about 14 m. The different subsystems are discussed in the text. The antiproton beam comes in from the left side of the figure. Tracking systems are shown in blue.

may end up in the beam together, meaning that large variations in the instantaneous luminosity will occur. The lateral spread can be reduced using skimmers (a type of cylindrical collimator with a fixed diameter), but as the maximal achievable luminosity is very sensitive to deviations from the target axis, the beam would still need to be widened. It is foreseen to implement a tracking system for the pellets using lasers and cameras to improve the position resolution of the primary interaction. The design dimensions for the pellets feature a radius of 25-40 μm , and a lateral RMS deviation in the stream of 1 mm. A resolution of 50 μm should then be achievable for the interaction point.

3.3.2.2 Micro Vertex Detector

The Micro Vertex Detector (MVD) [50] is a part of the detector used to help locate the interaction point, and the vertices (point of decay) of decaying particles. These include both primary and secondary vertices, such as for e.g. D mesons, which due to their relatively long lifetime can travel some distance before they decay. The MVD consists of a barrel and two caps for maximum acceptance. The barrel part of the MVD consists of 4 layers of radiation-hard silicon pixel detectors, while the caps combine this with double-sided silicon micro strip detectors. The MVD is designed to obtain a spatial resolution of 100 μm in the direction along the beam axis, and a few tens of μm in the transverse direction, and has an acceptance of $3^\circ - 150^\circ$ in polar angle⁴. This device also provides information to improve the momentum reconstruction.

3.3.2.3 Straw Tube Tracker

The Straw Tube Tracker (STT) [51] is a detector subsystem used for reconstructing tracks created by charged particles. The STT consists of 4,636 tubes, made of a thin foil and filled with an inert gas mixture ($\text{Ar} + 10\% \text{CO}_2$), and a central anode wire carrying a high voltage (several kV). The inner surface of the tubes is conductive and can thus act as a cathode. When a charged particle crosses the tube, it ionises the gas. The electrons that are freed drift to the central wire, producing a current, which signals the detector that a particle has passed. When a charged particle traverses the detector, it activates all the tubes in its path, allowing for a reconstruction of its track. The energy loss in the straw volume can be used to identify the particle species (particle identification, PID). The STT has an acceptance of $22^\circ - 140^\circ$ in polar angle, and a

⁴The polar angle is the angle with respect to the beam axis.

momentum resolution of $(\sigma_p/p) = 1 - 2\%$. To be able to accept tracks with polar angles smaller than 22° , three planar stations equipped with gaseous micro-pattern detectors based on Gas Electron Multiplier (GEM) foils will be placed in that region.

3.3.2.4 Cherenkov Light Detector

Cherenkov light is emitted when a particle would exceed the local speed of light. The light is emitted in a cone along the trajectory of the particle at an angle, which depends on the particle's velocity. When combined with momentum information from the tracker systems, this angle of emitted radiation can be used to determine the mass of these particles. This procedure is helpful to differentiate between charged pions and kaons. The detector, called 'Detection of Internally Reflected Cherenkov light system', or 'DIRC system' for short, consists of 1.7 cm thick fused-silica plates, placed in such a way as to enclose the target volume in a cylindrical shape. As can be seen in Figure 3.3, it does not form a closed box, but the larger size of the disk part of the DIRC system allows a continuous coverage by the complete system of $5^\circ - 140^\circ$, again in polar angle [50].

3.3.2.5 Electromagnetic Calorimeter

The Electromagnetic Calorimeter (EMC) [52] is a device which is designed to absorb all energy of the (electromagnetic) particles that enter the detector. In practice, the particles which are fully stopped by the EMC are mainly photons and electrons (and positrons), originating from e.g. the decay of some other particle, or bremsstrahlung processes. The particles are absorbed by lead tungstate ($PbWO_4$, or PWO for short) crystals. For a discussion on the interaction of particles with the crystals, see section 4.1.1. The EMC contains approximately 15,500 PWO crystals in total, which are cooled down to $-25^\circ C$ to improve the light yield by an additional factor 2, when compared to room temperature (section 4.1.2).

The target spectrometer EMC consists of three parts, the barrel (containing 11,360 crystals), and the forward and backward endcap (containing 3,600 and 592 crystals, respectively), and is complemented by a sampling ('shashlik'-type) calorimeter (see section 4.1.1) made of lead-scintillator sandwiches in the forward spectrometer to cover very small polar angles down to almost 0° . As the antiproton beam will hit the target at an angle (about 2° from the central axis), even particles emitted at a polar angle of 0° in the centre-of-mass frame can be detected. Thanks to this design, the calorimeter will have an acceptance of 99% of the full solid angle. The design energy

Figure 3.4. Measured time resolution as a function of the energy deposited in a crystal, performed at the electron accelerator MAMI at Mainz (Germany) with quasi-monochromatic photons from bremsstrahlung with energies up to 1.5 GeV, produced by a monoenergetic electron beam [53].

resolution of the target spectrometer EMC is $(\sigma_E/E) = 1\% \oplus 2\%/\sqrt{E/\text{GeV}}$ ⁵ over a large dynamic range of 20 MeV to 14.6 GeV, with σ_E the standard deviation in the energy E . The spatial resolution, given in terms of the polar angle, is expected to be $\leq 0.5^\circ$, $\leq 0.3^\circ$, and $\leq 0.1^\circ$, for the backward endcap, barrel, and forward endcap, respectively. The time resolution, which is a critical property at high interaction rates, is very good, less than 1 ns for energy depositions above 200 MeV, as can be seen in Figure 3.4.

Because of the forward-boosted particle production, the forward endcap is expected to endure the highest hit rate, which is the reason for the asymmetry between the forward and backward endcap. In the centre, individual crystal hit rates up to 500 kHz can be expected, dropping down to 100 kHz around the edge of the endcap, see Figure 3.5. Due to fluctuations in the target thickness of the pellets, spikes as high as 1 MHz can be expected to occur in the most forward region. This high hit rate, combined with the enormous variety of particles created, poses significant experimental and technical challenges. These include the occurrence of overlap between detector signals and mixing of events, and the need to process a very high data rate. About 275 Gbps of raw data is produced by the calorimeter alone; see section 4.3.3 for a breakdown of the calculation. In sections 4.2 and 4.3, it is discussed how the readout system deals with these challenges.

⁵The *direct sum*, $\oplus$, of two quantities a and b is $a \oplus b = \sqrt{a^2 + b^2}$.

Figure 3.5. Average hit rate, expected at the crystals in the forward endcap of the $\bar{P}ANDA$ EMC, at an antiproton beam momentum of 15 GeV/c with an interaction rate of 20 MHz, obtained using the DPM event generator (section 5.2.2).

3.3.2.6 Time-of-Flight System

Time-of-Flight (TOF) spectroscopy is based on the principle that heavier particles have larger moments of inertia. Combining the time it takes a particle to travel between two points in the detector, called the flight time, with the distance travelled allows to determine the particle's velocity. Through its momentum (determined using information from the tracker subsystems) the mass of the particles is obtained. The TOF subdetector is made of 5,760 scintillator tiles oriented in a cylindrical shape, surrounding the DIRC detector, with a time resolution of 100 ps. The tiles will be read out by silicon photomultipliers [50].

3.3.2.7 Solenoid Magnet

The large solenoid magnet provides a homogeneous 2 T magnetic field (with field strength variations less than 2%) that causes charged particles to follow curved trajectories by an induced Lorentz force. This is of vital importance for the functionality of most detector subsystems, most notably the tracker systems.

3.3.2.8 Muon Counters

Several leptonic decay modes involve the production of muons, which are generally not fully stopped by any of the subsystems. To make sure they don't elude detection, which would hamper particle reconstruction, muon counters are placed in the outermost shell of the detector. The muon counters do not measure the momenta or energy of the muons, but only register whether or not muons have passed through.

Muon identification in the reconstruction stage is performed using information from other detector subsystems; the muon counters provide information if one or more reconstructed tracks belong to muons. Placed on the yoke of the solenoid magnet, the device consists of 13 sensitive layers, alternating with absorber layers made of iron. The amount of material has been tuned to ensure all charged pions are absorbed in the momentum range covered by $\bar{P}ANDA$. A setup of Mini Drift Tubes will be used as sensors between the sensitive layers [50].

3.3.2.9 Forward Spectrometer

As particle production is strongly forward boosted, a separate part of the detector, the Forward Spectrometer, is placed next to the Target Spectrometer downstream of the beam to measure particles at small polar angles [48]. The Forward Spectrometer will consist of a dipole magnet with a bending power of 2 Tm, which will deflect charged particles whose momenta can then be measured using Mini Drift Chambers with a resolution of $(\sigma_p/p) = 0.2\%$. For particle identification, in particular to separate charged pions, kaons, and protons, two detector types will be used. The first is a Cherenkov-light imaging detector, similar to the one used in the HERMES experiment [54]. Using two radiators with different indices of refraction, to wit silica aerogel and C_4F_{10} gas, a $\pi/K/p$ separation in a broad momentum range of 2-15 GeV/c will be possible. To improve the separating power, a TOF wall is installed with a time resolution of 50 ps, that will allow a 3σ separation for π/K and K/p at momenta up to 2.8 GeV/c and 4.7 GeV/c, respectively. Behind this, the ‘shashlik’-type EMC is placed, as described in section 3.3.2.5, which has an energy resolution of $4\%/\sqrt{E/\text{GeV}}$. Muon counters in the Forward Spectrometer complement the ones in the Target Spectrometer, described in section 3.3.2.8, and allow discrimination of pions from muons, detection of charged-pion decays and, a low-resolution energy measurement of neutrons and antineutrons. At the beam exit, a luminosity detector is placed, that determines the number of impinging antiprotons with a resolution of about 3%. A precise measurement of the instantaneous luminosity is needed to extract e.g. absolute cross sections.

Chapter 4

The EMC: From Particle to Byte

The present work is primarily concerned with the Electromagnetic Calorimeter (EMC) of the $\bar{P}ANDA$ experiment. This device will be exploited as an example for the way the other subsystems process the produced raw data. The EMC was discussed in section 3.3.2.5, where some of the technical and experimental challenges that accompany it (and most other subdetectors) were unveiled. In this Chapter, a closer look will be taken at the materials and the readout of the EMC, and the solution that is foreseen to solve the challenges is presented. The path of the data from the first interaction of the physical particles to the offline storage will be discussed. The software packages used to simulate the detector properties and its response will be presented in chapter 5.

4.1 Scintillators

In section 3.3.2.5, the Electromagnetic Calorimeter of the $\bar{P}ANDA$ experiment, the detector subsystem of primary interest in this work, was introduced. The active material in the target spectrometer EMC is a scintillator, made of lead tungstate ($PbWO_4$, PWO). Some general properties of scintillators will be introduced in this section, including the working principle, and specific details of lead tungstate and why it was chosen.

4.1.1 General Properties

In order to detect particles, they must undergo interactions with the detector material. The type of interaction differs for charged and for neutral particles, but as the EMC measures particle energies, the interactions imply energy losses in the material. As the focus in this chapter lies on the EMC, also the descriptions of the interactions will focus on those that are most relevant for the EMC.

Charged Particles | When charged particles traverse matter, they can be deflected by the electric field of the nuclei in the material. This will cause them to radiate so-called *bremsstrahlung* photons (German for ‘breaking radiation’), depicted in Figure 4.1 (left). The process is proportional to E/m_p^2 , for an arbitrary charged particle of mass m_p with energy E . Hence, this process gives an important contribution to the energy loss of electrons and positrons, in particular. For relativistic electrons, the energy loss is given by

$$-dE/dx = E/L_R, \quad (4.1)$$

Figure 4.1. Feynman diagrams for bremsstrahlung emission (left) and pair production (right).

with dx the distance travelled in the material, and L_R the *radiation length*, which is inversely proportional to Z^2 (the atomic number squared), and the number density of atoms/cm³, n_a , in the medium. Integrating Eq. 4.1 reveals that L_R is the average distance a charged particle needs to travel through a medium until its mean energy is reduced by a factor e . As it is proportional to Z^2 , the radiation length is material-dependent, and is much shorter for high- Z material, which is why those materials are favoured in calorimeters. At energies below a few tens of MeV, Coulomb scattering from atomic electrons becomes the dominant process for energy loss.

Neutral Particles | Neutral composite particles will generally decay before they can reach the calorimeter¹, and of the neutral elementary particles, neutrinos are mostly not detected because of their small interaction cross section. This leaves photons, which, unlike charged particles, have a high probability to be absorbed or scattered through large angles by the atoms in the detector material. Analogous to the energy loss of charged particles, the number of photons (the intensity I) lost when travelling through a piece of material of thickness dx is given by

$$-dI/dx = I/\lambda, \quad (4.2)$$

with $\lambda = (n_a \sigma_\gamma)^{-1}$ the mean free path of a photon, and σ_γ the total interaction cross section. The mean free path, hence, plays the role of the radiation length in the case of photons. The main contributions to the cross section are

- *Rayleigh Scattering*, where a photon scatters coherently off an atom,
- the *photo-electric effect*, where the photon is absorbed by the atom, thereby freeing one of its electrons,
- *Compton Scattering*, in which the photon loses part of its energy by scattering off an atomic electron, and
- *pair production*, when the photon, provided its energy is sufficient, is converted to an electron-positron pair in the vicinity of the electric field of a nucleus or an atomic electron (see Figure 4.1 (right)).

Figure 4.2 shows a plot of the cross section for the different processes. As can be seen, above its threshold, pair production is the dominant contributor. To a good approximation, it is given by $\sigma_{pair} = 7/9 \cdot 1/(n_a L_R)$. Neglecting the other

¹A noteworthy exception is the neutron, which is unlikely to decay inside of the $\bar{P}ANDA$ detector setup. It may interact with setup material, but there is no way to unambiguously identify it.

Figure 4.2. Plot of the energy-dependent cross section for photon interaction processes on Carbon (open circles) [60]. σ_{Rayleigh} , σ_{Compton} , and $\sigma_{\text{p.e.}}$ are the contributions from Rayleigh scattering, Compton scattering, and the photo-electric effect, respectively, and κ_{nuc} and κ_e represent contributions from pair production in the electric field of nuclei and electrons, respectively (see text).

contributions to the cross section, which is a valid assumption at sufficiently high energies (as can be seen in Figure 4.2), Eq. 4.2 reveals that photon absorption, like electron radiation loss, is characterised by the radiation length L_R .

Energetic photons, electrons, and positrons, produced in the reactions, deposit their energy in the calorimeter by means of these processes. When such a particle enters, it will create a cascade of secondary particles, where electrons emit bremsstrahlung photons, and photons undergo pair production. This continues until the energy of the bremsstrahlung photon is no longer sufficient to initiate pair production, at which point the cross section for electron energy loss via ionisation equals that of bremsstrahlung. At this critical energy E_C the cascade development stops abruptly. The cascade of particles is called an *electromagnetic particle shower*. The main features of this shower can be described by a simple model, in which each electron with $E > E_C$ emits a bremsstrahlung photon with half of the electron initial energy after travelling one radiation length. Each photon with $E > E_C$ decays into an e^+e^- pair after one radiation length, with each particle having half of the photon energy. When $E < E_C$, these processes stop abruptly, and the electrons lose the remainder of their energy through ionisation losses, while these losses are negligible above E_C . The photons continue to interact according to the other processes depicted in Figure 4.2. It follows that after each radiation length, the number of particles in the shower doubles. If the initial particle's energy $E_0 \gg E_C$, there will be 2^n particles, each with average energy $E_n = E_0/2^n$, after n radiation lengths. The

shower development stops at $E = E_C$, after having continued for

$$n_{max} = \ln(E_0/E_C)/\ln 2 \quad (4.3)$$

radiation lengths. This shows that the maximum shower depth in a material grows only logarithmically with the primary energy, meaning that the physical size of a calorimeter need only increase slowly with the maximum energies of the particles to be detected. Though the shower will mainly develop in the longitudinal direction with respect to the momentum vector of the primary particle, scattering, pair production, and bremsstrahlung of the secondary particles forces the shower to expand in the transverse direction. How large this spread is, depends on the material. The resulting size of the image of the shower's end can be described by a radius, the so-called *Molière radius*, R_M . As the Molière radius and the radiation length both characterise how far a particle penetrates the material, albeit in different directions, they are related. The relation approximately follows $R_M = 0.0265 L_R (Z + 1.2)$. A small Molière radius and radiation length allow for a compact calorimeter design, and a high granularity, which increases the angular resolution of the device, and hence the momentum resolution.

The particles in the shower excite or ionise a large number of molecules in the calorimeter material. These molecules may de-excite to the ground state by the emission of photons in the visible light range. This process is called ‘radioluminescence’, or *scintillation*. To measure the energy of the incoming particle, the scintillation photons are collected and the amount of light (which is proportional to the energy of the incoming particle) is measured using photosensors. Hence, it is important that the detector material itself is transparent to photons with those energies (wavelengths). Most transparent materials will produce scintillation light to some extent, but the amount is generally too small to be measured reliably. In some materials, the conversion of excitation energy into scintillation light is more effective. These are called *scintillators*. In general, there are two types of scintillator: organic, in which the scintillation light comes from transitions in the energy levels of the molecules in the material, and inorganic, where light is produced by excitations in the crystal lattice sites. The latter material often has higher Z .

To reduce the physical size of the calorimeter, sometimes additional absorber material is sandwiched between layers of scintillator material. The absorber material has a shorter radiation length than the scintillator, allowing to reduce the overall size of

Figure 4.3. A lead tungstate crystal of the CMS ECAL with a Vacuum Photo Tetrode connected to it. Picture by W. Funk, CERN - Team from United Kingdom (D. Cockeill, B. Smith).

the device. Generally, the absorber material is also made of a cheaper material, allowing to reduce the cost of the device as well. This design is used when the physical size of the calorimeter becomes very large, where it is impossible or extremely costly to produce crystals of the sizes needed. This type of calorimeter is called a ‘sampling calorimeter’, because it samples the particle shower at the scintillator layers.

4.1.2 Lead Tungstate

A number of factors played a role in choosing the scintillator material [52] for the *PANDA* calorimeter. As the target spectrometer EMC is located inside the superconducting solenoid, the material is required to have a short radiation length and Molière radius. To handle the high hit-rates (see section 3.3.2.5), the decay time (i.e. duration of the produced light pulses) must be short — a few tens of ns or better. In addition, the scintillator should enable a sufficient energy resolution and detection efficiency with particle energies over a large dynamic range, and it needs to exhibit a decent radiation hardness. These requirements already rule out many candidates. Several remaining candidates fulfill these requirements (see also: Table 4.1), including:

- PWO, currently used by the CMS Collaboration at CERN for their EMC (called ECAL [61]). A picture of a single crystal is shown in Figure 4.3,
- CeF_3 , a fast scintillator with an extremely radiation-hard crystal matrix, and a short radiation length and Molière radius,
- Bismuth Germanate ($\text{Bi}_4\text{Ge}_3\text{O}_{12}$, BGO), a well-performing scintillator used in several experiments, like L3 at CERN [62], or GRAAL at Grenoble [63], or
- $\text{Lu}_{2(1-X)}\text{Y}_{2X}\text{SiO}_5$ (LYSO), which features a particularly high stopping-power and fast and bright luminescence.

Out of these, PWO was chosen, because of its short decay time. BGO has a much longer decay time, CeF_3 has never been produced with adequate quality for samples

Material	ρ (g/cm ³)	L_R (cm)	R_M (cm)	dE/dx (MeV/cm)	τ_{decay} (ns)	Relative LY	d(LY)/dT (%/°C)
NaI(Tl)	3.67	2.59	4.13	4.8	245	1	-0.2
BGO	7.13	1.12	2.23	9.0	300	0.21	-0.9
LYSO(Ce)	7.40	1.14	2.07	9.6	40	0.85	-0.2
PWO	8.30	0.89	2.00	10.1	6.5	0.003	-2.5
PWO-II	8.30	0.89	2.00	10.1	6.5	0.006	-3.0
CeF ₃	6.16	1.70	2.41	8.42	30	0.073	0

Table 4.1. Key properties of several scintillator materials [2, 52]. ρ is the density of the material, R_M is its Molière radius, dE/dx is the average energy loss of incoming particles per centimeter of scintillator material, τ_{decay} is the decay time, and LY is the light yield (see text). In some materials, the decay time has a slow and a fast component, depending on the mechanism of scintillation. Here, only the time of the slow component is quoted for those materials, as it determines the response time of the material, and is generally also the chief contributor to the light yield.

larger than 10 radiation lengths, and LYSO is a very expensive material. However, the light yield² of standard PWO is much lower, which is why a study was initiated to improve it [64]. A limiting factor was the introduction of impurities in the crystal lattice during production. New techniques allowed to increase the structural perfection of the crystals, and de facto the light yield, without compromising other favourable properties (for example, the use of doping can also increase the light yield, but often introduces long decay-time components into the material [65]). This improved the light yield of the PWO crystals, now referred to as PWO-II, by a factor 2. Cooling the crystals down to -25°C reduces the thermal quenching experienced at room temperature, increasing the yield by an additional factor 2 to about 100 photo-electrons/MeV for 1.2 GeV photons from a ^{60}Co source [52], but was shown to compromise the recovery capabilities of the crystals following radiation damage. New studies implied that this damage can be recovered by illuminating the crystals with UV light from LED pulsers in between time periods when the beam is on target [66].

²The *light yield* is the number of photons produced by an incoming particle per MeV of its energy. The higher the light yield, the lower the minimal photon energy that can be detected. Also, since the ability to detect a single photon follows a probability distribution, a larger number of photons increases the detection probability, and thereby the resolution.

4.2 Readout Electronics

The light produced in the PWO crystals is collected at the backface of the crystals, guided there by wrapping the crystals in reflective foil, and is converted into an analogue signal by photosensors. In the EMC of the $\bar{P}ANDA$ detector, two types will be used:

- Hamamatsu Silicon Large-Area Avalanche Photo Diodes (LAAPDs). These are reverse-biased semiconductor diodes. When a photon of scintillation light hits the diode, it excites an electron, which accelerates under the influence of the applied bias voltage. It then frees more electrons in the next semiconductor layer. This process continues, creating an avalanche of electrons, whose current can be measured. The internal gain is at least about 50; the $\bar{P}ANDA$ detector will apply a gain between 100–400. APDs feature a high quantum efficiency³ in the order of 70–80%, and cover a broad spectral range from the near infrared to ultraviolet and even to high-energy regions [67]. The APDs for the $\bar{P}ANDA$ experiment feature a large active area of $7 \times 14 \text{ mm}^2$, two of which will be fitted to each crystal to cover most of the backface. This increases the light collection efficiency, offers redundancy (see Section 4.3.1) and allows the detection of a nuclear counter event, which occurs when a shower particle deposits its energy directly in the photosensor (a so-called *punch-through*). With two APDs per crystal, likely only one will show a signal from this effect, allowing to correct for it.
- Hamamatsu Vacuum Photo TeTodes (VPTTs). The front electrode in Photo Tetrodes is a photocathode, that releases electrons when a photon hits. The electrons are accelerated to the anode, which consists of a fine mesh. Most electrons miss the anode and impinge on the dynode behind it, which releases secondary electrons that are accelerated back towards the anode and are collected there. A fourth electrode provides an electrostatic screen to reduce the internal capacitance. The quantum efficiency for VPTTs is around 20%. They are radiation hard and can handle high hit-rates, and will therefore be used in the most forward region of the EMC (i.e. at the centre of the Forward Endcap). A single VPTT is fitted to each crystal.

³The quantum efficiency is the relative number of electrons per photon that are released after it interacts with the photosensor.

Figure 4.4. Schematic overview of the readout chain for the EMC. A step-by-step description of the components is provided in sections 4.2 through 4.3.4.

The amplitude of the signal coming from the photosensors is too low for further processing, and requires amplification. Amplification is done by integrating the input pulse. The resulting waveform features a sharp leading edge of 50 ns and a long trailing edge, making it unsuitable for digitisation⁴. To remedy this deficiency, the signal is shaped to a bell form. Both actions are performed by a single, custom-made Application-Specific Integrated Circuit (ASIC). In the Forward Endcap of the EMC, the APDs will be fitted with Low Noise and Power (LNP) discrete component preamplifiers designed at the University of Basel [68]. The LNPs feature a noise of 160 keV and a (idle) power consumption of 45 mW. In the Barrel part of the EMC, the APDs will be read out by the APFEL ASIC [69] designed at the GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany, with a slightly higher power consumption of 56.4 mW. Both preamplifiers operate with two gains to allow them to cover a large dynamic range. This way, energy depositions from 2 MeV up to 12 GeV can be measured.

4.3 Readout of the Detector

As discussed in chapter 1, the physics programme of the $\bar{P}ANDA$ experiment rules out the possibility of using a triggered readout. However, the raw detector output would produce too much data to be stored (estimated to be $\sim 2,400$ Gbps, 275 Gbps of which will be produced by the EMC), and it would mainly contain background events and noise hits. To only store data possibly containing events of interest, a new readout concept, called *triggerless readout*, is being developed. The detector subsys-

⁴Digitisation starts from sampling the analogue signal; to do this properly, the sampling frequency is defined by the rise time of the leading edge of the pulse shape. Hence, the sharper the leading edge, the higher the required sampling frequency.

Figure 4.5. Photographic picture of the most recent SADC prototype, with a 1 Euro coin for scale. The most relevant components are highlighted. Note: only one optical output link will be used in the final design.

tems preprocess the data on-the-fly, and the final event selection is determined by a physics-based software ‘trigger’. To reduce and preprocess the data, the subsystems form an intelligent, self-triggering, dead-time-free readout system. To illustrate the concept, the readout chain will be discussed for the EMC, although most steps also apply to other detector subsystems. An overview is presented in Figure 4.4.

4.3.1 Digitisation

The output from the preamplifiers is digitised by *Sampling Analogue-to-Digital Converters* (SADCs, see Figure 4.5), which sample continuously at 80 MHz and feature a 14-bit resolution. The SADCs have 64 input channels, and read out both preamplifier gain outputs. If a clipping at the high-gain output signal is detected, the low-gain output will be used. As most crystals are equipped with two APDs, each digitiser can read out 16 crystals. In total, approximately 520 SADCs will be needed for the EMC (see Table 4.2 for a breakdown of the calculation). The two APDs will be read out by two different digitisers, so if one of the SADCs needs to reboot, the detector will not be blind in the corresponding section. Rebooting takes only ~ 300 milliseconds, but some data will be lost if this redundancy procedure is not used. The SADCs will be placed inside the solenoid magnet, because the feed-through to the outside of the detector is too small for the many cables of the APDs. This makes it very difficult to replace or bypass the SADCs, which is why this redundancy procedure was chosen. The power supply of the device has been custom designed to be able to operate inside the magnetic field. As they are also close to the cooled crystals, the digitisers must not only be compact, but also produce only

Figure 4.6. (left) Representation of the distorted waveform, resulting from two (partially) overlapping detector response pulses. Symbols: $A_{1(2)}$ is the amplitude of the first (second) pulse, I is the integral of the combined waveform. (right) Calculated and measured pile-up probabilities for different hit rates, for a pulse with a length of 800 ns. $N_{s(p)}$ is the number of single (pile-up) pulses [70].

little heat. The SADC is in its final stage of development at the time of writing. The hardware development is taking place at the University of Uppsala, Sweden.

Because of their location, the SADCs will be exposed to high levels of radiation. Several tests were performed on the prototypes to identify vulnerabilities. One solution is to implement so-called *triple modular redundancy*, where the data is processed three times, and the three output streams are compared. If at least two of them are equal, that result is pushed to the next stage. This helps prevent radiation-induced data corruption. The computing cores of the boards also have a built-in capability to detect and correct configuration errors. However, this will only work if a single bit in the configuration was changed (a so-called *Single Event Upset*). If multiple bits are changed, the device needs to be rebooted to recover the errors.

Preprocessing is done in two stages, by employing two Xilinx Kintex-7 Field-Programmable Gate Arrays⁵ (FPGAs). The preprocessing capabilities of the device in the first stage include a baseline follower, correcting for slowly varying offsets in the signal baseline, and pulse detection. By considering the ratio between the amplitude of the pulse and its area, the SADC is able to determine if a pile-up⁶ has occurred. The pile-up detection methods have been investigated by G. Tambave *et al.* [70], and the rate-dependent probability for pile-up is shown in the right part of Figure 4.6. A rough timestamp is assigned using the timestamp of the maximum amplitude. If a

⁵A Field-Programmable Gate Array (FPGA) is an integrated circuit containing an array of programmable logic blocks, connected by a collection of reconfigurable interconnects.

⁶If two or more particles hit the same crystal shortly after each other, the response signals can start piling up, leading to a distorted waveform, see the left part of Figure 4.6.

pile-up was detected, the complete waveform is sent to the next part of the readout chain. Single pulses are subjected to the second stage, which starts by calculating the precise timestamp using Constant Fraction Timing (CFT) [53]. In CFT, the time is calculated by adding an inverted fraction of the original signal, delayed by the original rise time, to the signal. The precise time is obtained from the first zero crossing point of the resulting waveform, which is extracted by linear interpolation. This way, a time resolution of ~ 1 ns can be achieved (see Figure 3.4). Finally, the energy is calculated from the amplitude of the pulse, and only that information, together with the timestamp, is pushed to the next part of the chain. This procedure greatly reduces the amount of data.

4.3.2 Data Concentration

FPGA-based devices called *Data Concentrators* (DCs) collect data from several digitisers of the EMC. In particular, they recombine the data from two digitisers reading out the two APDs of the same crystal. Adding the two signals together has the additional advantage of reducing the noise by a factor $\sqrt{2}$, and allows the detection of, and correction for, a nuclear counter hit (see section 4.2). Upon receiving a pile-up waveform from a SADC, the DC sets out to disentangle the signal. An efficient procedure for the pile-up recovery still needs to be implemented and tested on hardware; a possible algorithm has been developed and tested at KVI-CART, Groningen, the Netherlands [70]. The DCs perform an energy calibration, and sort the data according to their timestamp.

Each DC reads out up to 16 SADCs using high-speed 16-Gbps optical fibres, and thus has access to ~ 128 crystals⁷. This opens up the possibility to perform some more advanced processing, see e.g. section 6.3.1. To perform these tasks, the DCs are equipped with a Xilinx Kintex-7 Ultrascale FPGA. About 24 DCs are needed to read out the EMC digitisers (see section 4.3.3), which can be placed away from the detector thanks to the use of optical fibres. The firmware for the DCs is close to its final design, with the time-sorting and energy calibration algorithms already implemented. For the hardware, a couple of options are explored: using existing technologies developed at GSI, or using a custom design with a more compact form factor. The latter option is being investigated in collaboration with the University of Uppsala, Sweden.

⁷In the actual implementation, this number will vary. For example, in the most forward region, the high hit-rate increases the local data production rate in that area, limiting the available bandwidth.

The Data Concentrators also provide synchronisation to the SADCs (see section 4.3.4), and relay commands for e.g. rebooting, and changing thresholds and other settings (so-called *slow control*).

4.3.3 Event Building

Before the event⁸ selection can take place, the data from all detector subsystems in a given time period are collected through what is called a *Burst-Building Network* (BBN). To ensure that all data will be available, the length of this period will coincide with sixteen so-called *bursts*, the time that the beam and the target overlap plus the time gap between these periods of overlap, which is $2.4 \mu\text{s}$ in total. This so-called *superburst* thus spans a $38.4 \mu\text{s}$ period. No design exists for the network at the time of writing, but using estimates for the data rate produced by the Data Concentrators, it is possible to present some options for the EMC part of the network.

First, an estimate for the data rates and the number of devices needed is required. A simulation in PandaROOT was run, where the DPM generator (a description of which can be found in section 5.2.2) with a beam momentum of $5.61 \text{ GeV}/c$ was used to generate 10,000 events. The time span of the simulation, i.e. the time difference between the time of the first and the last hit, which is determined as described in sections 5.3 and 5.3.1, is $497,386.07 \text{ ns}$ at an interaction rate of 20 MHz . As the number of hits cannot be accurately determined after the simulation, because the pile-up recovery was disabled (the reason for this will be discussed in section 6.5), the number of hits with an energy above 2 MeV is used to estimate the data rate. There are 1,069,094 hits, yielding $2.15 \cdot 10^9 \text{ hits/s}$. With 64 bit per hit, the total rate comes out at 137.56 Gbps . As the different sections of the EMC will be read out independently in the first stage, the total rate is divided, based on the location of the hit in the calorimeter. The results are shown in Table 4.2. Each DC has 24 optical links, of which at most 23 can be used for data transfer (as one link is needed to relay control and synchronisation commands), implying that 13 DCs will be needed in the barrel and the backward endcap region, 10 in the forward endcap region, and 1 in the shashlik region.

If the Data Concentrators run the advanced processing algorithms proposed in

⁸An event, in this context, is defined as the information that follows from an interaction between the initial proton and antiproton.

Region	Total data rate (Gbps)	#SADCs	Average data rate per SADC (Gbps)
Backward Endcap and Barrel	37.777	280	0.135
Forward Endcap	52.756	217	0.243
Shashlik	47.029	22	2.138

Table 4.2. Estimated data rates for the different sections of the calorimeter, along with the number of SADCs needed. SADC number estimates are based on the number of crystals with APDs and VPTTs in the forward endcap. In the barrel and backward endcap region, only half of the APDs will need the dual gain readout, so the number of SADCs can be estimated by $0.5 \times \#crystals / (32 \text{ channels (dual range)}) + 0.5 \times \#crystals / (64 \text{ channels (single range)})$. For the Shashlik region, which has 1,404 modules, that are read out by photomultiplier tubes (PMTs) [52], $1,404 \text{ modules} / (64 \text{ channels (single range)}) = 22$ SADCs are needed.

section 6.3.1, and we assume the worst-case scenario where each hit becomes a cluster⁹, then the amount of data will be doubled by the DCs. This can be seen by examining the proposed structure of the data. Each hit is stored in a 64-bit word¹⁰, containing its timestamp, energy, and the channel number on the SADC that read out the crystal. The latter can be used to look up the hit's location in the calorimeter. Clusters are also represented by a 64-bit word, containing its mapped position, diameter, timestamp, and the number of hits in that cluster. Hence, the worst-case scenario occurs when each hit becomes a cluster. As the cluster information is stored in addition to the hit data information, in this worst case, the amount of data doubles. Using these estimates, there will be ~ 13 optical links with a total rate of ~ 75.6 Gbps exiting the DCs of the barrel and the backward endcap, ~ 10 optical fibres with a combined rate of ~ 106 Gbps flowing out of the DCs in the forward endcap region, and ~ 10 optical links with a total rate of ~ 94 Gbps exiting the shashlik DC. This is depicted schematically in Figure 4.7. A conservative estimate for the total output data rate then comes out at 275 Gbps.

Now that the rate estimates are known, three options for the BBN to process the data flow will be discussed.

1. **Using a large data collection network** | The first option is to collect everything in a large, interconnected network, that may or may not do any additional advanced processing, but at least it should collect and preferably sort the data.

⁹Defined in section 6.1; however, its nature is not relevant for the current discussion.

¹⁰A data container with a specific size and structure.

Figure 4.7. Schematic picture showing the expected worst-case data rates produced by the EMC Data Concentrators, and the number of optical links needed to send these data to the Burst-Building Network. The output produced by the digitisers is also shown. Notation: $n \times m$ Optical Links, where n is the number of DCs the data flows to/from, and m is the number of inputs/outputs.

The description that follows will assume that:

- Each network node will have in total 24 SFPs (Small Form-factor Pluggable, a connector for the optical links), around 20 of which will be used as inputs and outputs to free resources for additional processing.
- Equating the number of inputs to the number of outputs, the maximum input data rate will be $10 \text{ inputs} \times 16 \text{ Gbps} = 160 \text{ Gbps}$.
- If the actual data rate falls below this maximum, then the device will have sufficient resources available to perform additional processing.

With these rates and the available inputs, the network nodes can access 13 of the barrel and backward endcap DCs, 10 of the forward endcap DCs, or the shashlik DC. The situation including connections drawn is shown schematically in Figure 4.8. The Figure shows that 3 Level 1 (L1) network nodes will be needed. Each node has (at least) 10 inputs, so the combined data from the L1 nodes can be collected into a single L2 network node, see the centre of Figure 4.8. However, as the full data rate far exceeds the bandwidth of a node, the data will need to be distributed over several nodes. The idea is to let each node collect the data from the same superbunch, as mentioned above. For the so-called phase 1 and phase 2 of the project, the experiment will not have the full luminosity available due to financial limitations on the construction of the FAIR complex and on the $\bar{P}ANDA$ detector itself. At this reduced luminosity, the amount of data that is produced will be considerably less. At an interaction

Figure 4.8. Schematic overview of the topology of the Burst Building Network. The number displayed in the blue (green) arrows for Data in (out) is the number of inputs (outputs) needed, while the number next to those arrows shows the expected data rate of the L1 Nodes. The output rate will be less than or equal to the input rate, depending on whether the nodes perform additional advanced processing as suggested in section 6.3.1. The L1 nodes merge the input data stream chronologically from the time-sorted data they receive from the EMC DCs. The output data is distributed over the L2 Network Nodes depending on the superburst number of the data objects. The L2 nodes can then optionally perform the same additional processing as the L1 nodes on the larger dataset, and send the final data to the Event Building Network, consisting of Compute Nodes. As the L1 nodes of the forward endcap and the shashlik need to have 10 outputs, there should be 10 L2 nodes, adding up to a total of 13 network nodes.

rate of 2 MHz, the data production rate will be about a factor 10 less than what is shown in Figure 4.7 — assuming that the data production rate scales linearly with the interaction rate. However, in the high rate scenario, depicted in Figure 4.8, the number of L1 nodes is already as low as it can be, based on the number of devices that need to be connected. The number of L2 nodes can be reduced to 2 when running at lower rates, because the input rates are in this case low enough to accommodate the data stream, but the number of connections is still too high for everything to be ported to a single node. The L2 nodes may still need to distribute the data over the Compute Nodes by their superburst number depending on the latency induced by the time it takes to build events in them. For phase 1 and 2, the layout in this scenario should be sufficient.

Note: The reality is a bit more complex, as the amount of data that each DC outputs is dependent on the portion of the EMC it is reading out; in the more

forward region, the data rate will be much higher as a result of the forward-boosted particle production in the experiment. The picture sketched here assumes an average data production rate. In the more forward region, network nodes will be able to take less inputs than the ones at larger azimuthal angles. It may be possible to distribute the load on the network more evenly by connecting, for example, forward endcap DCs and backward endcap DCs to the same L1 nodes. However, this will compromise the merging of data, as the different regions are not directly connected. In the future, it should be investigated which distribution delivers the optimal results.

2. **Using a backplane-connected network** | In the final construction, it is considered to interconnect the network nodes via a backplane. Each device would then have 8 inputs, 8 outputs, and 8 connections to the backplane. However, the number of fibres would then exceed the number of available SFPs, so the complete calorimeter cannot be read out by a single backplane-connected network node ensemble. That means that multiple ensembles would need to be used, which would also need to be interconnected at some point. This makes this option undesirably complicated, and will therefore not likely be used.
3. **Using a single high-performance FPGA-board** | Xilinx is currently developing FPGA boards that feature 64 32-Gbps optical links. Provided they have sufficient resources to perform complex tasks on this high-speed data flow, one such board may be capable of handling the output of the entire EMC. All 33 output links from the DCs can be connected to the FPGA-board, as the bandwidth is not pushed to its maximum. However, this board is not yet available at the time of writing, and the current version may not have sufficient resources to perform additional operations on the data. Therefore, it will not be considered for phase 1 and 2, but for phase 3, which should start around 2030, it is conceivable that these devices will be readily and affordably available with sufficient resources (extrapolating the development shown in Figure 4.9), and, hence, may become the preferred option.

4.3.4 Synchronisation

Experiments that rely on a hardware trigger, can also use the trigger to provide synchronisation to the detector subsystems. However, as the $\bar{P}ANDA$ experiment lacks such a trigger, synchronisation needs to be provided by a different means. The high interaction rate further stresses the importance of synchronisation to ensure proper

Figure 4.9. Development of the number of logic cells incorporated in FPGA devices [72].

online event reconstruction. In the $\bar{P}ANDA$ experiment, this is realised by employing a dedicated clock source, called *SODANET* (Synchronisation Of the Data Acquisition NETwork). *SODANET* distributes its clock signal, which has a precision (jitter) of 20 ps, to the Data Concentrators of the detector subsystems using bidirectional links. This allows to monitor the propagation time of the clock signal, so that it can be corrected if e.g. the length of the cables changes due to thermal expansion. The DCs then propagate the synchronisation to the digitisers and other readout devices inside the detector, also called ‘Front-End Electronics’. A recent test with two SADCs and two DCs has shown that the clock signal is stable, as its precision (20 ps) is better than that of the detector subsystem with the highest precision, the TOF detector, which has a precision of ~ 100 ps. In addition, the test proved that the hardware and firmware are ready to use [73].

Chapter 5

Simulation of the $\bar{P}ANDA$ EMC

In order to develop data processing and reconstruction algorithms, and assess (or at least estimate) the performance of the detector subsystems and the detector setup as a whole, a realistic simulation of the detector setup is required. The common way to proceed is to generate data samples on an event-by-event basis, taking possible particle decays, and the detector geometry and response into account. However, this simulation approach cannot deal with certain detector response features, induced by the 20 MHz collision rate, including the occurrence of pile-up and event mixing. To enable these features in the simulation, and to process the resulting data, a second simulation type is developed for the $\bar{P}ANDA$ experiment. In this Chapter, both simulation types and the general path to analyse the data produced will be discussed.

5.1 Software Package

The basic software package that is used for reconstruction and analysis in the $\bar{P}ANDA$ experiment is ROOT. ROOT is an object-oriented framework, developed at CERN using the C++ programming language. Version 5.34 was used in the simulations and data analysis. FAIRROOT and PandaROOT extend ROOT's functionality by adding a large variety of additional base classes (e.g. objects representing the detector hits) and functions (e.g. for modifying new objects and performing subdetector-specific reconstructions). For more details on FAIRROOT and PandaROOT, see [74] and [75]. For FAIRROOT, version 15.11 was used, and the final revision of PandaROOT was 28955 for the non-EMC code ('trunk' version), and 29974 for the (modified) EMC code. Because PandaROOT was under constant development during this research (including contributions by the author), no version can be quoted for use for the entire duration of the research.

5.2 Event-based Simulation

5.2.1 General Description

The standard approach to simulate the detector response and the following expected data flow is to generate an interaction between the initial-state particles. This may lead to the production of a pre-specified (meta-)stable intermediate state, that traverses the detector and undergoes a decay into final-state particles, generating an initial detector response. There may be several more intermediate states, or *resonances*, appearing in the decay chain. The initial detector response in this description envelops a registered hit in a detector element, such as a single photon hitting an EMC crystal or a straw tube registering an ionisation, without any further digitisation or signal processing. This will be referred to as a *hit*. The information that follows from a single interaction between the initial-state particles is called an *event*. An event may therefore constitute elastic scattering, or, more interesting, the production of a resonance. The standard simulation approach hence consists of the sequential generation of events with their corresponding decay chains and hits. Because of this structure, this is referred to as an *event-based simulation* type. The event-based simulation is widely used for the tasks outlined in this chapter's preface.

5.2.2 Technical Details

General | The general type of simulation used is Monte Carlo, commonly abbreviated to ‘MC’. In a MC simulation, results are generated by taking a large number of random samples from a probability distribution for each variable that is relevant in the problem under investigation. For the event generators, a large number of all possible decay modes is generated, according to the parameters of the generator. In the context of the passage of the created particles through the detector, the MC simulation amounts to generating a large set of possible tracks for each particle, that is random as long as it is allowed by e.g. energy and momentum conservation; along with the tracks and the interactions with detector material that may occur, as dictated by the appropriate physical models.

Event Generators | To generate events, different event generators can be used, depending on the type of reaction of interest. The simplest is the Box Generator, which creates particles of a specified type, with a tuneable momentum and multiplicity. The particles created can be set to disperse isotropic, or in a pre-determined range of the solid angle to illuminate a specific part of the detector. There is no limit on the type of particle that can be simulated — it is even possible to define a specific particle. However, for a more realistic test, it is common practice to use particles that are listed in the PDG¹ tables.

A commonly used generator type is EvtGen [76], which uses a table that specifies the decay chain for the particles of interest. If no decay products are specified, the complete set of allowed decays is taken from a PDG table containing all known decays, with corresponding branching fractions. It is possible to force a particle to decay exclusively to one decay mode², allowing to e.g. optimise selection criteria before proceeding to search for a rare decay.

In EvtGen, it is still required to specify which particles need to be created, meaning a clean signal of the desired type will be generated. To add unrelated background contributions to the simulation, making it more realistic, two event generators can be selected in PandaROOT, DPM (Dual Parton Model [77]) and FTF

¹The Particle Data Group, or PDG, is an international collaboration that strives to maintain a complete database containing all that is known about subatomic particles, including hadrons, gauge bosons, leptons, but also some hypothetical particles such as Dark Matter [2].

²This is often referred to as an exclusive study, in contrast to an inclusive study that contains the complete set of allowed decays for the particle of interest. See also: section 2.8.

(FRITIOF [78]). The DPM makes use of the partonic structure of hadrons, treating each Feynman diagram via a two-step process: separation of colour in the collision, and fragmentation of the coloured products. This way, new hadrons are produced from the available initial four-momentum, providing a useful and realistic way to add background contributions to the events of interest. FTF is based on the DPM, and extends it to include some large transverse-momentum-transfer³ processes, like Rutherford parton scattering (elastic parton interaction via the strong force). An important difference is that, although more processes are included, resonances generated with FTF do not have any width.

Transportation | After being generated, the dynamics of the particles and their decay products needs to be calculated. This includes the particle trajectories through the detector, and their interaction with any material they may encounter (including the detector sensitive volumes). The part of PandaROOT that handles this uses Geant4 (v9.2), a toolkit for simulating the passage of particles through matter [79], although the use of the Virtual Monte Carlo packages allows to select other engines without the need to rewrite the code [74]. It is necessary to have the complete geometry of the detector implemented, including any passive material. PandaROOT instructs Geant4 to generate interactions with this material, and then create objects that can be used in the framework, e.g. hits in the EMC.

5.3 Time-based Simulation

The $\bar{P}ANDA$ experiment will feature high collision rates up to 20 MHz, i.e. $2 \cdot 10^7$ interactions per second. Most detector elements require some time to recover⁴ after they registered a hit, and if they are hit a second time within this period, this will either not be registered (dead time) or overlap with the signal of the preceding hit, creating a so-called *pile-up signal*. A pile-up signal hence contains information on more than one hit; this will need to be disentangled if the information on the individual hits and corresponding events is to be retained. Another type of pile-up that may occur is the pile-up of complete events, rather than individual hits. This type of pile-up, which is called *event mixing*, can occur in two ways:

1. A second event follows the first before it has finished (when e.g. the first event creates long-living secondary particles), or starts (almost) simultaneously with

³In this context, transverse with respect to the beam axis.

⁴See section 4.1.1 for a discussion on interactions of particles with detector material that illustrates this statement.

the first one, i.e. it starts within the detector response time of ~ 200 ns, see e.g. Figure 6.1;

2. An event features a high multiplicity of final-state particles that hit the detector at different points with vastly different flight times, like the Forward Endcap and the Shashlik part of the calorimeter.

The second can be recovered by correcting for the flight time before attempting to reconstruct an event. The first one is more challenging. Since there is no way to separate such events *a priori*, it will be up to the event reconstruction algorithms to identify them, taking this effect (i.e. the possibility that there are multiple events in the subset of data under investigation) into account. This behaviour also needs to be incorporated into the simulation, but is not part of the event-based simulation by default, which is why another type of simulation is introduced.

It is clear that for these features, the time stamp of each hit plays an important role. For this reason, this simulation type is referred to as *time-based*. The time-based simulation first generates hits for all events in the standard (event-based) way. After this, the signals of the hits within the same detector element are allowed to interfere for a set period of time, simulating the occurrence of pile-up signals (see Figure 5.2). Event times are allocated from a Poisson distribution, allowing the events themselves to mix as well. More details will be discussed in the next section. An overview of the complete time-based simulation for the EMC is presented in the flow chart in Figure 5.1.

5.3.1 Digitisation and Signal Pre-processing

Analogous to the way the hardware readout system handles sensor output, the simulated hits are run through a digitisation task, `PndEmcFWEndcapDigi` [80], written in the ROOT framework. This mimics digitisation of the signal, and the Feature Extraction and Pile-up Recovery algorithms introduced in section 4.3.1. The process is depicted in Figure 5.2 below. The start time of the first hit of the first event is taken from the MC event information. Assuming that the number of events per unit time is Poisson distributed, the time of the first hit of each following event, t_{event} , is determined using

$$t_{event,i+1} = t_{event,i} + \Delta t_{event}, \quad (5.1)$$

Figure 5.1. Flow chart of the time-based simulation of the EMC. Extensive details on the different parts of the chart can be found in sections 5.3.1 and 5.4. The chart starts with the detector hits that are created by the transport engine (section 5.2.2), from which waveform objects are created by the PndEmcFWEndcapTimebasedWaveforms task. In that task, first raw waveforms are created, which are then modified to include noise and pile-up effects. The PndEmcFWEndcapDigi task then creates digitised hits, digis, out of the waveforms. From the digis, cluster objects are constructed (described in chapter 6), which need to be sorted and packed in bunches again after the process because their time stamps are corrected for the flight time. After this, the PndPidCorrelator tasks assigns the cluster objects to neutral-particle candidates, which the event reconstruction macro uses to reconstruct the original event(s).

with the time difference Δt_{event} between two events obtained from the corresponding exponential distribution:

$$f(\Delta t_{event}) = \frac{1}{\tau} e^{-\frac{\Delta t_{event}}{\tau}}, \quad (5.2)$$

where τ is the mean time between two events (50 ns at an interaction rate of 20 MHz). If there are secondary decays, the corresponding decay times are added to t_{event} . Small values of Δt_{event} may lead to event mixing. Next, to make the time more realistic, to each hit a time stamp t is assigned in the following way:

$$t = t_{event} + t_{tof} + t_{response}, \quad (5.3)$$

where t_{tof} is the time it takes a particle to travel from its creation point to the crystal, and $t_{response}$ includes the time it takes for the shower to develop and the signal pulses to develop and be integrated.

Next, an ideal (i.e. noise-free) waveform is generated for each hit. The points on the digitised waveform are calculated, based on the following formula, that was

Figure 5.2. The generation of simulated digitised hits, inspired by the way the envisaged readout hardware handles the sensor output (section 4.3.1). From the event start time, t_{event} , which follow a Poisson distribution, an event and its decay(s) are simulated. The lifetime of the intermediate state(s), t_{state} , the flight time of the final-state particle, t_{tof} , and the detector response time, $t_{response}$, are added to t_{event} (not shown to scale in the picture). If the active times of the generated waveforms overlap, a pile-up signal is created, which is resolved by using a Moving Window Deconvolution (MWD) filter [81]. The recovered times, t_1 and t_2 , are also added to t_{event} to form the final time stamp, t , of the digis.

obtained by a fit to waveforms that were recorded in an experimental run [70]:

$$f(x) = A \left(\frac{x - t}{\tau_{decay}} \right)^N e^{-N \left(\frac{x - t}{\tau_{decay}} \right)}. \quad (5.4)$$

Here, A and N determine the pulse's amplitude, and τ_{decay} is the decay time. In the simulation, these parameters are set to the values that were determined in the experiment: $N = 1.667$ and $\tau_{decay} = 68.7$ ns. A time-window, t_{active} , is defined for each waveform from the expected pulse length, which is about 700 ns [82]. The waveforms are then convoluted with other waveforms when their respective time windows overlap, and if they lie within the same detector element, allowing the creation of a pile-up signal. After this, noise at a level as determined by experiments with EMC prototypes is added to the thus modified waveforms, and they are subjected to the mimicked Feature Extraction algorithm, which starts with pile-up detection. If no pile-up was found, the waveform is converted back into a (now digitised) hit, referred to as a *digi* to distinguish it from a non-digitised hit. If the algorithm flags the current waveform to have a pile-up, it is run through the Pile-up Recovery algorithm (see section 4.3.1). If successful, the separated waveforms are converted to digis. The waveforms are generated by the `PndEmcFWEndcapTimebasedWaveforms` task. The resulting stream of digis is sorted in time after the first layer of the readout chain, to match the structure of the expected experimental data flow.

If the pile-up recovery failed to separate the overlapping pulses, they will be seen as a single pulse. The first rise time and the maximum amplitude will be used to determine its time stamp and energy, respectively. If this also fails, the waveform is discarded. In either scenario, this will mean that some information will be lost, which may cause the event reconstruction to fail as well. Consequently, the information in that event will not be stored.

5.4 Event Reconstruction

To reconstruct an event from the (simulated) data using the PandaROOT framework, a series of steps is needed. The general procedure is depicted in Figure 5.1. As an example, the reconstruction of the h_c in the decay chain $p\bar{p} \rightarrow h_c \rightarrow \gamma\eta_c \rightarrow \gamma\eta\pi^0\pi^0 \rightarrow 7\gamma$ will be discussed. Because it only has photons in the final state, this channel can be reconstructed using the information from the EMC only. The first step, after the digis have been generated from the simulated data as described above, is to group them into clusters to form the photon candidates (see chapter 6). A flight time correction is applied to the clusters to counter any dispersal of clusters from the same event due to the varying flight distances to the different sections of the detector. For example, the flight time for a photon from the interaction point to the barrel part of the EMC is about 3 ns, while it would take about 30 ns to reach the shashlik calorimeter. After re-bunching the clusters using `PndEmcPackClusters`, which is needed as the time stamps have been modified after the preceding step, the cluster objects are run through the PID task, `PndPidCorrelator`, which takes the output of the reconstruction task(s) and creates neutral- and charged-particle candidates. In general, it will use the output of other subsystems to obtain this information. However, for the channel at hand, since there are only photons expected, all cluster objects are assumed to represent photon candidates, and are added to the neutral-particle candidate list.

In the final step, the decay chain is reconstructed. This is done from the bottom up (section 2.8), i.e. starting from the photon candidates and reconstructing their parents, then their grandparents, etc. The `Rho` package in PandaROOT is used to achieve this, which allows to recombine particles in a straightforward way, taking double counting into account. Specifically, it starts by searching for π^0 and η candidates, by looking at all combinations of two photons. These are identified by their masses from the invariant mass of the two-photon combinations, see for example Figure 5.3. The in-

Figure 5.3. Invariant mass spectrum of two-photon candidates. The peaks corresponding to the π^0 (at $M = 135 \text{ MeV}/c^2$) and the η (at $M = 548 \text{ MeV}/c^2$) are clearly visible.

variant mass is calculated using $M = \sqrt{E_{tot}^2 - p_{tot}^2}$ (chapter 2.8), with E_{tot} the sum of the photon energies, and p_{tot} the sum of their momenta, whose components can be calculated⁵ using $p_i = i \cdot E / \sqrt{x^2 + y^2 + z^2}$, with $i = \{x, y, z\}$ the coordinates of the cluster in the laboratory frame. The energy and momentum are often combined into the four-momentum (p_x, p_y, p_z, E) ; the calculation of M is then done using the components of this vector. To be eligible for the procedure, photon candidates must have a minimal energy of 30 MeV. This minimal energy was chosen to reject unrelated noise hits. This cut will not negatively impact the result of the reconstruction, as the least-energetic photons will come from a π^0 , and will have an energy of at least $\sim 65 \text{ MeV}$ (half of the π^0 rest mass). π^0 and η candidates are selected from the spectrum in Figure 5.3 by imposing a mass window. This window covers 20 MeV around the nominal mass for π^0 s, and 25 MeV around the nominal mass for η s, indicated by red boxes in Figure 5.3. Once the π^0 and η candidates have been found, η_c candidates are searched for by handling all possible combinations of two π^0 s and an η in the same way. Again, a mass window is applied to the obtained invariant-mass spectrum to select the η_c candidates, shown in Figure 5.4. A window of 500 MeV was chosen to perform a loose selection. Finally, from the η_c candidates and the remaining photon candidates, h_c candidates are formed. A kinematic fit is applied to clean up the spectrum, as can be seen in Figure 5.5, which rejects h_c candidates whose four-momentum does not match that of the initial $p\bar{p}$ system.

⁵ Assuming the cluster comes from a photon, its trajectory is a straight line from the interaction point.

Figure 5.4. Invariant mass spectrum of η_c candidates, from combinations of an η and two π^0 candidates. The red box indicates the mass window that is applied to select η_c candidates.

Figure 5.5. Invariant mass spectrum of h_c candidates, from combinations of η_c and photon candidates. Because many photon combinations are possible, the h_c peak (at $M = 3.525$ GeV/c^2) is quite broad. From the very low energetic photon candidates, an artifact at the η_c mass (at $M = 2.983$ GeV/c^2) can be seen when taking all combinations, shown in the black histogram. After applying the kinematic fit, as seen in the filled green histogram, the spectrum gets ‘cleaner’.

The time-based simulation as described above will be the simulation type that is used whenever simulated data in the context of the $\bar{P}ANDA$ experiment is referred to, unless specified otherwise.

Chapter 6

Cluster Finding

Before the events themselves can be reconstructed for the event selection, the final-state particles need to be reconstructed. The need for this becomes evident when considering that e.g. photons are detected by the EMC by depositing their energy through creating a particle shower, that will in general extend out through several crystals. Therefore, single-crystal hit information is not sufficient to access the required information on the final-state photon that hit the EMC, calling for a reconstruction of this photon.

As emphasized in Chapter 1, one of the goals of this work is to develop an algorithm to perform this type of reconstruction in real-time, i.e. while the detector is gathering data. Because data are produced at a rate of 138 Gbps in the EMC (see section 4.3.3), the algorithm is required to be fast and to use as little resources as possible. The output of the algorithm, in the form of e.g. reconstructed photon candidates, provides vital information for the online event selection process. Hence, it is of paramount importance that a reasonable resolution and efficiency are obtained in addition, i.e. close to the design values (see section 3.3.2.5).

6.1 Final-State EM Particle Reconstruction

The focus of this chapter will be on particles that are fully stopped by the EMC, such as photons, electrons and positrons, referred to as *EM particles* in this work. Some other relatively long-living particles such as charged pions and muons, that decay via the weak interaction, can also be detected, but will often not deposit all their energy in the EMC. In those cases, the energy depositions in the EMC can be used to assist in particle identification (PID). When an EM particle enters the calorimeter, it deposits its energy by creating a particle shower in the scintillator crystals (see also section 4.1.1). For most final-state particles of interest, the shower will span multiple crystals, leading to groups of hits in the EMC. Such a group of hits is called a *cluster*. Hence, reconstructing a final-state EM particle amounts to combining the correct group of hits in the detector. This process is called *cluster finding*. The occurrence of pile-up signals and geometrical overlap between clusters as a result of the 20 MHz interaction rate and forward-boosted particle production complicates this task significantly (see also: Figure 4.6).

6.2 Cluster Finding - Input

The pile-up of signals has already been taken care of, at the point that the cluster-finding algorithm sets to work, by the pile-up recovery algorithm for digitised hits, digis, in the EMC front-end electronics. The Poissonian nature of the event generation should allow a grouping of digis according to their time stamp, see Figure 6.1. Hits belonging to the same event have an increased probability to end up in such a group, called a *time-bunch*.

As an input, the cluster-finding algorithm takes a stream of digis grouped in time

Figure 6.1. Formation of the bunched structure in a stream of digis, schematically indicated by vertical lines. Different events are indicated by different colours, and the resulting stream is shown in black.

(see the start of the flow diagram in Figure 6.3). This stream of digis can either be generated from a MC simulation or be comprised of actual detector data. The bunched structure shows up clearly in Figure 6.1. It is clear that the division of the data stream into bunches requires an optimisation. The cutting into bunches will be done by checking the time difference, $\Delta\tau_{diff}$, between any two consecutive digis. Choosing a lower limit value, $\Delta\tau$, for this difference will result in smaller time-bunches. A too low value of $\Delta\tau$ may cause parts of the same cluster to be spread over multiple time-bunches, hindering event reconstruction. A too high value will compress multiple events into a single time-bunch, increasing the probability of geometrical overlap, hindering assignment of clusters to their proper parent event, and possibly overflowing the buffers in the readout hardware by loading too many hits. One way to prevent the latter is to implement cluster finding in a distributed way.

6.3 Distributed Cluster-Finding

6.3.1 Motivation

Distributed cluster-finding in the context of this work amounts to search for clusters in subsets of the calorimeter, instead of in the complete device. This distributes the load of finding clusters, freeing resources, and thereby allowing (in principle) faster event reconstruction at the final stage. In the readout chain, as explained in section 4.3.2, each Data Concentrator reads out a part of the calorimeter. Using the digi-data it has at its disposal, the Data Concentrators could already look for clusters in the geometrical portion of the EMC that they are reading out. As it is possible that clusters are spread out over multiple of these portions (see the left part of Figure 6.2), the clusters found at this stage, which will be called *preclusters*, may need to be merged later on to find the complete clusters. That can be done by the Burst Building Network, or at the Compute Node stage (see section 4.3.3).

6.3.2 Precluster Merging

Once preclusters have been constructed by the Data Concentrators, they can be merged in the same way as the individual digis are merged into preclusters (see section 6.4.1), i.e. by searching for neighbouring preclusters. Hence, in order to do this, the 4D position (t, x, y, z) of each precluster is needed. To speed up the calculations, the crystals will be mapped onto a 2D grid, where the granularity is such that each crystal is a ‘pixel’ in this map. This way, only three numbers need to be calculated for the position of the preclusters: its timestamp, t , and its mapped

Figure 6.2. Schematic view of a section of the EMC, divided into two subsections that are read out by different Data Concentrators (DC 1 and DC 2). Left: The red/blue cluster falls into two DCs, and will hence form two preclusters, that will need to be merged later on, while the black ones formed completely in one DC. Right: The size of a precluster in the simulation can be calculated in three ways: (A) Form the circumcircle corresponding to the polygon shape; (B) Take the smallest box that fits the precluster; (C) Take a square box with sides equal to the largest dimension of the precluster. For the circle, the size is equal to the radius of the circle, and for the boxes, the size is equal to half the height (or width, depending on the method) of the box. In each scenario, the bounding shapes (indicated by a yellow line) of the two shown preclusters touch or overlap, and hence they will be merged to one larger cluster, even though this might not always be desirable.

position, (X, Y). For the time stamp, the time stamp of the most energetic digi¹ in the precluster is used. If the size (see Figure 6.2 (right)) is specified in addition, the precluster merging algorithm can be implemented in a very straightforward way — by merging any two preclusters for which the spatial distance between them is less than or equal to the sum of their respective sizes, and by requiring them to be close in time as well. Figure 6.2 (right) shows available methods for obtaining the precluster size. Methods (A) and (C) will often overestimate the size, as they take the largest dimension of the precluster as its size; this may affect the final performance most notably with asymmetrically-shaped preclusters. However, the error made in this procedure, e.g. the unjust merger of two preclusters, will not lead to any disadvantages that were not already present in the original (i.e. non-distributed) algorithm, where two digis are assigned to the same cluster if they are neighbouring. To get a more accurate candidate for later reconstructions, the spatial position of the final clusters is calculated from the energy-weighted centroid coordinates of the participating crystals using a logarithmic weighting. It is noteworthy that the result of this procedure relies on a proper matching of the gain of the photosensors (section 4.2), and on a proper energy calibration.

The effect on the yield for the channel $h_c \rightarrow \gamma\eta_c \rightarrow \gamma\eta\pi^0\pi^0 \rightarrow 7\gamma$, determined following the method described in section 5.4, when using the different size methods

¹The crystal with the highest energy deposition.

Coordinate sys- tem + geometry	Real circle	Mapped circle (A)	Mapped rect- angle (B)	Mapped square (C)
Yield	740	456	461	413
Number of clusters found	35,104	35,866	35,775	35,329

Table 6.1. The number of clusters found and the yield (from a total of 5,000 events) in the channel $h_c \rightarrow \dots \rightarrow 7\gamma$ at an interaction rate of 200 kHz, for the different precluster size methods in the right side of Figure 6.2. The yield was the number of h_c mesons, obtained using the method described in section 5.4.

for preclusters is shown in Table 6.1. The *yield* is the number of states of interest that were recovered. The significant drop in yield when changing to the mapped coordinate system can most likely be attributed to the imprecise mapping that is currently implemented in PandaROOT. In particular, the edges of some sections are not always properly connected. This can be seen in the number of clusters that were found in Table 6.1. The methods that use the mapped coordinates (X, Y) find a larger number of clusters, while showing a lower efficiency. This implies that some preclusters, that ought to be merged, were not merged. Breaking up these clusters leads to the drop by almost 40%, when switching from the real (laboratory) coordinate system, which does not have most of these discontinuities, to the mapped coordinate system. Therefore, in the future, with an improved crystal mapping, where the edges of the sections are properly connected, results are expected to improve. For the time being, a circle in the real coordinate system in the laboratory frame (‘Real circle’ in Table 6.1) will be used to evaluate the performance of the algorithms discussed in this chapter.

6.4 Cluster Finding - Algorithms

The art of cluster finding is not new. The algorithm that is currently implemented in the PandaROOT software package performs well, but was not designed for online usage. It was not optimised to run fast and to use as little resources as possible. For this reason, new, so-called ‘online’ algorithms have been developed. The online algorithms and the existing algorithm will be described in this section.

Figure 6.3. Flow chart of the online cluster-finding algorithm, describing each step taken by the algorithm. First, digis that have an energy below some threshold (in the current setting 3 MeV) are discarded. Then, a loop over all pairs of digis determines neighbourhood relations and creates the digi ‘map’ in the top right of the Figure (in reality this has been flattened to a 1D array). The algorithm proceeds to assign digis to clusters using this map, creating the cluster map in the bottom right, and then uses this map to build the cluster objects. In the final step, the cluster properties are determined.

6.4.1 Description and Parameters

The online cluster-finding process is depicted graphically in Figure 6.3. Once the input stream of digis is passed to the algorithm, it starts to loop over all pairs of digis to search for neighbouring ones. Digis are considered to be neighbours if they are not only close in space, but also in time. It follows that the distance, Δd , up to which digis are considered to be neighbours and the time separation, Δt , are tuneable parameters in the algorithm and need to be optimised. It is important to note that Δt is the time difference between a pair of digis, that are not necessarily consecutive in the time domain, as is the case for $\Delta \tau$. Tuning Δd allows to include digis that are not directly neighbouring to another, so that so-called *split-offs* can be reabsorbed into the cluster. Split-offs can be categorised in two types:

1. A relatively low-energy particle in the particle shower scatters out of the crystal where the shower is developing, skips² one or more crystals, and starts a new particle shower in a neighbouring crystal, or

²Here, skipping can mean no interaction in that crystal (because low-energy particles with an energy above the detection threshold (~ 3 MeV) enjoy a longer mean free path), or an interaction that leads to an energy deposition below the energy detection threshold.

2. Closer to the interaction point, a charged particle emits a bremsstrahlung photon, some particle interacts with passive material in the detector to emit an electron, or a photon undergoes pair production.

The second type technically does not fit the conventional definition of a split-off, but produces a similar signal in the EMC. They may be identified by using information from other subsystems. Type 1 split-offs can be (partially) recovered by setting Δd to include next-to-nearest neighbours. On the other hand, increasing Δd will also increase the probability that two distinct clusters, that happen to be close, will unjustly be merged. Optimising Δt is necessary to ensure that clusters that were spawned closely after one another in the EMC, leading them to pile up, can be disentangled. Summarising, Δd can be used to separate clusters that are close in time (or even simultaneous), and Δt to disentangle those that are geometrically overlapping. An optimisation study on the parameters will be discussed in section 6.5.

6.4.2 Algorithm Implementations

The algorithm described above (called `PndEmcMakeClusterOnline` in the PandaROOT package) takes the input time-bunch and loops over all member digis to identify clusters, as shown in Figure 6.3. For distributed cluster-finding, there are two options: single-pass and double-pass:

- In the single-pass case (called `PndEmcDistributedClustering` in the PandaROOT package), the algorithm takes the input time-bunch, and loops per (virtual³) Data Concentrator over its member digis to identify preclusters. Then, within the same time-bunch, it loops over the preclusters from all DCs that are active in that time-bunch and merges them (if needed) into clusters.
- In the double-pass algorithm (called `PndEmcMakePreclusters` and `PndEmcMergePreclusters` in the PandaROOT package), the first step, i.e. identifying preclusters, is the same as for the single-pass case. For the second step, all preclusters from all time-bunches are put in a stack, and new time-bunches are formed using the time stamps of the preclusters. Within these new time-bunches, the algorithm loops over all member preclusters and merges them into clusters if they are neighbouring (section 6.3.2).

The single-pass case will be difficult to realise in the real hardware, as each DC will define its own time-bunch, based on the subset of the data it receives. The

³A ‘virtual’ Data Concentrator (DC) is a subsection of the EMC in the simulation, that would correspond to one DC.

only foreseeable way would be to not let the DCs run any cluster-finding algorithm; i.e. only start looking for clusters when the data from the complete calorimeter have been collected somewhere. The double-pass algorithm, however, fits this design and can easily be implemented in the DCs and the Burst Building Network or Compute Nodes (see also section 4.3.3). All algorithms described above follow the flow chart in Figure 6.3, where in the distributed cluster-finding case, preclusters are created in the fashion described there, instead of clusters. In that case, clusters are formed in the same way in the next step, using preclusters instead of digis, but without setting an energy threshold.

The ‘default’ algorithm that is currently implemented in PandaROOT (called `PndEmcMakeCluster` in the PandaROOT package) loops over all digis in the input time-bunch, and checks if they should be added to an existing cluster. If so, then they are added to that cluster, and if not, then they form the seed for a new cluster. Checking whether a digi should be added to an existing cluster requires a loop over all current members of that cluster, for each new digi in the time-bunch. This large number of loops would create more latency in the online environment, making this cluster-finding method less suited for online usage. Therefore, it will not be used in the readout system, but can still be used in the offline analysis of data.

6.4.3 Bump Splitting

Even if Δd is set to its minimum, separate clusters that ‘touch’ each other will be identified as a single large one by the cluster-finding algorithm, as their digis are neighbouring to each other. To rectify this problem, at a later stage, a procedure called *bump splitting* is used. In this framework, two joined clusters form a single geometrical entity with (usually) local maxima, or ‘bumps’. In PandaROOT, the bump-splitting algorithm searches for these local maxima, and then re-assigns energy contributions from participating crystals in accordance with the energy depositions in those maxima. Like the existing cluster-finding algorithm, the existing bump-splitting algorithm was not designed for online use. It loops over all digis and checks if they are a local maximum by looping over all other digis in the cluster and checking their energy deposition. In addition, the energy gradient is calculated to strengthen the claim that the current digi is a local maximum in the cluster. Although untested, it stands to reason that the large number of loops that the algorithm needs, in addition to the calculation steps that are required to obtain the result, makes the current implementation unsuited for use in the readout chain. One suggestion for a faster, online replacement is to test

whether the two most energetic digis in a cluster are neighbours. The logic behind this is that the central digi in a cluster will have the highest energy deposition, and the energy depositions decrease radially outwards. Hence, if the two highest energy depositions in a cluster are not adjacent, it is likely that it consists of two clusters that overlap. In that case, the cluster can be split up, with the energy distributed over the daughters as

$$E_{d1} = \left(\frac{E_{h1}}{E_{h1} + E_{h2}} \right) E_{clus} \quad \text{and} \quad E_{d2} = \left(\frac{E_{h2}}{E_{h1} + E_{h2}} \right) E_{clus}, \quad (6.1)$$

where $E_{d1,2}$ is the energy assigned to the daughter clusters, $E_{h1,2}$ is the energy of the two most energetic digis (h_1 and h_2), and E_{clus} is the total energy of the unsplit cluster. The position and timestamp of the two daughter clusters can be set using the position and timestamp of h_1 and h_2 , respectively. In principle, it is possible to expand this method to search for three or even more overlapping clusters. However, it is highly unlikely that those situations will occur, as the overlapping of clusters will generally only occur when e.g. a π^0 with a high momentum decays to two photons, which will then have a very small opening angle. Otherwise, clusters are spread isotropically. This algorithm has not yet been implemented into the PandaROOT framework, because it falls outside of the scope of this work. It is, however, recommended to be investigated in future works.

6.5 Performance Assessment by PandaROOT Simulation

6.5.1 Parameter Optimisation

As mentioned before, the performance of the algorithm depends on the three parameters $\Delta\tau$, Δt , and Δd . They are checked using one or both of these two exclusive decay channels:

- $p\bar{p} \rightarrow \gamma\gamma$ with a beam momentum of 1.5 GeV/c, so $\sqrt{s} = 2.251$ GeV/c², where s is the square of the centre-of-mass energy. This simple system allows to check the performance on a very basic level, i.e. it is possible to see if the cluster energies are constructed properly, without relying on complicated reconstruction algorithms. Because particle production is boosted in the forward direction, the cluster energies themselves cannot be used to check the performance of the cluster-finding algorithms. Therefore, this simple two-particle channel was chosen, because it allows the construction of a quantity that is independent on the frame of reference.

Figure 6.4. Two-photon invariant-mass histogram of the channel $p\bar{p} \rightarrow \gamma\gamma$ at a beam momentum of 1.5 GeV/c, comparing the results when different neighbour distances Δd are taken, in units of the number of crystals. In the inset, the fit results are quoted for the different cases, with the yield being the number of entries in the main peak, and σ its width. The interaction rate was set to 20 MHz. No background subtraction has been applied, but a minimal energy of 30 MeV was imposed on the clusters.

- $p\bar{p} \rightarrow h_c \rightarrow \gamma\eta_c \rightarrow \gamma\eta\pi^0\pi^0 \rightarrow 7\gamma$, which has 7 photons in the final state. This channel was chosen because it can be fully reconstructed using information from the EMC only, and because it features a high photon multiplicity. This increases the probability for pile-up, so the recovery performance for those events can also be checked. It must be noted, however, that this decay channel will present an extreme scenario, in which each event produces (at least) 7 photons.

Unless stated otherwise, all spectra and results are obtained using the single-pass distributed cluster-finding algorithm, with logarithmic energy-weighting and the 4D position of hits, which is expected to be the most accurate, and can, hence, serve as a benchmark for the online algorithms. Although the double-pass version fits the design of the readout chain better, the single-pass version is easier to use in testing, and delivers very comparable results to the double-pass algorithm. (e.g. Figures 6.13 and 6.14 demonstrate this). The first parameter to be optimised is Δd , the distance up to which hits are considered neighbours in space. Figure 6.4 shows how the two-photon invariant-mass spectrum of the two-photon channel depends on this parameter. This spectrum was obtained by adding the four-momenta of all combinations of two clusters within a time-bunch. The single-pass distributed cluster-finding algorithm was used with $\Delta\tau$ set to 10 ns to construct the clusters, each time using the same data set as an input, while setting Δd to different values. As can be seen in the spectra, the effect of changing Δd is negligible. This is also reflected in the yield and the width of the peak, which slightly drop when increasing Δd . Considering that

choosing a larger value of the parameter may lead to an increase in the number of overlapping clusters, as discussed in section 6.4.1, the conclusion is that Δd should be set to 1.

With only two photons, and hence only two high-energetic clusters, the two-photon channel does not provide adequate complexity to identify the optima for the other parameters. Looking at Table 6.2, which shows the number of reconstructed h_c mesons (section 5.4) out of 5,000 that were simulated for an example set of combinations for Δt and $\Delta \tau$, it can be seen that the success of the reconstruction critically depends on the choice of Δt and $\Delta \tau$. This follows naturally from the higher hit-multiplicity in that channel; the more hits, the higher the chance of having a couple of them close together in time.

The way the time stamps of the digis are constructed was described in section 5.3.1. Although the pile-up recovery algorithm was tested on a prototype of the EMC [70], and, hence, conceptually works, its current implementation in the simulation framework is not functioning properly and has been disabled. To compensate for the occurrence of pile-up at higher rates, where the probability for pile-up is high (about 9% at 20 MHz), an idea was tested to artificially reduce the length of the waveforms by setting τ_{decay} to 5.27 ns in Eq. 5.4, reducing the length of the waveform from ~ 700 ns to ~ 150 ns. The shorter pulses naturally have a smaller probability to pile up, and so this ‘trick’ mimics the recovery of pile-up. To ensure that the shape of the shorter pulses is still properly recorded, the sampling frequency was increased from 80 MHz to 400 MHz. However, the move to shorter waveforms negatively impacted the efficiency of the algorithms, reducing the yield

$\Delta \tau$ (ns) \ Δt (ns)	5	10	15	20
5	24	29	26	24
10	49	63	78	63
15	148	181	174	124

Table 6.2. Yields (out of 5,000 events) in the channel $p\bar{p} \rightarrow h_c \rightarrow \eta_c \gamma \rightarrow \pi^0 \pi^0 \eta \gamma \rightarrow 7\gamma$ at an interaction rate of 20 MHz, for some range of combinations of Δt and $\Delta \tau$ (incomplete, see Figure 6.10 for a more complete picture). Yields are obtained by applying the reconstruction method in section 5.4 to the collection of constructed clusters, which are assumed to be photon candidates.

by as much as a factor two, and therefore, it was decided to abandon this approach. It was found that the digitisation code does not function properly at different sampling frequencies and values for the pulse shape at the level of pulse detection. Initial tests did not reveal these flaws, as the code produces reasonable results for the default set of parameters, and was reconciliated with experimental data at those settings [80]. Only after changing the parameters, as was done for the pulse length, did the bugs start to manifest themselves. Hence, when using other parameters than the default ones, the results from the simulation may not be very reliable. In the near future, this part will require extensive revision. It must be noted, that the digitisation code itself was assumed to be functioning, and falls outside the scope of this work. Therefore, in all the situations that follow, the default set of parameters is used.

To assist in the assignment of clusters to their correct event, an idea was tested to apply a flight-time correction to the digis. Assuming all clusters to be created by photons, originating from the interaction point, the correction is $t_{tof} = d/c$, with d the distance from the digi to the interaction point. This way, clusters from the same event should group in the time domain. However, this correction proved to severely compromise the cluster reconstruction efficiency, and was therefore abandoned. Most probably, this drop is a result of the reduction in the temporal separation between digis within an event, which increases the probability to assign them to a different cluster than they belong to. In addition, it is difficult to implement such a correction in the readout hardware, because it involves a resorting of the data. Abandoning the correction will preserve the time sorting in the data stream. Therefore, it was decided to not further investigate this option. The flight-time correction on the cluster level is still applied (section 5.4).

To find the value to choose for $\Delta\tau$, it is possible to look at the time difference between consecutive digis. Assuming digis from different events to be more separated, the optimum can be found by determining the number of hits having $\Delta\tau_{diff} > \Delta\tau$, and tune $\Delta\tau$ until the resulting integral approximates the number of events. To account for some inherent loss of efficiency in this ‘automatic’ procedure, the cut was set to 98% of the number of events generated. At this setting, the efficiencies at lower interaction rates equal that of the event-based approach after reconstruction. Two time-difference spectra at low interaction rates are shown in Figure 6.5. In the plot in the left side of the figure, a peaking structure is clearly visible, which is absent in the right-hand side. Further investigation revealed that these peaks are exactly 12.5 ns

Figure 6.5. Time difference between consecutive digis from 5000 events at an interaction rate of 2 kHz (left) and 200 kHz (right), for the decay channel $h_c \rightarrow \dots \rightarrow 7\gamma$. The spectra consist of a convolution of two contributions; a ‘slow’ part, corresponding to the mean time between events (see section 5.3.1), and a ‘fast’ part, originating from the time-distribution within an event. As can be seen in the ‘overflow’ box, at a rate of 2 kHz, most of the slow part of the distribution is shifted beyond the range of the histogram, whereas the same part of the 200 kHz distribution starts at smaller time differences, as expected.

apart, the inverse of the sampling frequency. The reason for its preferred occurrence at the lower rate can be attributed to rounding errors in the calculations — at lower rates, the time stamps quickly assume large values, because for the same number of events, the total run time of the simulation increases drastically. The floating-point numbers used to store these numbers are not long enough to store them with the required accuracy. However, as stated before, the digitisation part of the simulation code is up for extensive revision. Because the peaking structure manifests itself only in the simulation at a low interaction rate, where events are well-separated, its impact on the results is expected to be minimal. This can be checked by comparing the results at an interaction rate of 200 kHz with the 2 kHz case. At 200 kHz, the probability for pile-up is still very low; hence, the results should be nearly identical. Figure 6.6 shows that this is indeed the case. The values for $\Delta\tau$, that were automatically determined using the integration method described at the beginning of this paragraph, are shown in Table 6.3 for the two channels. The variation in the structure in the time spectra at different interaction rates (Figure 6.5) indicates that $\Delta\tau$ naturally depends on the interaction rate. Hence, the value to use for this parameter will be rate-dependent.

Channel	2 kHz	200 kHz	2 MHz	20 MHz
$p\bar{p} \rightarrow \gamma\gamma$	385	239	97	13
$h_c \rightarrow \dots \rightarrow 7\gamma$	257	173	64	9

Table 6.3. Automatically-determined values for $\Delta\tau$, in ns, at different interaction rates, for the two channels.

Figure 6.6. Yields for the channel $h_c \rightarrow \dots \rightarrow 7\gamma$ for different methods to vary Δt , after $\Delta\tau$ was automatically determined (see Table 6.3). The numbers quoted inside the bars show the absolute values of Δt that were used to obtain the results. Yields were obtained using the method in section 5.4, as before.

A case may be built to claim that the optimal value to choose for Δt would follow from the complement of the integral, i.e. slightly smaller than $\Delta\tau$, but the time stamps of digis belonging to the same cluster need not follow each other chronologically in the data stream. Because the time resolution is energy dependent, the time stamps of low-energetic digis, in particular those at the edge of the shower, have a large uncertainty (see Figure 3.4). Clusters within the same time-bunch should have their time stamps closely packed, compared to the time between the time-bunches, at least at lower rates. Because a cut on such ‘events’ has already been made by creating the time-bunches, and the clusters as a whole will exhibit geometrical dispersion, one would want to take Δt large, to ensure that all digis are absorbed in the clusters. However, a too large value will increase the probability to absorb unrelated noise hits into the cluster. The bar histogram in Figure 6.6 shows how the yields depend on variations of Δt . The maximal yields for the 2 and 200 kHz case match with what is recovered in an event-based simulation. Judging from the figure, the optimal choice for Δt is $\Delta\tau + 25$ ns. This optimum was chosen by applying a weighted average to the yields of the two channels, with weight factors 5 for the $h_c \rightarrow \dots \rightarrow 7\gamma$ channel and 1 for the $p\bar{p} \rightarrow \gamma\gamma$ channel (not shown in Figure 6.6), respectively. The reason for assigning a much higher weight to the $h_c \rightarrow \dots \rightarrow 7\gamma$ channel is that the choice of Δt and $\Delta\tau$ affects the yield for channels with a higher photon-multiplicity more strongly. The weight factor itself is chosen arbitrary, but assigning an even higher weight to the $h_c \rightarrow \dots \rightarrow 7\gamma$ channel does not influence the result.

Figure 6.7. Yields out of 5000 events for the channels $h_c \rightarrow \dots \rightarrow 7\gamma$ and $p\bar{p} \rightarrow \gamma\gamma$ for variations of $\Delta\tau$, with $\Delta t = \Delta\tau + 25$, where for the former channel, two independent data sets were used to check the consistency of the results, shown in red and blue. The yields for the latter channel are shown with an offset, for easier comparison. Yields were obtained using the methods described in the text.

Figure 6.8. Method to correct for combinatorial background for the channel $p\bar{p} \rightarrow \gamma\gamma$ (see text). The solid red line indicates the linear fit, the dashed black line shows the extrapolation of the linear fit to the starting point of the area of the peak to be integrated, and the gray triangle shows the correction that will be subtracted from the results of the integral.

The results in Table 6.3 suggest that a photon-multiplicity-dependent cut should be used; however, this is highly undesirable, as it would mean that particle decay modes that do not exhibit the EM-particle multiplicity of the selected optimisation would be discriminated against – even before any reconstruction took place. To investigate if a more global optimum exists, the results were re-evaluated at different time cuts. These include the optima in Table 6.3, and some intermittent values. The results are shown in the bar chart in Figure 6.7. The data in Figure 6.7 imply that the combination $\Delta t = 282, \Delta\tau = 257$ ns works well for all tested channels at an interaction rate of 200 kHz and below, and $\Delta t = 122, \Delta\tau = 97$ ns is the optimum at an interaction rate of 2 MHz. The yield of the $p\bar{p} \rightarrow \gamma\gamma$ channel was corrected for combinatorial background. This was done by applying a linear fit to the spectra up to $1.9 \text{ GeV}/c^2$, just before the peak is expected to start, and then subtracting a triangular area under the peak, as depicted schematically in Figure 6.8. The integration bounds in the spectra were fixed to $\{2.12, 2.34\} \text{ GeV}/c^2$. The yield of the $h_c \rightarrow \dots \rightarrow 7\gamma$ channel was obtained following the recipe of section 5.4.

Note that the 20 MHz scenario is excluded in most of the above discussion. This is because in this case, the time-cuts, that were automatically determined, do not yield acceptable results, due to pile-up and event mixing effects. Figure 6.9 shows the distribution of time stamps of hits for the interaction rate of 20 MHz. The absence of a tail at larger time separations is indicative of the imminent failure of the automatic determination of the time cut. Because the time-distribution spectra do not allow to set a ‘natural’ cut, the only way to determine the optima for Δt and $\Delta\tau$ is to try different combinations of them. The results are shown in Figure 6.10, for the two channels

Figure 6.9. Time difference between consecutive digis of 5000 events at an interaction rate of 20 MHz, for the decay channel $h_c \rightarrow \dots \rightarrow 7\gamma$. Unlike at lower rates, see Figure 6.5, the slow part of the time-difference distribution, which was used to automatically determine the time-cuts, fully overlaps with the fast part, corresponding to the mean time between events, which depends on the interaction rate.

mentioned in the beginning of this section.

From the maxima in Figure 6.10, the optimal combination of values for $\Delta\tau$ and Δt is determined to be 15 and 40 ns, respectively, although values close to these numbers yield comparable results. It must be noted that the pile-up recovery features were disabled due to the unreliable performance of the current implementation; therefore, this result may not be entirely accurate. It is recommended to update and re-evaluate the results in Figure 6.10 after pile-up recovery has been restored.

Summarising, the optimal values for the parameters of the cluster-finding algorithms are rate-specific, except for the geometrical parameter, which should be set to $\Delta d = 1$. The other parameter optima are found to follow $\Delta t = \Delta\tau + 25$, with values $\Delta t = 40$ ns, and $\Delta\tau = 15$ ns for 20 MHz, $\Delta t = 122$ ns, and $\Delta\tau = 97$ ns for 2 MHz, and $\Delta t = 282$ ns, and $\Delta\tau = 257$ ns for lower rates, which will be the values that are used in the following analysis, unless stated otherwise. In the actual experiment, the instantaneous rate may vary due to fluctuations in the target thickness. Depending on the time scale of these fluctuations, it may be feasible to use the automatic procedure described in this section to update the values of $\Delta\tau$ and Δt online, by filling time-difference histograms and integrating them.

6.5.2 Performance of the Cluster-Finding Algorithms

The first performance check is the efficiency for single-photon identification. To test this, data sets of 10,000 events were generated, where mono-energetic single photons were fired isotropically at the forward part of the calorimeter with energies of 300 MeV, 500 MeV, 1 GeV, or 2 GeV using the Box generator (see section 5.2.2), as the forward-boosted particle production forces most decays in that direction. The

Figure 6.10. 2D histograms of the optimisation study for combinations of the parameters Δt and $\Delta \tau$ at an interaction rate of 20 MHz. Yields are shown as a percentage of the number of events generated for the two channels. For $(\Delta \tau, \Delta t) = (50, 5)$, $(40, 5)$, and $(30, 5)$, the yields were set to zero, because the shape of the spectra did not allow for a proper background correction.

efficiencies are obtained in the following manner. First, photons are reconstructed using the single-pass distributed cluster-finding algorithm, and their energies are stored in a histogram. The histogram is fitted with a double-exponential function, and the area under the peak is integrated from 3σ to the left of the mean to 2σ to the right, as the shape is asymmetric. The standard deviation, σ , was extracted from the Full-Width-at-Half-Maximum (FWHM) from the fit. The resulting yield is divided by the total number of events generated to obtain the efficiency. The final results are shown in the left part of Figure 6.11 for the event-based and the time-based scenario. The efficiency overall lies around 80%, with the efficiency of the event-based simulation slightly lower, but well within error bars. The resolution, σ , is shown in the right side of Figure 6.11. As desired, all values fall below the design resolution, indicated by the green dashed line, except for the 2 GeV case, which lies slightly above the design value. At other interaction rates, the single-photon efficiency for 1 GeV photons was determined to be 76.5%, 76.5%, and 74.6%, for rates of 2 kHz, 2 MHz, and 20 MHz, respectively, using the spectra in Figure 6.12. All numbers have an error of about 1.6%. The probability for pile-up was determined using the Poisson distribution, $P[k; (r \cdot l)] = \frac{(r \cdot l)^k e^{-(r \cdot l)}}{k!}$ with $k = 1$ (one more waveform within the time of the

Figure 6.11. Single-photon efficiency (left) and resolution (right) at different photon energies. The error bars are statistical only. The green dashed line in the right figure indicates the desired resolution (section 3.3.2.5). The right figure shows only the results of the time-based simulation.

first waveform), r the average hit-rate in a single crystal, and $l = 700$ ns, the length of the waveform. It was found to be less than 0.01% for rates of 2 kHz, 200 kHz, and 2 MHz, respectively, and 1.1% at a rate of 20 MHz. The occurrence of pile-up can, hence, account mostly for the drop at 20 MHz. It is unclear why the efficiency at 200 kHz is higher than at the other rates. The peaking structure in the time distribution of the 2 kHz scenario (Figure 6.5 (left)) may be connected to its apparent tendency to shift the cluster energies towards lower energies, as an ill-defined time difference may cause hits to be excluded from a cluster. Correcting the rounding errors problem should improve the results at this interaction rate.

For a more advanced test, the four cluster-finding algorithms described in section 6.4.1 are pitted against each other for the two decay channels and five interaction rates mentioned above (event-based, 2 kHz, 200 kHz, 2 MHz, and 20 MHz). Because speed is key, also the processing time is compared, next to the yield.

Comparison of Efficiency | Naively, one would expect the efficiencies for the $p\bar{p} \rightarrow \gamma\gamma$ and $h_c \rightarrow \dots \rightarrow 7\gamma$ channels to be around $\{58.5\%, 64.3\%, 58.5\%, 55.6\%\}$ (ϵ_γ^2) and $\{15.3\%, 21.3\%, 15.3\%, 12.9\%\}$ (ϵ_γ^7), respectively, starting from the single-photon reconstruction efficiency, ϵ_γ , for rates of $\{2 \text{ kHz}, 200 \text{ kHz}, 2 \text{ MHz}, 20 \text{ MHz}\}$. This assumes that the single-photon efficiency is independent on the photon energy, and that the detected photons are uncorrelated. The results of the simulation are shown in Figure 6.13 and Figure 6.14 for the two decay channels. The cluster-finding algorithms are referred to by ‘DEF’ for the currently implemented, default algorithm, ‘ONL’ for the online algorithm, ‘DIST’ for the single-pass distributed method, and ‘2STEP’ for the double-pass distributed

Figure 6.12. Cluster energy spectra for 10,000 events of single 1 GeV photons at different interaction rates. A minimum cluster energy of 30 MeV was required.

algorithm. The yields are given as a percentage of the original 5000 events that were generated. The yields in the $p\bar{p} \rightarrow \gamma\gamma$ channel are higher than the expectation from ϵ_γ . This can be attributed to the correlation between the two photons: to conserve energy and momentum, the angle between the two photons is correlated: they are emitted back-to-back in the centre-of-mass frame. Both photons will be detected by the EMC, effectively leading to two single photons, rather than a single double-photon event, to be detected. This makes the efficiency approach that of single-photons. At multiplicity 7, the correlations appear to be small, although correlations due to conservation of energy, momentum, and mass do exist. In this case, the yields seem to conform more to the naive prediction. The small deviations are likely due to an efficiency loss in the reconstruction algorithm. The 20 MHz rate forms an exception, where yields are compromised by pile-up and event-mixing effects. Especially for that rate, much may be gained when pile-up recovery is re-enabled.

Apart from this global trend, there is little difference in the yield of the four methods for the $p\bar{p} \rightarrow \gamma\gamma$ channel. Most differences are smaller than the statistical error, which is about 1.4%. The largest deviation between the online algorithms, ONL, DIST, and 2STEP, occurs at the 20 MHz scenario, and is 3.2%, or 2.3σ . At the same rate, also the largest deviation from the default method, DEF, occurs, which lies 2.2%, or 1.6σ , below the worst-performing online algorithm, DIST. For the $h_c \rightarrow \dots \rightarrow 7\gamma$

channel, the DEF and ONL methods yield similar results up to a rate of 2 MHz, and perform about 5%, or 3.6σ , better than the distributed algorithms, DIST and 2STEP. At the highest rate, ONL performs about 4%, or 2.9σ , better than DIST and 2STEP, and the yield of DEF drops significantly. This is because DEF does not take the time between digis into account when building clusters. Restoring pile-up recovery is expected to improve the efficiency of the online algorithms. In particular, it would be interesting to see how the difference between the algorithms is affected by .

Speed Comparison | For the processing time, all times are taken relative to the DEF method. This procedure enables a more independent comparison, as the absolute processing time needed depends on the hardware configuration of the device used to run the simulation. The quoted times are measured for the algorithms only, assuming they have already been given the list of digis to process. The single- and double-pass versions of the distributed cluster-finding algorithm (DIST and 2STEP) each have two numbers: the average time it took to form preclusters per participating (virtual) DC, obtained by dividing the total time needed to form preclusters by the number of participating (virtual) DCs, and the time needed to merge them into the final clusters. Times were measured using ROOT's TStopwatch function, and the obtained CPU time was used in the results.

As can be seen in Figures 6.13 and 6.14, the distributed cluster-finding algorithms, DIST and 2STEP, are considerably faster in most cases. This holds most notably for the double-pass version, 2STEP, which continues to emerge victorious when speed is concerned, as desired. They still exhibit a decent efficiency, close to the naive expectation, based on the single-photon efficiency, and outperform the DEF algorithm at the 20 MHz interaction rate. As stated before, the drop in efficiency at this rate for the DEF algorithm can be attributed to the failure to take the time into consideration when constructing clusters. Although the online algorithms already start to suffer at lower rates when the time of the digis is neglected, the critical need for accurate timestamping manifests itself most strongly in the high-rate case, as can be seen in Figure 6.15. The online cluster-finding algorithm, ONL, generally performs the same as the default one, and better at higher rates, however, then the processing time explodes. Investigation revealed that this is because at the chosen optimal values for $\Delta\tau$ and Δt , the time-bunches become very large. Specifically, the average number of time-bunches created per event, based on the number of events generated and the number of time-bunches formed from the digi data stream, are 1.3, 0.85, 0.47, and

Figure 6.13. Performance test results for $p\bar{p} \rightarrow \gamma\gamma$, at five interaction rates, for the four cluster-finding methods: DEF = default method, ONL = online method, DIST = single-pass distributed cluster-finding method, 2STEP = double-pass distributed cluster-finding method (see section 6.4.1). Quoted yields are the percentages of successful reconstructions relative to the number of events generated. ‘Successful’ in this context means a two-photon combination that ends up in the peak around $\sqrt{s} = 2.251 \text{ GeV}/c^2$ ($\pm 3\sigma$) in the invariant-mass spectrum. Processing times are taken relative to DEF.

Figure 6.14. Same as Figure 6.13, except for the channel $h_c \rightarrow \dots \rightarrow 7\gamma$. The number of h_c mesons was obtained by using the reconstruction method described in section 5.4.

Figure 6.15. Relative drop in yield, when time is not taken into consideration by the single-pass distributed cluster-finding algorithm, as a function of the interaction rate.

0.02 at interaction rates of 2 kHz, 200 kHz, 2 MHz, and 20 MHz, respectively. It must be noted that, after clusters are formed and the flight-time correction is applied, the number of time-bunches comes closer to the number of events generated. Hence, the event reconstruction algorithm does not have to deal with such high multiplicities. As the neighbour-finding part of the online algorithms grows with the number of digis that need to be processed, so does the processing time. This also explains the large contribution from the precluster merging section of the single-pass distributed cluster-finding algorithm. The double-pass case is less influenced, because it creates new time-bunches after preclusters have been formed and, at the forming stage, it runs the computations in parallel on subsets of the data.

In summary, the newly developed online cluster-finding algorithms perform well, and in particular, the double-pass distributed cluster-finding algorithm, which fits the design of the readout hardware, performs especially well in terms of speed. There are some important points to take into consideration regarding these conclusions:

- These results were obtained by performing the (more advanced) offline analysis directly after the cluster-finding step. In reality, there will be an event selection step in between these two steps, which will also effect the final efficiency.
- The time stamp generation of the digitised hits is not necessarily very reliable at this stage, and the very low rate case suffers from rounding errors.
- Pile-up recovery was disabled in obtaining these results.

In order to arrive at an unambiguous conclusion on the performance of the cluster-finding algorithms, these notes must be taken into account. However, all algorithms

Circumstances	Ideal situation	Including background	Including full detector	Including full detector and background
Relative yield (%)	20.7	12.6	6.2	5.8

Table 6.4. Relative yield of signal events for the channel $h_c \rightarrow \dots \rightarrow 7\gamma$ at an interaction rate of 200 kHz. The effect of including different, more realistic circumstances is showcased. Third column: Including 10,000 background events, created with the DPM generator at a beam momentum of 5.61 GeV/c (same as needed to produce the h_c resonance). Fourth column: Including the full detector setup, with the exception of the SciTil detector (section 3), as there were problems with the geometry file. Last column: Including both situations.

should more or less be equally affected by this. Therefore, the relative performance, and, hence, the conclusions drawn here, are not expected to change dramatically.

6.6 Performance Under More Challenging Circumstances

So far, all tests were performed under special circumstances, i.e. only events of the channel under investigation (‘signal events’) were generated, and all detector subsystems that are nested inside the EMC were not simulated, giving the produced particles a free path to the EMC crystals. A natural follow-up question is, how the algorithms perform under less ideal circumstances, e.g. by including unrelated background contributions, using the DPM generator, to the ‘signal’ events and materials of nested subsystems. In the simulation, one ‘signal’ event and one DPM background event are created simultaneously when that background type is enabled. This creates the (unrealistic) scenario where a signal event and a background event are maximally mixed, thereby putting the cluster-finding algorithm and (more so) the following event reconstruction algorithms to test. The results are shown in Table 6.4 for the single-pass distributed cluster-finding algorithm (which was also used for benchmarking before). Results are obtained in the timebased framework, running at an interaction rate of 200 kHz.

As expected, the yields drop when more realistic circumstances are introduced. The drop in efficiency, when adding background contributions from the DPM generator, is a consequence of the large cluster multiplicity in that background, which can be seen in Figure 6.16. However, the largest effect comes from including the other detector subsystems. This can likely be attributed to scattering effects and

Figure 6.16. Cluster multiplicities versus cluster energies, for 10,000 events generated with the DPM model at a beam momentum, p , of 15 GeV/c (top panel), 5,000 DPM generated events at $p = 5.61$ GeV/c (middle panel), and 5,000 events of the $h_c \rightarrow \dots \rightarrow 7\gamma$ channel at $p = 5.61$ GeV/c (bottom panel). The average multiplicity (‘Mean x’ in the figure) drops with p , as expected. The amount of pile-up at an interaction rate of 200 kHz, and with $\Delta\tau$ set to 257 ns, is in total about 5%, calculated using the Poisson distribution. This agrees with the result shown in the bottom panel, where the average cluster multiplicity should be seven; multiples thereof, hence, correspond to mixed events.

Figure 6.17. Cluster energies, obtained by the single-pass distributed cluster-finding algorithm at an interaction rate of 200 kHz. The results from three simulations are shown: one with both signal and background events ($h_c + \text{DPM}$), where results from a simulation with only background events (DPM) have been subtracted (blue line), and one with only signal events (h_c only', green line).

other interactions that take place in the material of the other subsystems. A future investigation may reveal which systems cause the largest effect. It is important to note that the drop in yield is not necessarily due to the cluster-finding algorithm, and more likely due to the event reconstruction algorithm.

To see how the algorithm itself is affected when e.g. background is included, the cluster energies in the situation where only background was simulated (DPM) has been subtracted from the case where both the signal events and background events ($h_c + \text{DPM}$) are included. The result is compared to the situation where only signal events (h_c only') were simulated in Figure 6.17. The colourful shades around the histograms indicate a 3σ statistical-error band. At most places in the Figure, the bands overlap, indicating that there is no statistically significant difference between the two cases. However, in particular at lower energies, deviations do occur. This can likely be attributed to the fact that for each data set, the digitisation task was rerun, which introduced a small randomness to the digis that were found. This fluctuation will most severely impact the lower-energetic clusters, that are inherently more susceptible to the influence of noise. The cluster-finding algorithms' ability to properly identify clusters is not affected within statistical uncertainty by the inclusion of non-ideal factors, judging from Figure 6.17.

Figure 6.18. Schematic overview of the setup that was used to test the implementation of the double-pass cluster-finding algorithm on a prototype of the readout hardware. The Compute Node prototype has a Xilinx Virtex 5 FX70 FPGA with 11,200 slices and 5,328 kB RAM. The test board has a Xilinx Kintex 7 XC7K325 FPGA with 50,950 slices and 16,020 kB RAM. Further details are described in the text.

6.7 Performance Assessment by Hardware Simulation

6.7.1 Performance of the Algorithm on a Hardware Prototype

To support the claim that the double-pass distributed cluster-finding algorithm will run well on the readout system, its implementation on a prototype of the readout hardware is currently ongoing. Important properties to be checked are the time and resources the algorithm will need, to see if sufficient resources are available on the FPGAs in the Data Concentrators (DCs). The test is carried out at KVI-CART, Groningen, the Netherlands, where the SADCs are also debugged, and the other algorithms mentioned in section 4.3.1 and 4.3.2 are developed and implemented. The algorithm has been implemented on an FPGA in the VHDL programming language using Vivado v2017.1 for the test board, running on a Kintex 7 XC7K325 FPGA, and ISE v14.7 for the Compute Node, based on a Virtex 5 FX70 FPGA. The setup that was used to test the output is shown schematically in Figure 6.18. First, digitised hit (digi) data were simulated for the $h_c \rightarrow \dots \rightarrow 7\gamma$ decay channel using PandaROOT. The simulated, digitised data were reformatted to mimic the data coming out of the SADCs. These data were mapped onto virtual Data Concentrators (as described in section 6.4.1), and digi data from four adjacent DCs were written in the correct data format (see section 4.3.3) using a LabView program. Using the TRB protocol [83], the data were sent via a Gb network and a TRB3 board to a Xilinx test board with a Kintex 7 FPGA over a 2 Gb optical link. There, the data were stored in a memory block on the FPGA, before the command was given to send all the data to a prototype

Figure 6.19. Comparison between the output of the double-pass distributed cluster-finding algorithm run on an FPGA in VHDL, and on a CPU in PandaROOT. The small difference corresponds exactly to the number of times some data was not stored because the number of objects exceeded the buffer size. This difference occurs mainly in the bin with the highest number of counts.

of the Compute Node, which also used 2 Gb optical links. At its input, the CN first ran the first step of the double-pass algorithm to identify preclusters in the virtual DCs. After this, the second step of the double-pass algorithm was run to merge preclusters if needed. The final cluster data, which consisted of a header, specifying each cluster's position, timestamp, energy, and the number of digis, followed by the participating digis, were sent over a Gb ethernet connection back to a regular desktop computer. This data format is foreseen to be used in the final design of the detector setup as well. A new version of the Compute Node, with upgraded hardware (Kintex 7 FPGAs and new high-speed links), is currently being designed at IHEP, Beijng, China.

The result from the VHDL implementation is shown in Figure 6.19 for the mapped X-coordinate, comparing it to the result of a PandaROOT simulation. The figure demonstrates that there is practically no difference between the VHDL implementation and the PandaROOT simulation. The small difference (about 400 counts) corresponds exactly to the number of times some data were not stored because the number of objects exceeded the buffer size. This has to be corrected by obtaining reliable estimates for the number of digis per time-bunch, such that the size of the buffer can be adjusted to minimize this effect. This allows to fix the buffer size. The result suggests that the output of the hardware prototype can be trusted, and its

Figure 6.20. Block diagram of the FPGAs in the L1 nodes. Each FPGA has four modules, and each module has each four chains. One of the chains is shown. Each chain processes a different part of the data from the *same eight inputs*, based on their superburst number, but has its *own output*. Hence, the FPGAs use 32 inputs and 16 outputs. More details can be found in the text.

performance can be optimised using simulated data. In addition, the DCs appear to have sufficient resources to run the precluster forming algorithm.

6.7.2 Simulation of the Burst Building Network

Next, a virtual version of the Burst Building Network, discussed in section 4.3.3, was simulated on a desktop computer, to check if the time needed to process the data is sufficient to run the precluster merging algorithm. In the simulation, there are 64 DCs each with 16 inputs and 2 outputs, 8 L1 Burst Building Network nodes each with 4 modules, where each module has 8 inputs and 4 outputs. There are 8 L2 Burst Building Network nodes each with 2 modules, every module having 8 inputs and 8 outputs. All modules are fitted with a Xilinx Kintex Ultrascale (XCKU115-1FLVB1760C-ND) FPGA with a clock speed of 150 MHz. A schematic overview of the processing chain in an L1 node is depicted in Figure 6.20. The node's FPGA is divided into four modules, which can operate independently of each other. Each module has four of the chains depicted in Figure 6.20, which draw different data from the same eight inputs of the module, based on their superburst⁴ number. The data are, then, chronologically combined to a single data stream. The Gap Finder divides the new precluster data stream into bunches and distributes these bunches over the available IP

⁴Recall that a superburst is 16 bursts, and each burst is the time that the beam and the target overlap, plus the time gap between the overlap time, totalling to 38.4 μ s.

cores⁵, which perform the merging of preclusters. Eight of these were used; however, this number can be changed depending on the available resources. The output of the IP cores is buffered, and then collected. The chain follows a push architecture, meaning that the next data set cannot be processed until the current one has finished. The L2 nodes follow the same architecture, albeit with a different number of modules.

Due to limited resources on the desktop computer, only two out of four outputs of the Data Concentrators were simulated, meaning $(8 \text{ L2 nodes}) \times (2 \text{ modules}) \times (4 \text{ outputs}) = (64 \text{ CNs})$ were simulated. Data are collected in the L2 nodes, based on their superburst number, and each data packet with a given superburst number is sent sequentially to one of the 128 Compute Nodes. Data with superburst number 1 are sent to the first CN, data with superburst number 2 to the second, and so on. Hence, the time available to a CN to process the data, until the next superburst data packet is offered, is $128 \times 38.4 \mu\text{s} = 4.92 \text{ ms}$. In the current simulation, half of the superbursts were not simulated, and the fact that only half of the network was used in the simulation, therefore, does not influence this time estimate.

The test data suite, generated with PandaROOT running the DPM generator, leads to an estimate of a total data rate of 137.56 Gbps (see section 4.3.3 for a breakdown of the calculation). The data are processed in parallel on the L1 nodes, and then pushed to the L2 nodes. Before the L2 nodes can set to work, they need to wait until they have received all of the data from the same superburst from the L1 nodes. Hence, the processing speed of the network is limited by the ‘slowest’ L1 node, i.e. the one that takes the longest to process the data it is being fed. According to the simulation, the longest time that is needed to process the data is 5.32 ms. As the available time is 4.92 ms, the speed of the algorithm is insufficient to be run on the network. However, each Data Concentrator received output from the same number of SADCs, although the forward-boosted particle-production causes the data rates in SADCs in the forward direction to lie substantially higher than that of their counterparts in the backward direction. Distributing the data over the DCs, by taking this asymmetry into account, will likely improve the throughput of the network to the extent that it is able to process the data stream. In addition, the fact that the 20 MHz interaction rate will not be attained in the initial stages of the experiment further aspires confidence that the envisaged network will be able to handle the data stream. In later stages, it is

⁵An Intellectual Property (IP) core is a reusable unit of logic blocks with a specific (often patented) function.

foreseen that new developments in hardware will be able to handle the data produced in any future upgrades of the setup.

Concluding, the double-pass distributed cluster-finding algorithm shows a good performance, while consuming few resources. It can, and should, therefore, run on the currently envisaged devices for the readout system.

Chapter 7

Outlook and Conclusion

The future $\bar{P}ANDA$ experiment will be searching for rare decays. To find the needle in the proverbial haystack, an extensive readout concept has been devised, where each detector subsystem prepares the data for filtering. For the EMC, this amounts to identifying the correct clusters of hits, to build the photon and electron candidates (among others). As the data will be produced at a very high rate, the algorithms that perform this task need to work efficiently at high speed. In this Chapter, the results and points of improvement on the algorithms that have been developed, which were discussed in the previous chapters, will be summarised. Suggestions on how to proceed and how to use the obtained results will be coined.

7.1 Outlook

As emphasized throughout this work, time information is crucial. However, in the course of the investigation, several problems with the time stamp generation and pile-up recovery procedure were found. In the future, these problems will need to be fixed in order to get a more reliable estimate of the performance of the algorithms. However, these problems are expected to affect all of the tested algorithms equally, and should, therefore, not affect the conclusions, but the optima that were determined for the time cuts may change. Another point that requires refinement is the mapping of the crystals in PandaROOT. In particular, the boundaries of the sections of the EMC are not (properly) connected, leading to a loss of performance when using the algorithms with the mapped coordinates. In addition, the results were obtained by performing the (more advanced) offline analysis directly after the cluster-finding step. In reality, there will be an event selection step in between these two steps, which will also effect the final efficiency. To obtain the final efficiencies of the algorithms, these points must be taken into consideration.

The inclusion of other subsystems in the simulation leads to a noticeable decrease in the event reconstruction performance, but does not critically affect the cluster-finding algorithms themselves. It may be worth investigating which of the subsystems contributes the most to this drop in efficiency, so that a correction may be applied. For example, the detection of bremsstrahlung by electrons allows an energy correction to the reconstructed electron. Also, electrons that are knocked out from the detector material (also known as ‘delta electrons’) by energetic charged particles may be identified by a track reconstruction, depending on where they were created. These corrections require data from other detector subsystems and may, hence, only be applied near the end of the readout chain. At the earliest, this can be done in the Burst Building Network, depending on the implementation that is chosen. In the topology described in section 4.3.3, the Compute Nodes will be the first to have access to sufficient data to perform such corrections. The cluster-finding algorithm can also be used in combination with the tracking system to identify photon candidates in a very early stage. This allows to reconstruct neutral pions online, and, as their mass is well-known, they can be used to calibrate the energy while the detector is running. It is strongly recommended to implement this procedure, as the use of uncalibrated energies may cause the event reconstruction to fail.

7.2 Conclusion

The $\bar{P}ANDA$ experiment, currently under development, has a great discovery potential in several regimes of QCD. The rare decay channels that are searched for feature a final-state topology that has a very rich background, which excludes a conventional hardware trigger to suppress this background. Therefore, a novel data acquisition method has been developed, in which events are reconstructed in real-time. To make this approach feasible, algorithms must run on the readout hardware of the detector subsystems to reconstruct information in real-time, such as charged-particle tracks and photon and electron energy depositions. Online algorithms to perform the latter task, so-called cluster-finding algorithms, have been developed and tested.

To allow a high through-put, these algorithms need to be fast and consume only little resources. At the foreseen high interaction rate in the experiment, pile-up of signals and event mixing complicate the task of identifying clusters and assigning them to the proper event. Furthermore, these features lead to the absence of a clearly-separated event structure in the data stream, which is why an approach is used in which the stream is cut into bunches. The parameter used to realise this division in the example of the EMC has been optimised. It naturally depends on the interaction rate, and optima of $\Delta\tau = 257, 97$, and 15 ns were determined for rates of 200 kHz or less, 2 MHz, and 20 MHz, respectively. As time information is relevant, this information must also be used to assign hits to clusters, and the optima for this parameter were found to be $\Delta t = 282, 122$, and 40 ns at the aforementioned rates, following the global trend $\Delta t = \Delta\tau + 25$ ns. Once $\Delta\tau$ has been determined, consecutive digis will have a time difference smaller than $\Delta\tau$. However, digis belonging to the same cluster need not be consecutive in time in the data stream, as all digis from all clusters are thrown on one big pile. Hence, a larger value for Δt is needed to collect all the digis belonging to the same cluster. The value found, 25 ns higher than $\Delta\tau$, is related to the time resolution, which drops drastically for the lower energies of the digis on the periphery of the cluster (Figure 3.4). The time difference with the central crystal extends to over 30 ns, see Figure 7.1.

Several algorithms were pitted against each other to evaluate their performance. The distributed cluster-finding algorithms, which fit the design of the readout system, were found to perform the best in terms of speed, while only performing a few percent less in terms of yield than the online algorithm, which showed the highest efficiency. Furthermore, investigations if these algorithms can run on the envisaged

Figure 7.1. Absolute time difference between a digi in a cluster and the most central digi in the cluster.

readout hardware turned out promising. Although the online algorithm performed the best at an interaction rate of 20 MHz in terms of yield, the fact that it takes about 800 times longer to obtain this result compared to the double-pass algorithm, in combination with the fact that the double-pass algorithm already pushes the available resources in the network to their limit, would render the system less usable. Hence, it is recommended to use the double-pass distributed cluster-finding algorithm in the final design of the detector.

A summary of the performance of the double-pass algorithm is shown in Table 7.1. The single-photon efficiency lies around 80%, and drops for higher interaction rates. This drop is much higher than the amount of pile-up that is expected at 20 MHz, which is 1%. This estimate was obtained from the Poisson distribution, $P[k; (r \cdot l)] = \frac{(r \cdot l)^k e^{-(r \cdot l)}}{k!}$ with $k = 1$ (one more waveform within the time of the first waveform), $r = 15.242$ kHz (the average hit-rate in a single crystal in this configuration), and $l = 700$ ns, the length of the waveform. The actual amount of pile-up in the simulation, was found to be $\sim 3\%$ at 20 MHz using Figure 6.12. At a multiplicity of 2, the extracted efficiency is higher than what would be naively expected from the single-photon efficiency. This is again more than what is expected from pile-up and event mixing effects, which contribute 1.5% at these rates. The higher yields can be attributed to the correlation between the two photons. This makes the efficiency approach that of single-photons. At multiplicity 7, the correlations appear to be small, although correlations due to conservation of energy, momentum, and mass do exist. In this case, the yields seem to conform more to the naive prediction. The small deviations are likely due to an efficiency loss in the reconstruction algorithm. The 20 MHz interaction rate forms an exception, likely due to an underestimation of

Photon multiplicity	Interaction rate			
	2 kHz	200 kHz	2 MHz	20 MHz
1 ($E = 1$ GeV)	78.6	81.0	77.8	65.3
2 ($\sqrt{s} = 2.251$ GeV)	77.2 (61.8)	77.6 (65.6)	74.7 (60.5)	68.9 (42.6)
7 ($h_c \rightarrow \dots \rightarrow 7\gamma$)	18.2 (18.5)	18.8 (22.9)	16.5 (17.3)	8.5 (5.1)
$\sigma_E(1$ GeV)	2.53	2.38	2.07	2.18

Table 7.1. Overview of the efficiencies (in %) of the double-pass distributed cluster-finding algorithm at different $p\bar{p}$ interaction rates and at different photon multiplicities. The ‘expected’ efficiency, based in the results for multiplicity 1, is shown in brackets. The efficiencies shown have an error of around 2%. Numbers for multiplicity 2 and 7 are taken from Figures 6.13 and 6.14. The last row shows the resolution (σ_E , in %) for single photons with an energy of 1 GeV. Results for multiplicity 1 are obtained from the first paragraph of section 6.5.2.

the single-photon efficiency. The drop is more pronounced at higher rates, indicating pile-up effects, which is about 9% at 20 MHz. In this case, pile-up correction may improve the efficiencies at a higher rate significantly. Table 7.1 can be used to determine whether it is worth to pursue decay modes with a certain photon (or electron) multiplicity at the highest interaction rate. However, the correct implementation of a pile-up recovery algorithm (into the PandaROOT code) may radically increase the efficiencies at the highest rate, and thereby affect this conclusion. It is, therefore, vital that pile-up recovery is properly implemented for higher interaction rates, both in the simulation and in the readout hardware.

Following the study if the algorithms can run on the readout system, a concept for the data collection network was set up using estimates of the data rate. This includes the number of nodes, connections, and layers needed in this network. The design of this network is useful in determining the final requirements for the complete data acquisition system, and can serve as a guideline in the design of the network for the other detector subsystems. The foreseen network design will have FPGAs, and can, therefore, perform advanced processing tasks, in addition to collecting data of the complete detector subsystem, as long as the bandwidth and resources permit this. Two topologies were discussed. In section 4.3.3, the number of devices needed was based on an estimate for the data rates and the number of devices that need to be connected to each other. The implementation in section 6.7.2 took an initial guess of the number of devices needed based on the available resources. In future calculations, it is probably best to let the first solution guide the second.

The analysis on the ability of the cluster-finding algorithm to run on the readout system was based on the worst-case scenario of a 20 MHz interaction rate. As this rate will not be attained in the first years of operation, the amount of data that the readout system has to process will be lower than this estimate. This implies that the speed of the algorithm will certainly be high enough to run online on the data stream. In the future, advances in hardware development are expected to allow the algorithms to run at the higher interaction rates as well. In conclusion, it is recommended to proceed with the implementation and testing of the double-pass distributed cluster-finding algorithm, considering its performance and the fact that its design fits well in the concept of the triggerless readout system.

Summary

After Antonie van Leeuwenhoek first observed bacteria using his microscope in 1673, a quest was started to find the smallest building blocks of nature. Molecules and atoms were discovered, and the atoms themselves turned out to contain a compact core of protons and neutrons. The arrival of ever more powerful particle accelerators opened up the possibility to investigate nature on an even smaller scale. The protons and neutrons turned out to have been built up from even smaller particles, the quarks. The model that describes the interactions between these quarks and the forces they experience, the Standard Model, depicted in Figure 7.2, could also be used to describe the other composite particles that had been observed.

Figure 7.2. Schematic overview of the fundamental building blocks of matter. Colour-shaded areas indicate which forces act on which set of particles by the ‘shadow’ casted from the base of the participating blocks. Differently shaded blocks show the family they are part of. The red blocks show the carriers of the different forces: the photon for the electromagnetic force, the W and Z boson for the weak nuclear force, and the gluon for the strong nuclear force.

However, the Standard Model cannot account for all facets of the observed particles. The quarks and leptons obtain their mass by interacting with the Higgs field. The story changes for composite particles, like the proton. The mass of the whole is much more than the sum of its parts. This can be attributed to the interaction between the quarks and the gluons, and the gluons with each other, within the proton. However, it remains unclear how this process works. According to the Standard Model, particles that cannot be categorised as conventional hadrons, i.e. hadrons consisting of two or three quarks, should have already been observed in abundance, which is not the case. These so-called ‘exotic particles’ have only been observed recently, and their structure still remains a mystery. These phenomena show that nature has not yet been fully understood at the smallest scale (femtometer, 10^{-15} m).

Several experiments are working to unravel these mysteries. They set their sights on the energy range where the charmonium meson, consisting of a charm quark and its antiquark, can be formed. The relatively high mass of the charm quark makes it possible to connect theory and experiment, using non-relativistic perturbation theory. However, current experiments are either limited to producing particles with a specific set of quantum numbers, or cannot directly produce all particles of interest. In order to study rare processes in detail, the $\bar{P}ANDA$ (antiProton ANnihilations at DArmstadt) experiment is being developed. Here, an energetic antiproton beam collides on a stationary proton target (consisting of frozen hydrogen droplets) at momenta between 1.5 GeV/c and 15 GeV/c. The antiproton beam is produced by the main accelerator at FAIR, the Facility for Antiproton and Ion Research, as a secondary beam. Because the (anti)proton is a composite particle, all (from an energy balance point-of-view accessibly) particles, in particular the exotic particles, can be directly and copiously produced.

Like all detector systems of its kind, the $\bar{P}ANDA$ detector is a layered system. Tracking systems are nested in time-of-flight systems and the electromagnetic calorimeter (EMC). The whole is situated within an electromagnet, forcing charged particles to follow curved trajectories, so they can be identified. To measure the rarest processes, a high interaction rate will be used, in which 20 million interactions per second take place (interaction rate of 20 MHz). Because the rare processes leave a similar footprint in the detector as background signals, the use of a conventional hardware trigger will not be sufficient to reduce the data stream, about 2,400 Gbps, to a manageable size. For this reason, a new approach was created, in which the

interactions that lead to the formation of new particles, and their decays, so-called ‘events’, will be reconstructed online by the readout system. In this way, potentially interesting events can be distinguished from background events, so they can be stored for later analysis. This should reduce the amount of data by a factor 1,000. To enable this concept, the subsystems of the detector must prepare the data. In the case of the EMC, which is exploited here as an example, the particles that deposit their energy in the material of the detector must first be reconstructed, before the readout system can reconstruct the complete event. To achieve this, a so-called cluster-finding algorithm has been developed. This algorithm will run in an early stage of the readout system to identify the correct clusters of energy depositions in the EMC crystals (so-called ‘hits’). The high interaction rate leads to the pile-up of signals and the mixing of events. This complicates the task of finding clusters considerably. Besides this, the data stream no longer has a clear and unambiguous structure in time. Therefore, the data stream must be divided to the best of our knowledge in so-called time-bunches using the temporal distance between two consecutive hits.

A well-functioning cluster-finding algorithm already existed, but this was not designed for online usage; it was not designed to operate as quickly as possible and to use as little resources (computation power, memory) as possible. In this algorithm, each hit starts out as a cluster, and it is checked if each following hit should become part of the cluster that is currently being constructed. However, checking for each hit if it is neighbouring to an existing cluster leads to a large number of loops in the algorithm. For this reason, this algorithm is not suited for online usage, and so-called ‘online’ algorithms were developed. All the online algorithms follow the same principle: the list of hits is looped over to check which ones are neighbouring in space and time. Then, this list is used to assign hits to the correct cluster. The online algorithm loops over the entire data set to build clusters using this principle.

Because the EMC is made up of a large quantity (circa 15,552) lead tungstate crystals (PbWO_4 , about $2 \times 2 \times 20 \text{ cm}^3$, tapered), it needs to be read out in sections. Each Data Concentrator (DC) reads out one of these sections, consisting of about 128 crystals. To reduce the load on the readout system, each DC can already start looking for clusters in the section it is reading out. These clusters are called ‘preclusters’, because the ones that are on the edge of a section may need to be merged later on. This approach is called ‘distributed cluster-finding’. Two implementations have been investigated: single- and double-pass versions. The difference is that in de

single-pass version, first, all data must be collected, so it can be used in parallel. In the double-pass version, first preclusters are formed, and later new time-bunches are built from the precluster data stream, after which preclusters are merged, if needed. This way, the data can be processed in parallel from the first DC, without the need to first collect the complete data set. The double-pass version, therefore, fits the design of the readout system, and is the version of choice.

As time plays an important role, because of complications caused by the high interaction rate, the optimal settings were investigated. Two parameters are important: the time used to divide the data stream into bunches, $\Delta\tau$, and the time used to assign hits to clusters, Δt . As expected, the optima depend on the interaction rate. To determine the optima, two decay channels with different cluster multiplicities have been used: $p\bar{p} \rightarrow \gamma\gamma$ (two photons in the final state) and $p\bar{p} \rightarrow h_c \rightarrow \gamma\eta_c \rightarrow \gamma\pi^0\pi^0\eta \rightarrow 7\gamma$ (with seven photons in the final state). The latter required a more extensive reconstruction, and could be used to determine the influence of geometrical overlap because of its cluster multiplicity. The single-pass version of the distributed cluster-finding algorithm was used for testing on all instances. The optima were found to be $\Delta\tau = 257, 97, \text{ and } 15 \text{ ns}$ and $\Delta t = 282, 122, \text{ and } 40 \text{ ns}$ for interaction rates of 200 kHz and less, 2 MHz, and 20 MHz, respectively. The optimal neighbour distance in space was found to be one crystal, and a threshold of 3 MeV was imposed on the hits.

Using the aforementioned optima, the performance of the different cluster-finding algorithms could be compared. First, the efficiency for the detection of single photons was determined. The results are shown in Figure 6.11 for four energy values. Assuming that the efficiency is independent of the energy, and that the photons are uncorrelated, the expectation for the two- and seven-photon channels are $\{61.8\%, 65.6\%, 60.5\%, 42.6\%\}$ (ϵ_γ^2) and $\{18.5\%, 22.9\%, 17.3\%, 5.1\%\}$ (ϵ_γ^7), respectively, with ϵ_γ the efficiency for the detection of single photons. The actual results are shown in Figure 6.13 and 6.14. The obtained values for the two-photon channel are higher than expected, because the two photons are correlated. In general, the efficiency fluctuates little, staying within the statistical error. Only at the highest interaction rate does a difference of 3.2% occur between the online algorithm and the distributed algorithms, and does the efficiency of the existing algorithm drop with an additional 2.2%. At the seven-photon channel, the efficiencies of the distributed cluster-finding algorithms are 5% lower than the others, but at the 20 MHz interaction

rate, the difference with the online version drops to 4%. However, the online algorithm required a substantial amount of time to reach this highest efficiency, which would make the system less usable. A test showed that the DCs are capable of running the precluster finding part of the distributed cluster-finding algorithms. For this reason, the use of the double-pass version of the distributed cluster-finding algorithm is recommended.

Finally, it was investigated if the data collection network is capable of running the precluster merging algorithm, next to collecting data. The simulation contained the number of nodes, connections, and layers that this network should have. The design of this network can be used to determine the final requirements for the complete data processing system, and can be used as a guideline when designing the network for other subsystems. Two topologies were investigated: an estimate for the number of devices needed based on the expected data rates, and an estimate based on the available resources on the computation blocks (Field-Programmable Gate Arrays, FPGAs). In summary, it is recommended to proceed with the development of the double-pass distributed cluster-finding algorithm, as it delivers a good performance at a low cost, and it can run on the foreseen hardware of the readout system.

Samenvatting

Na de eerste waarneming van bacteriën door Antonie van Leeuwenhoek in 1673 met zijn microscoop, begon een zoektocht naar de kleinste bouwstenen van de natuur. Moleculen en atomen werden ontdekt, en de atomen zelf bleken een compacte kern van protonen en neutronen te hebben. Met de komst van steeds krachtigere deeltjesversnellers, werd het mogelijk om de natuur op nog kleinere schaal te onderzoeken. De protonen en neutronen bleken zelf opgebouwd te zijn uit nog kleinere deeltjes, de quarks. Het model dat het samenspel beschrijft tussen deze quarks en de dragers van de krachten die zij ervaren, het Standaard Model, schematisch weergegeven in Figuur 7.3, kon tevens gebruikt worden om alle andere samengestelde deeltjes die werden waargenomen te beschrijven.

Figure 7.3. Schematisch overzicht van de fundamentele bouwstenen der materie. De gekleurde gebieden geven aan welke krachten op welke deeltjes werken middels de ‘schaduw’ vanaf de voet van de blokken. Verschillende kleurtonen geven de verschillende families aan. De rode blokken geven de dragers van de verschillende krachten weer: het foton voor de elektromagnetische kracht, het W en het Z boson voor de zwakke kernkracht, en het gluon voor de sterke kernkracht.

Echter kan het Standaard Model niet alle facetten van de waargenomen deeltjes verklaren. De quarks en leptonen krijgen hun massa door de wisselwerking met het Higgsveld. Voor samengestelde deeltjes, zoals het proton, is het verhaal anders. De massa van het geheel is veel hoger dan de som van de delen. Dit is te wijten aan de wisselwerking tussen de quarks en de gluonen, en de gluonen onderling, binnen het proton, maar hoe dit precies werkt is nog niet duidelijk. Volgens het Standaard Model zouden deeltjes die niet ondergebracht kunnen worden onder de noemer van conventionele hadronen, dat wil zeggen, hadronen bestaande uit twee of drie quarks, reeds vaker waargenomen moeten zijn, hetgeen niet het geval is. Enkele exemplaren van deze zogeheten ‘exotische deeltjes’ zijn pas recentelijk waargenomen, maar hun structuur is nog niet bepaald. Deze fenomenen laten zien dat de natuur op de kleinste schaal (femtometer, 10^{-15} m) nog niet helemaal begrepen is.

Er zijn meerdere experimenten die zich richten op het ontrafelen van deze mysteries. Deze onderzoeken hebben hun oog laten vallen op het energiegebied waarin het charmoniummeson, bestaande uit een charmquark met zijn antiquark, gevormd kan worden. De relatief hoge massa van de charmquark maakt een koppeling tussen theorie en experiment mogelijk, middels niet-relativistische storingsrekening. Echter, zijn de huidige experimenten ofwel beperkt tot het produceren van deeltjes met een specifieke set kwantumgetallen, ofwel kunnen ze niet rechtstreeks alle deze deeltjes produceren. Om de zeldzame processen in detail te kunnen bestuderen, is het $\bar{P}ANDA$ (antiProton ANnihilations at DArmstadt) experiment in het leven geroepen. Hierin wordt gebruik gemaakt van een energetische antiprotonenstraal, die op een stilstaand protonendoelwit (bestaande uit bevroren waterstofdruppels) botst met een impuls tussen de 1,5 GeV/c en de 15 GeV/c. De antiprotonenstraal wordt gemaakt als secundaire straal met behulp van de hoofdversneller van FAIR, de Facility for Antiproton and Ion Research. Doordat het (anti)proton een samengesteld deeltje is, kunnen alle (vanuit een energieoogpunt beschikbare) deeltjes, in het bijzonder de exotische deeltjes, rechtstreeks en in grote hoeveelheden geproduceerd worden.

De $\bar{P}ANDA$ detector is, net als vrijwel alle detectorsystemen van zijn soort, een gelaagd systeem. Traceersystemen zijn genesteld in vluchttijdmeetsystemen en de elektromagnetische calorimeter (EMC). Het geheel bevindt zich binnen een grote elektromagneet, zodat geladen deeltjes kromme banen zullen volgen en aan de hand daarvan geïdentificeerd kunnen worden. Om de meest zeldzame processen te kunnen meten, zal een hoge interactiesnelheid gebruikt worden, waarbij 20 miljoen

interacties per seconde (interactiesnelheid van 20 MHz) plaatsvinden. Omdat de zeldzame processen een soortgelijke vingerafdruk achterlaten in de detector als achtergrondsignalen, zal het gebruik van een conventionele hardware ‘trigger’ niet voldoende zijn om de datastroom, zo’n 2.400 Gbps, tot een behapbare grootte te verkleinen. Om deze reden is er gekozen voor een nieuwe aanpak, waarbij de interacties die leiden tot de formatie van (nieuwe) deeltjes en hun verval, zogeheten ‘gebeurtenissen’, online door het uitleessysteem gereconstrueerd worden. Op deze manier kan een onderscheid gemaakt worden tussen achtergrondgebeurtenissen en potentieel interessante gebeurtenissen, die vervolgens opgeslagen worden voor latere analyse. Dit zou de hoeveelheid data met een factor 1.000 moeten verkleinen. Om dit concept mogelijk te maken, moeten de subsystemen van de detector de data voorbereiden. In het geval van de EMC, moeten de deeltjes die hun energie in het detectormateriaal deponeren eerst gereconstrueerd worden, voordat het uitleessysteem de volledige gebeurtenis kan reconstrueren. Om dit te bereiken, is een zogeheten clusterzoekalgoritme ontwikkeld. Dit algoritme draait in een vroeg stadium van het uitleessysteem om de juiste clusters van energiedeponeringen in de EMC kristallen (zogeheten ‘treffers’) te identificeren. De hoge interactiesnelheid leidt tot ophoping van signalen en het door elkaar lopen van gebeurtenissen. Dit bemoeilijkt het proces om clusters te vinden aanzienlijk. Bovendien bevat de datastroom op het eerste gezicht geen duidelijke structuur meer in de tijd, waardoor deze op een intelligente manier opgedeeld moet worden in zogenoemde tijdsgroepen met behulp van de tijdsafstand tussen twee opeenvolgende treffers.

Er bestond reeds een goed functionerend clusterzoekalgoritme, maar deze was niet ontworpen voor online gebruik; het was niet ingericht om zo snel mogelijk en met zo weinig mogelijk bronnen (rekenkracht, geheugen) te werken. In dit algoritme begint elke treffer als een cluster, en wordt van alle volgende treffers bekeken of deze onderdeel moeten worden van het cluster dat momenteel opgebouwd wordt. Echter, is het nodig dat voor elke treffer wordt onderzocht of deze naburig is aan een bestaand cluster, hetgeen leidt tot grote hoeveelheden lussen in het algoritme. Om deze reden, is dit algoritme niet geschikt voor online gebruik, en zijn zogeheten ‘online’ algoritmes opgesteld. Al deze algoritmes volgen hetzelfde basisprincipe: er wordt geïtereerd over de lijst met treffers om te kijken welke naburig zijn in de ruimte en in de tijd. Hierna wordt deze lijst gebruikt om de treffers toe te kennen aan het juiste cluster. Het online algoritme itereert over de gehele dataset om volgens dit principe clusters te vormen.

Omdat de EMC is opgebouwd uit een grote hoeveelheid (circa 15.552) lood-wolframaat kristallen (PbWO_4 , ongeveer $2 \times 2 \times 20 \text{ cm}^3$, taps-lopend), moet deze in secties worden uitgelezen. Elke Data Concentrator (DC) leest één sectie uit, bestaande uit ongeveer 128 kristallen. Om de last op het uitleessysteem te verlagen, is het mogelijk om elke DC alvast naar clusters te laten zoeken binnen de sectie die uitgelezen wordt. Deze clusters worden ‘preclusters’ genoemd, omdat degenen die op de rand van een sectie liggen, later mogelijk samengevoegd moeten worden. Deze aanpak heet ‘gedistribueerd clusterzoeken’. Er zijn twee implementaties onderzocht: enkele- en dubbele-doorgeefversies. Het verschil is dat bij de enkele-doorgeefversie alle data eerst verzameld moeten worden, en deze vervolgens parallel verwerkt kunnen worden, terwijl bij de dubbele-doorgeefversie eerst preclusters gevormd worden, en later uit de preclusterdatastroom nieuwe tijdsgroepen worden gemaakt, waarna, indien nodig, preclusters samengevoegd worden. Hierdoor kunnen de data vanaf de eerste DC parallel verwerkt worden, zonder dat eerst de volledige dataset verzameld hoeft te worden. De dubbele-doorgeefversie past goed binnen het ontwerp van het uitleessysteem, en geniet daarom de voorkeur.

Omdat, door de complicaties die de hoge interactiesnelheid met zich meebrengt, de tijd een belangrijke rol speelt, is onderzocht wat de optimale instellingen zijn. Er zijn twee parameters belangrijk: de tijd om de datastroom op te delen in tijdsgroepen, $\Delta\tau$, en de tijd om treffers aan een cluster toe te kennen, Δt . Zoals verwacht, hangen de optima af van de interactiesnelheid. Om de optima te bepalen, zijn twee vervalkanalen met verschillende clusterveelvouden gebruikt: $p\bar{p} \rightarrow \gamma\gamma$ (twee fotonen in de eindtoestand) en $p\bar{p} \rightarrow h_c \rightarrow \gamma\eta_c \rightarrow \gamma\pi^0\pi^0\eta \rightarrow 7\gamma$ (met zeven fotonen in de eindtoestand). De tweede vereiste een uitgebreidere reconstructie, en kon door zijn hogere clusterveelvoud goed gebruikt worden om de invloed van geometrische overlap te bepalen. De enkele-doorgeefversie van het gedistribueerde clusterzoekalgoritme werd telkens gebruikt. De optima die bepaald zijn, zijn $\Delta\tau = 257, 97$ en 15 ns en $\Delta t = 282, 122$ en 40 ns voor respectievelijke interactiesnelheden van 200 kHz en minder, 2 MHz , en 20 MHz . De optimale burenafstand in de ruimte was één kristal, en een drempelwaarde van 3 MeV werd opgelegd aan de treffers.

Gebruikmakende van de voorgenoemde optima, konden de prestaties van de verschillende clusterzoekalgoritmes met elkaar vergeleken worden. Als eerste werd onderzocht wat de efficiëntie is voor de detectie van enkele fotonen. Fotonen met vier

energiewaardes werden gebruikt, en de resultaten zijn weergegeven in Figuur 6.11. Er vanuit gaande dat de efficiëntie onafhankelijk is van de energie, en dat de fotonen niet gecorrigeerd zijn aan elkaar, is de verwachting voor de efficiëntie van de twee- en zeven-fotonenkanalen respectievelijk $\{61.8\%, 65.6\%, 60.5\%, 42.6\%\}$ (ϵ_γ^2) en $\{18.5\%, 22.9\%, 17.3\%, 5.1\%\}$ (ϵ_γ^7), met ϵ_γ de efficiëntie voor de detectie van enkele fotonen. De daadwerkelijke resultaten van de simulatie staan in Figuur 6.13 en 6.14. De gevonden waarden voor het twee-fotonenkanaal zijn hoger dan verwacht, omdat de twee fotonen gecorrigeerd zijn. Over het algemeen fluctueert de efficiëntie weinig, en blijft het binnen de statistische fout. Alleen bij de hoogste interactiesnelheid treedt een 3,2% verschil op tussen het online algoritme en de gedistribueerde algoritmes, en zakt de efficiëntie van het bestaande algoritme met nog eens 2,2%. Bij het zeven-fotonenkanaal liggen de efficiënties van de gedistribueerde clusterzoekalgoritmes 5% lager dan de anderen, maar bij de 20 MHz interactiesnelheid zakt het verschil met de online versie naar 4%. Echter had het online algoritme erg veel tijd nodig om tot deze beste efficiëntie te komen, waardoor deze het systeem minder bruikbaar zou maken. Een test heeft uitgewezen dat de DCs in staat zijn om het preclusterzoekgedeelte van de gedistribueerde clusterzoekalgoritmes uit te voeren. Om deze reden wordt het gebruik van de dubbele-doorgeefversie van het gedistribueerd-clusterzoekalgoritme aangeraden.

Als laatste is onderzocht of het dataverzamelingsnetwerk in staat is om, naast data te verzamelen, ook het preclustersamenvoegalgoritme uit te voeren. De simulatie omvatte het aantal knooppunten, verbindingen, en lagen dat dit netwerk zou moeten hebben. Het ontwerp van dit netwerk kan gebruikt worden om de uiteindelijke vereisten voor het complete dataverwerkingssysteem te bepalen, en kan gebruikt worden als richtlijn in het ontwerp van het netwerk van andere subsystemen. Twee topologieën hebben de revue gepasseerd: een schatting voor het aantal apparaten dat nodig is gebaseerd op de verwachte datasnelheid, en een schatting, gebaseerd op de beschikbare bronnen op de rekeneenheden (Field-Programmable Gate Arrays, FP-GAs). Samenvattend, wordt het aangeraden om verder te gaan met de ontwikkeling van het dubbeledoorgeefgedistribueerde-clusterzoekalgoritme, gezien het goede prestaties tegen lage kosten levert, en het kan draaien op de voorziene apparatuur van het uitleessysteem.

Acknowledgements/Dankwoord

Now comes (to some, at least) one of the most important parts of the thesis, the acknowledgements. Let me start by extending my gratitude and respect for my promotor, Nasser. You provided insightful and useful comments on my work, and always had a keen eye for detail, something that I can appreciate. You supported me, both professionally and personally, which I will not forget. I enjoyed the discussions you started at the lunch table, and found the opinions and insights shared by the participants to be most enlightening. I also thank you and your wife for the hospitality, when you invited the group over for dinner or drinks.

Next, I thank my first co-promotor, Myroslav. You were always there to answer my questions, and I find you to be very knowledgeable and professional. You helped me find my way, and you provided many useful comments along the way. I also enjoyed the many trips we made together. In particular, I remember our first trip, which was to Alba in Italy. It was raining heavily, and we shared a poncho trying to find a place to eat (which we did). Of course, one of our last trips, to Strassbourg, France, was also memorable, as we were not allowed to go into our hotel by the military, because someone had left a bag unattended at a tram stop nearby. I also like to thank you and your wife, Oksana, for inviting us over for a barbecue, and for baking all those lovely cakes.

Of course, I cannot leave out my other co-promotor, Johan, who also has provided many useful insights and comments, which helped shape this thesis. Apart from this, you were always a pleasant company, and I enjoyed all the times you joined the group to go out for dinner, or to see a movie or the latest Game of Thrones episode. Later on, you and Ali found out that I have a PS4, and you convinced me to play Destiny with you, which I liked very much. You are very open, and I enjoyed our conversations about, well, things that were not always relevant. Mohammad Babai often played a role in these conversations. Speaking of which, on a more professional note, I am thankful for his help with ROOT and PandaROOT.

I am grateful for my reading committee members (in alphabetical order), prof. dr. Herbert Löhner, prof. dr. Wolfgang Kühn, and prof. dr. Antonio Pellegrino, for taking the time to carefully read my thesis, and for providing useful comments that helped to improve this thesis. In particular: Herbert, thanks for catching some potentially embarrassing mistakes.

I also like to extend a special thanks to Peter Schakel, who very patiently helped me test my algorithm on the FPGA prototype. You were always willing to discuss my problems (about the project) with me. Thank you for the suggestions to improve the algorithm, and for all the work you put into it.

The next person has to be Alex, a.k.a. ‘Greece man’, with whom I had many adventures. Our trip to Rauischholzhausen with Zahra and Solmaz (the latter didn’t always appreciate us during this week), the many *PANDA* meetings, Lunteren; we saw many places. Any, of course, there was Veldhoven, where we... eh... actively participated. You were always motivating everyone (I will leave out what for), and provided great support in calling everyone to lunch. And, last, but not least, you agreed to be one of my paranimphs. Thank you for everything. I will move on to your office mate, Zahra, who started her PhD around the same time as me. I will never forget our trip to China, with the crazy taxi ride to the Great Wall, and the many stairs in the Summer Palace. Solmaz, thank you for always being interested. I remember our race who could eat the most dessert in Veldhoven, which you (unfortunately) won.

Then there’s Josbert, our sunshine. Thank you for all the nice lunch discussions, and the advice about my thesis near the end of my PhD. Julia, thank you for all the nice conversations and support, and for introducing me to belayed climbing, which I enjoy doing very much. And thank you for Monster :) And the waffles. Olav, thank you for all the insightful questions during everyone’s talks, and for helping me with my motivation letter. Ali (Sr.), I enjoyed the many nights we spend trying to shoot down those stupid Cabal with Johan, and all the nice discussions we had (not only about Destiny). Brankica, a.k.a. ‘Bra’, thank for being such a nice office mate and my paranimph. And you brought nice chocolate with you. Of course, I also thank my other office mate, Ali (Jr.); like a true Iranian, you always had plenty of candy to share. I also enjoyed your nuts, which were nice and salty.

Christiaan. My first encounter with you was while I was grading your homework, which was about an order of magnitude longer than most others. Thank you for always being patient and ready to help, and of course for the chocolate fountain. Karol, thank you for the wonderful meat and sauces you brought to the barbecue, and for the times we went climbing. Ikechi, thank you for the nice discussions during lunch and the barbecue.

I also like to thank my former colleagues. Stefano, thank you for being such a nice office mate. We had a lot of fun together. Thank you Olga, Olena, Wouter, Soumya, Keri, and Faruk for all the nice activities we did together, like going to the movies or pain(t)ball.

Thank you, Maisam, Mohammad, Reza, and Hajar for the nice conversations. And thank you, Elwin, Olivier, Aernout, Artem, Manisha, Tom, my fellow PhD students. Thank you all! For those of you that were wondering: yes, the order in which I'm going through the list of people is mostly random. Please don't feel bad if you were one of the last people to be mentioned.

Last, but certainly not least, I would like to thank all my friends. In particular Gerjan and Lars, for always being there for me, and for all the good times we had (and hopefully continue to have). And thank you, father, mother, and brother, for your support.

Finally, I would like to thank the GSI Helmholtzzentrum für Schwerionenforschung for their financial support, which made this work possible.

With the risk of putting too much icing on the cake, I end with the following. To all that have stood by me, thank you for creating some wonderful memories for me. I will cherish them.

List of Acronyms

2STEP	Double-pass distributed version of cluster-finding algorithm
ASIC	Application-Specific Integrated Circuit
BBN	Burst-Building Network
BEPC	Beijing Electron-Positron Collider
BESIII	Beijing Spectrometer III
BGO	Bismuth Germanate
CFT	Constant-Fraction Timing
CKM	Cabibbo-Kobayashi-Maskawa
CMS	Compact Muon Solenoid
CN	Compute Node
CPU	Central Processing Unit
CR	Collector Ring
DC	Data Concentrator
DEF	Default version of cluster-finding algorithm
DIRC	Detection of Internally Reflected Cherenkov light
DIST	Single-pass distributed version of cluster-finding algorithm
DPM	Dual Parton Model
ECAL	Electromagnetic Calorimeter
EFT	Effective Field Theory
EM	Electromagnetic
EMC	Electromagnetic Calorimeter

ESR	Experimental Storage Ring
FAIR	Facility for Antiproton and Ion Research
FPGA	Field-Programmable Gate Array
FTF	FRITIOF
FWHM	Full Width at Half Maximum
GEM	Gas Electron Multiplier
GSI	Gesellschaft für Schwerionenforschung
HESR	High Energy Storage Ring
HLT	High-Level Trigger
HL	High Luminosity
HR	High Resolution
IP	Intellectual Property
KVI-CART	KVI - Center for Advanced Radiation Technology
(LA)APD	(Large-Area) Avalanche Photo Diode
L0/1/2	Level 0/1/2
LED	Light-Emitting Diode
LHC	Large Hadron Collider
LNP	Low Noise and Power
LQCD	Lattice Quantum Chromodynamics
LY	Light Yield
MC	Monte Carlo
MDC	Main Drift Chamber
MVD	Micro Vertex Detector
MWD	Moving Window Deconvolution
NRQCD	Non-Relativistic Quantum Chromodynamics
ONL	Online version of cluster-finding algorithm
<i>PANDA</i>	$\bar{p}$ Annihilations at Darmstadt
PDG	Particle Data Group
PID	Particle Identification

pNRQCD	potential Non-Relativistic Quantum Chromodynamics
Ps	Positronium
PWO	Lead Tungstate
QCD	Quantum Chromodynamics
QED	Quantum Electrodynamics
RMS	Root Mean Square
SADC	Sampling Analogue-to-Digital Converter
SFP	Small Formfactor Pluggable
(S)FRS	(Super) Fragment Separator
SIS	Super Ion Synchrotron
SODANET	Synchronisation Of the Data Acquisition NETwork
STT	Straw Tube Tracker
TOF	Time Of Flight
VHDL	VHSIC Hardware Description Language
VHSIC	Very High Speed Integrated Circuit
VPTT	Vacuum Photo Tetrode Tube

Bibliography

- [1] ATLAS Collaboration, Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, *Phys. Lett. B*, **716** (2012).
- [2] K.A. Olive *et al.* (Particle Data Group), *Chin. Phys. C*, **38**, 090001 (2014) and 2015 update.
- [3] BESIII Collaboration, The construction of the BESIII experiment, *Nucl. Instr. and Meth. in Ph. Research A*, **598** (2009).
- [4] X. Liu *et al.*, An overview of XYZ new particles, arXiv 1312.7408 [hep-ph] (2014).
- [5] S. Chatrchyan *et al.* (CMS Collaboration), *J. High Energy Phys.*, DOI:10.1007/JHEP04(2013)154, (2013).
- [6] Z. Q. Liu *et al.* (Belle Collaboration), Study of $e^+e^- \rightarrow \pi^+\pi^- J/\psi$ and Observation of a Charged Charmoniumlike State at Belle, *Phys. Rev. Lett.*, **110** (2013).
- [7] T. Xiao, S. Dobbs, A. Tomaradze, Kamal K. Seth, Observation of the charged hadron $Z_c^\pm(3900)$ and evidence for the neutral $Z_c^0(3900)$ in $e^+e^- \rightarrow \pi\pi J/\psi$ at $\sqrt{s} = 4170$ MeV, *Phys. Lett. B*, **727** (2013).
- [8] H.W. Ke, Z.T. Wei, and X.Q. Li, *Eur. Phys. J. C*, **73**: 2561, DOI:10.1140/epjc/s10052-013-2561-0 (2013).
- [9] M. Ablikim *et al.* (BESIII Collaboration), Observation of a Charged Charmoniumlike Structure in $e^+e^- \rightarrow \pi^+\pi^- J/\psi$ at $\sqrt{s} = 4.26$ GeV, *Phys. Rev. Lett.*, **110** (2013).
- [10] M. Ablikim *et al.* (BESIII Collaboration), Observation of a Charged $(D\bar{D})^\pm$ Mass Peak in $e^+e^- \rightarrow \pi D\bar{D}$ at $\sqrt{s} = 4.26$ GeV, *Phys. Rev. Lett.*, **112** (2014).

-
- [11] <http://physics.aps.org/articles/v6/139>, *Physics*, **6**: 139 (2013).
 - [12] S.-K. Choi *et al.* (Belle Collaboration), Observation of a Narrow Charmonium-like State in Exclusive $B^\pm \rightarrow K^\pm \pi^+ \pi^- J/\psi$ Decays, *Phys. Rev. Lett.*, **91** (2003).
 - [13] M. Ablikim *et al.* (BESIII Collaboration), Observation of $e^+e^- \rightarrow \gamma X(3872)$ at BESIII, *Phys. Rev. Lett.*, **112** (2014).
 - [14] BELLE Collaboration, The BELLE Detector, *Nucl. Instr. and Meth. in Ph. Research A*, **479** (2002).
 - [15] E. Prencipe, J.S. Lange, and A. Blinov, New Spectroscopy with $\bar{P}ANDA$ at FAIR: X, Y, Z and the F-Wave Charmonium States, *AIP Conference Proceedings*, **1735** 060011 DOI:10.1063/1.4949447 (2016).
 - [16] CMS Collaboration, The CMS Experiment at the CERN LHC, *JINST*, **3** (2008).
 - [17] LHCb Collaboration, The LHCb Detector at the LHC, *JINST*, **3** (2008).
 - [18] N. Brambilla *et al.*, Heavy Quarkonium: Progress, Puzzles, and Opportunities, *Eur. Phys. J. C*, **71**, DOI: 10.1140/epjc/s10052-010-1534-9 (2011).
 - [19] K. Hagiwara *et al.*, Review of Particle Physics, *Phys. Rev. D*, **66**, 010001 (2002).
 - [20] http://www.drcruzan.com/Chemistry_Electrons.html
 - [21] J.J. Aubert *et al.*, Experimental Observation of a Heavy Particle J, *Phys. Rev. Lett.*, **33** (1974).
 - [22] J.-E. Augustin *et al.*, Discovery of a Narrow Resonance in e^+e^- Annihilation, *Phys. Rev. Lett.*, **33** (1974).
 - [23] K. Pachucki, and S.G. Karshenboim, Complete Results for Positronium Energy Levels at Order $m\alpha^6$, *Phys. Rev. Lett.*, **80** (1998).
 - [24] V. Khachatryan, A.M. Sirunyan, A. Tumasyan *et al.* (CMS Collaboration), Measurement of the Inclusive 3-Jet Production Differential Cross Section in Proton-Proton Collisions at 7 TeV and Determination of the Strong Coupling Constant in the TeV Range, *Eur. Phys. J. C*, **75** DOI:10.1140/epjc/s10052-015-3376-y (2015).
 - [25] S. Bethke, The 2009 World Average of α_s , *Eur. Phys. J. C*, **64** DOI:10.1140/epjc/s10052-009-1173-1 (2009).

- [26] N. Brambilla, A. Pineda, J. Soto, and A. Vairo, Effective Field Theories for Heavy Quarkonium, *Rev. Mod. Phys.* **77**, DOI:10.1103/RevModPhys.77.1423 (2005).
- [27] K. G. Chetyrkin, J. H. Kühn, A. Maier, P. Maierhöfer, P. Marquard, M. Steinhauser, and C. Sturm, Charm and Bottom Quark Masses: an Update, *Phys. Rev. D* **80**, DOI:10.1103/PhysRevD.80.074010 (2009).
- [28] A. Signer, The Charm Quark Mass from Non-Relativistic Sum Rules, *Phys. Lett. B*, **672**, DOI:10.1016/j.physletb.2009.01.028 (2009).
- [29] A. Pineda, A. Signer, Renormalization Group Improved Sum Rule Analysis for the Bottom Quark Mass, *Phys. Rev. D* **73**, DOI:10.1103/PhysRevD.73.111501 (2006).
- [30] N. Brambilla, Y. Sumino, and A. Vairo, Quarkonium Spectroscopy and Perturbative QCD: Massive Quark-Loop Effects, *Phys. Rev. D*, **65**, DOI:10.1103/PhysRevD.65.034001 (2002).
- [31] A. Abulencia *et al.* (CDF Collaboration), Evidence for the exclusive decay B_c^{+-} to $J/\psi\pi^{+-}$ and measurement of the mass of the B_c meson, *Phys. Rev. D*, **96**, DOI:10.1103/PhysRevLett.96.082002 (2006).
- [32] B.A. Kniehl, A.A. Penin, A. Pineda, V.A. Smirnov, and M. Steinhauser, $M(\eta b)$ and α_s from Nonrelativistic Renormalization Group, *Phys. Rev. Lett.*, **92** (2004).
- [33] Z. Hadaddi, Ground State Studies of Charmonium with BESIII, (Ph.D. Thesis), University of Groningen (*to be published*) (2017).
- [34] A.A. Penin, A. Pineda, V.A. Smirnov, and M. Steinhauser, $M(B_c^*) - M(B_c)$ Splitting from Nonrelativistic Renormalization Group, *Phys. Lett. B*, **593** (2004).
- [35] P. Colangelo, P. Santorelli, E. Scrimieri, Determining α_s from the Hyperfine Splitting $M(v(1S)) - M(\eta_b)$, *Phys. Rev. D*, **83**, DOI:10.1103/PhysRevD.83.014016 (2011).
- [36] N. Brambilla, A. Pineda, J. Soto, and A. Vairo, The QCD Potential at $O(1/m)$, *Phys. Rev. D*, **63**, (2000).
- [37] A. Bazavov, C. Bernard, C. DeTar, S. Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, P.B. Mackenzie, M.B. Oktay, R. Sugar, D. Toussaint, and R.S. Van de Water, Full Nonperturbative QCD Simulations

- with 2+1 Flavors of Improved Staggered Quarks, *Rev. Mod. Phys.*, **82**, DOI:10.1103/RevModPhys.82.1349 (2010).
- [38] Z. Fodor, C. Hoelbling, S. Krieg, L. Lellouch, Th. Lippert, A. Portelli, A. Sastre, K.K. Szabo, and L. Varnhorst (Budapest-Marseille-Wuppertal collaboration), Up and down quark masses and corrections to Dashen's theorem from lattice QCD and quenched QED, *Phys. Rev. Lett.*, **117** 082001 (2016).
- [39] E. Follana, Q. Mason, C. Davies, K. Hornbostel, G.P. Lepage, J. Shigemitsu, H. Trotter, and K. Wong, Highly Improved Staggered Quarks on the Lattice, with Applications to Charm Physics, *Phys. Rev. D* **75**, DOI:10.1103/PhysRevD.75.054502 (2007).
- [40] S. Okubo, Consequences of quark-line (Okubo-Zweig-Iizuka) rule, *Phys. Rev. D* **16** 2336 (1977).
- [41] L. Liu, G. Moir *et al.* (Hadron Spectrum Collaboration), Excited and Exotic Charmonium Spectroscopy from Lattice QCD, *J. High Energ. Phys.*, DOI:10.1007/JHEP07(2012)126 (2012).
- [42] H.H. Gutbrod *et al.*, FAIR - Baseline Technical Report, ISBN 3-9811298-0-6 (2006).
- [43] Th. Stöhlker *et al.*, Atomic Physics with Highly-Charged Ions at the Future FAIR Facility: A Status Report, *Nucl. Instr. and Meth. in Ph. Research B*, **261** (2007).
- [44] P. Senger *et al.* (CBM Collaboration), Compressed Baryonic Matter: Experiments at GSI and FAIR, *Phys. Part. Nucl.* **39** (2008).
- [45] R. Krücken *et al.* (NuSTAR Collaboration), The Nustar Facility at FAIR, *J. Phys. G* **31** (2005).
- [46] K. Föhl *et al.* (PANDA Collaboration), The Panda Detector at the Future FAIR Laboratory, *Eur. Phys. J. ST* **162** (2008).
- [47] A. Lehrach *et al.*, *Int. J. Mod. Phys. E* **18**, **420**, DOI:10.1142/S021830130901246X (2009).
- [48] PANDA Collaboration, Physics Performance Report for PANDA: Strong Interaction Studies with Antiprotons, arXiv:0903.3905 [hep-ex] (2009).

- [49] PANDA Collaboration, Technical Design Report for the PANDA Internal Targets: the Cluster-Jet Target and Developments for the Pellet Target, arXiv:1404.5988 (2012).
- [50] PANDA Collaboration, Technical Design Report for: PANDA Micro Vertex Detector, arXiv:1207.6581 [physics.ins-det] (2012).
- [51] W. Erni, I. Keshelashvili, B. Krusche, *et al.*, *Eur. Phys. J. A* **49**, DOI:10.1140/epja/i2013-13025-8 (2013).
- [52] PANDA Collaboration, Technical Design Report for: PANDA Electromagnetic Calorimeter (EMC), arXiv:0810.1216 [physics.ins-det] (2008).
- [53] M. Kavatsyuk, D. Bremer, V. Dormenev, P. Drexler, T. Eissner, W. Erni, E. Guliyeu, T. Hennino, B. Krusche, B. Lewandowski, H. Löhner, M. Moritz, R.W. Novotny, K. Peters, J. Pouthas, P. Rosier, M. Steinacher, G. Tambave, A. Wilms (PANDA Collaboration), Performance of the Prototype of the Electromagnetic Calorimeter for PANDA, *Nucl. Instrum. and Meth. in Ph. Research A*, **648** (2011).
- [54] N. Akopov *et al.*, *Nucl. Instrum. Meth. A* **479**, **511** (2002).
- [55] *Nucl. Instr. and Meth. in Ph. Research A*, **598** (2009).
- [56] M. Ablikim *et al.* (BESIII Collaboration), Searches for Isospin-Violating Transitions $\chi_{c0,2} \rightarrow \pi^0 \eta_c$, *Phys. Rev. D*, **91**, 112018 (2015).
- [57] X.-D. Cheng, K.-L. He, H.-B. Li, Y.-F. Wang, and M.-Z. Yang, *Phys. Rev. D* **75**, 094019 (2007).
- [58] N. Cabibbo, *Phys. Rev. Lett.*, **10**, (1963) 531.
- [59] M. Kobayashi, T. Maskawa, *Prog. Theor. Phys.* **49**, **652** (1973).
- [60] C. Amsler, M. Doser *et al.* (Particle Data Group), Review of Particle Physics, *Phys. Letters B*, **667** (2008).
- [61] CMS Collaboration, ECAL Technical Design Report, CERN/LHCC 97-33 (1997).
- [62] R. Sumner, *Nucl. Instrum. Meth. A* **265**, **252** (1988).
- [63] F. Ghio, *et al.*, *Nucl. Instr. and Meth. A*, **404** (1998).

- [64] A. Annenkov, A. Borisevitch, A. Hofstätter, M. Korzhik, V. Ligun, P. Lecoq, O. Missevitch, R. Novotny, and J.P. Peigneux, Improved Light Yield of Lead Tungstate Scintillators, *Nucl. Inst. and Meth. in Ph. Research A* **450** (2000).
- [65] M. Böhm *et al.*, Influence of Mo Impurity on the Spectroscopic and Scintillation Properties of PbWO₄ Crystals, *Phys. Stat. Sol. A* **167**, **243** (1998).
- [66] V. Dormenev, T. Kuske, R.W. Novotny, A. Borisevich, A. Fedorov, M. Korjik, V. Mechinski, O. Missevitch, and S. Lugert, Stimulated Recovery of the Optical Transmission of PbWO₄ Scintillation Crystals for Electromagnetic Calorimeters after Radiation Damage, *Nucl. Inst. and Meth. in Ph Research A*, Volume **623**, (2010).
- [67] <http://www.hamamatsu.com/eu/en/product/category/3100/4001/index.html>
- [68] I. Keshelashvili, E. Werner, M. Steinacher, and B. Krusche, Preamplifier development for high count-rate, large dynamic range readout of inorganic scintillators, *Verhandlungen der Deutschen Physikalischen Gesellschaft* (Dresden2013issue), (2013).
- [69] P. Wiczorek and H. Flemming, Low noise preamplifier ASIC for the PANDA EMC, *Conf. Record, IEEE Nuclear Science Symposium*, October 30 - November 6, 2010, DOI:10.1109/NSSMIC.2010.5873982, <http://ieeexplore.ieee.org> (2010).
- [70] G. Tambave *et al.*, Pulse Pile-up Recovery for the Front-End Electronics of the PANDA Electromagnetic Calorimeter, *JINST*, **7** 11001 (2012).
- [71] W. Kühn, *et al.*, FPGA-Based Compute Nodes for the PANDA Experiment at FAIR, *Conf. Record, IEEE Nuclear Science Symposium*, 29 April - 4 May, 2007, DOI:10.1109/RTC.2007.4382729, <http://ieeexplore.ieee.org> (2007).
- [72] <https://www.xilinx.com/support.html#documentation>
- [73] M. Kavatsyuk *et al.*, *Conf. Record, IEEE Nuclear Science Symposium*, October 23-29, 2011, DOI:10.1109/NSSMIC.2011.6154360, <http://ieeexplore.ieee.org> (2011).
- [74] D. Bertini, M. A.-Turany, I. König and F. Uhlig, *IOP Publishing - Journal of Physics: Conf. Series*, Volume **119**, Part 3 (2008).
- [75] S. Spataro (PANDA Collaboration), *IOP Publishing - Journal of Physics: Conf. Series*, Volume **331**, Part 3 (2010).

- [76] D.J. Lange, *Nucl. Inst. and Meth. in Ph. Research A* 462 (2001).
- [77]] A. Capella, U. Sukhatme, C.I. Tan, J. Tran Thanh Van, *Phys. Rep.* 236, **225** (1994).
- [78] B. Andersson, G. Gustafson, H. Pi, *Z. Phys. C* 57, **485** (1993).
- [79] S. Agostinelli *et al.*, *Nucl. Instr. and Meth. in Ph. Research A* 506 (2003).
- [80] Ph. Mahlberg, Timebased Simulated Waveforms for the Forward Endcap of the PANDA Electromagnetic Calorimeter (Master Thesis), Helmholtz-Institut für Strahlen- und Kernphysik, Friedrich-Wilhelms-Univ. Bonn (2013).
- [81] MWD filter reference
- [82] M. Kavatsyuk *et al.*, *Conf. Record, IEEE Nuclear Science Symposium*, October 27 - November 3, 2012, DOI:10.1109/NSSMIC.2012.6551420, <http://ieeexplore.ieee.org> (2012).
- [83] J. Michel, Ph.D. Thesis, Johann Wolfgang Göthe-Universität, Frankfurt am Main (2012).