


Technical Design Report for the: \bar{P} ANDA Data Acquisition and Event Filtering

(AntiProton Annihilations at Darmstadt)

Strong Interaction Studies with Antiprotons

\bar{P} ANDA Collaboration

August 4, 2020



The $\overline{\text{PANDA}}$ Collaboration

2019-12-16 14:47:34

Università Politecnica delle Marche-Ancona, **Ancona**, Italy

G. Barucca, F. Davì, G. Lancioni, P. Mengucci, L. Montalto, P. P. Natali, N. Paone, D. Rinaldi,
L. Scalise

Universität Basel, **Basel**, Switzerland

W. Erni, B. Krusche, M. Steinacher, N. Walford

Institute of High Energy Physics, Chinese Academy of Sciences, **Beijing**, China

N. Cao, Z. Liu, C. Liu, B. Liu, X. Shen, S. Sun, J. Tao, X. A. Xiong, G. Zhao, J. Zhao

Ruhr-Universität Bochum, Institut für Experimentalphysik I, **Bochum**, Germany

M. Albrecht, S. Bökelmann, F. Feldbauer, M. Fink, J. Frech, V. Freudenreich, M. Fritsch, R. Hagdorn,
F.H. Heinsius, T. Held, T. Holtmann, I. Keshk, H. Koch, B. Kopf, M. Kuhlmann, M. Kümmer,
M. Kückner, A. Mustafa, M. Pelizäus, A. Pitka, J. Reher, G. Reicherz, M. Richter, C. Schnier, L. Sohl,
C. Sowa, M. Steinke, T. Triffterer, T. Weber, C. Wenzel, U. Wiedner

Department of Physics, Bolu Abant İzzet Baysal University, **Bolu**, Turkey

H. Denizli, N. Er

Rheinische Friedrich-Wilhelms-Universität Bonn, **Bonn**, Germany

R. Beck, C. Hammann, J. Hartmann, B. Ketzer, J. Müllers, M. Rossbach, B. Salisbury, C. Schmidt,
U. Thoma, M. Urban

Università di Brescia, **Brescia**, Italy

A. Bianconi

Institutul National de C&D pentru Fizica si Inginerie Nucleara "Horia Hulubei", **Bukarest-Magurele**,
Romania

M. Bragadireanu, D. Pantea

University of Technology, Institute of Applied Informatics, **Cracow**, Poland

M. Domagala, G. Filo, E. Lisowski, F. Lisowski, J. Płazek

IFJ, Institute of Nuclear Physics PAN, **Cracow**, Poland

K. Korcyl, A. Kozela, P. Lebedowicz, K. Pysz, W. Schäfer, A. Szczurek

AGH, University of Science and Technology, **Cracow**, Poland

T. Fiutowski, M. Idzik, K. Swientek, P. Terlecki

Instytut Fizyki, Uniwersytet Jagiellonski, **Cracow**, Poland

G. Korcyl, R. Lalik, A. Malige, P. Moskal, K. Nowakowski, W. Przygoda, N. Rathod, Z. Rudy,
P. Salabura, J. Smyrski

FAIR, Facility for Antiproton and Ion Research in Europe, **Darmstadt**, Germany

I. Augustin, R. Böhm, I. Lehmann, L. Schmitt, V. Varentsov

GSI Helmholtzzentrum für Schwerionenforschung GmbH, **Darmstadt**, Germany

M. Al-Turany, A. Belias, H. Deppe, R. Dzhygadlo, H. Flemming, A. Gerhardt, K. Götzen, A. Heinz,
R. Karabowicz, S. Koch, U. Kurilla, D. Lehmann, J. Lühning, U. Lynen, H. Orth, K. Peters, J. Rieger,
T. Saito, G. Schepers, C. J. Schmidt, C. Schwarz, J. Schwiening, A. Täschner, M. Traxler, B. Voss,
P. Wieczorek

Joint Institute for Nuclear Research, **Dubna**, Russia

V. Abazov, G. Alexeev, V. A. Arefiev, V. Astakhov, M. Yu. Barabanov, B. V. Batyunya, V.
Kh. Dodokhov, A. Fechtchenko, A. Galoyan, G. Golovanov, E. K. Koshurnikov, Y. Yu. Lobanov, A.
G. Olshevskiy, A. A. Piskun, A. Samartsev, S. Shimanski, A. N. Skachkova, E. A. Strokovsky,
V. Tokmenin, V. Uzhinsky, A. Verkheev, A. Vodopianov, N. I. Zhuravlev

University of Edinburgh, **Edinburgh**, United Kingdom
D. Branford, D. Watts

Friedrich-Alexander-Universität Erlangen-Nürnberg, **Erlangen**, Germany
M. Böhm, W. Eyrich, A. Lehmann, D. Miehling, M. Pfaffinger, S. Stelter

Northwestern University, **Evanston**, U.S.A.
N. Quin, L. Robison, K. Seth, T. Xiao

Università di Ferrara and INFN Sezione di Ferrara, **Ferrara**, Italy
D. Bettoni

Goethe-Universität, Institut für Kernphysik, **Frankfurt**, Germany
A. Ali, A. Hamdi, M. Himmelreich, M. Krebs, S. Nakhoul, F. Nerling

Frankfurt Institute for Advanced Studies, **Frankfurt**, Germany
A. Belousov, I. Kisel, G. Kozlov, M. Pugach, M. Zyzak

INFN Laboratori Nazionali di Frascati, **Frascati**, Italy
N. Bianchi, P. Gianotti, V. Lucherini

Dept of Physics, University of Genova and INFN-Genova, **Genova**, Italy
G. Bracco

Justus-Liebig-Universität Gießen II. Physikalisches Institut, **Gießen**, Germany
S. Bodenschatz, K.T. Brinkmann, L. Brück, S. Diehl, V. Dormenev, M. Düren, T. Erlen,
E. Etzelmüller, K. Föhl, C. Hahn, A. Hayrapetyan, J. Hofmann, S. Kegel, M. Kesselkaul, I. Köseoglu,
A. Kripko, W. Kühn, J. S. Lange, V. Metag, M. Moritz, M. Novova, R. Novotny, P. Orsich,
J. Pereira-de-Lira, M. Peter, M. Sachs, M. Schmidt, R. Schubert, H. Stenzel, M. Straube, M. Strickert,
U. Thöring, T. Wasem, C. Will, B. Wohlfahrt, H.G. Zaunick

IRFU, CEA, Université Paris-Saclay, **Gif-sur-Yvette Cedex**, France
E. Tomasi-Gustafsson

University of Glasgow, **Glasgow**, United Kingdom
D. Glazier, D. Ireland, B. Seitz

Birla Institute of Technology and Science, Pilani, K K Birla Goa Campus, **Goa**, India
P.N. Deepak, A. Kulkarni

KVI-Center for Advanced Radiation Technology (CART), University of Groningen, **Groningen**,
Netherlands

R. Kappert, M. Kavatsyuk, H. Loehner, J. Messchendorp, V. Rodin, P. Schakel, S. Vejdani

Gauhati University, Physics Department, **Guwahati**, India
K. Dutta, K. Kalita

University of Science and Technology of China, **Hefei**, China
G. Huang, D. Liu, H. Peng, H. Qi, Y. Sun, X. Zhou

Universität Heidelberg, **Heidelberg**, Germany
M. Kunze

Fachhochschule Südwestfalen, **Iserlohn**, Germany
H. Sohlbach

Department of Physics, Faculty of Arts & Sciences, Dogus University, **Istanbul**, Turkey
K. Azizi

Forschungszentrum Jülich, Institut für Kernphysik, **Jülich**, Germany
L. Bianchi, A. Derichs, R. Dossdall, A. Erven, W. Esmail, A. Gillitzer, F. Goldenbaum, D. Grunwald,
L. Jokhovets, J. Kannika, P. Kulesa, A. Lai, S. Orfanitski, G. Pérez Andrade, D. Prasuhn, E. Prencipe,
J. Pütz, J. Ritman, E. Rosenthal, S. Schadmand, R. Schmitz, A. Scholl, T. Sefzick, V. Serdyuk,
G. Sterzenbach, T. Stockmanns, D. Veretennikov, P. Wintz, P. Wüstner, H. Xu, Y. Zhou

Chinese Academy of Science, Institute of Modern Physics, **Lanzhou**, China
X. Cao, Q. Hu, Z. Li, H. Li, Y. Liang, X. Ma

INFN Laboratori Nazionali di Legnaro, **Legnaro**, Italy
V. Rigato

Lunds Universitet, Department of Physics, **Lund**, Sweden
L. Isaksson

Johannes Gutenberg-Universität, Institut für Kernphysik, **Mainz**, Germany

P. Achenbach, A. Aycock, O. Corell, A. Denig, M. Distler, M. Hoek, W. Lauth, H. H. Leithoff, Z. Liu,
H. Merkel, U. Müller, J. Pochodzalla, S. Schlimme, C. Sfienti, M. Thiel, M. Zambrana

Helmholtz-Institut Mainz, **Mainz**, Germany

S. Ahmed, S. Bleser, M. Bölting, L. Capozza, A. Dbeyssi, A. Ehret, P. Grasemann, R. Klasen,
R. Kliemt, F. Maas, S. Maldaner, M. Michel, C. Morales Morales, C. Motzko, O. Noll, S. Pflüger,
D. Rodríguez Piñeiro, F. Schupp, M. Steinen, S. Wolff, I. Zimmermann

Research Institute for Nuclear Problems, Belarus State University, **Minsk**, Belarus

A. Fedorov, D. Kazlou, M. Korzhik, O. Missevitch

Institute for Theoretical and Experimental Physics named by A.I. Alikhanov of National Research
Centre "Kurchatov Institute", **Moscow**, Russia

P. Balanutsa, V. Chernetsky, A. Demekhin, A. Dolgolenko, P. Fedorets, A. Gerasimov, A. Golubev,
V. Goryachev, A. Kantsyrev, D. Y. Kirin, A. Kotov, N. Kristi, E. Ladygina, E. Lushevskaya, V.
A. Matveev, V. Panjushkin, A. V. Stavinskiy

Moscow Power Engineering Institute, **Moscow**, Russia

A. Balashoff, A. Boukharov, O. Malyshev

Nuclear Physics Division, Bhabha Atomic Research Centre, **Mumbai**, India

K. N. Basant, H. Kumawat, B. Roy, A. Saxena, S. Yogesh

Westfälische Wilhelms-Universität Münster, **Münster**, Germany

D. Bonaventura, C. Fritzscher, S. Grieser, C. Hargens, A.K. Hergemöller, B. Hetz, N. Hüskens, J. Kellers,
A. Khoukaz

Suranaree University of Technology, **Nakhon Ratchasima**, Thailand

C. Herold, K. Khosonthongkee, C. Kobdaj, A. Limphirat, T. Nasawad, T. Simantathammakul,
P. Srisawad, Y. Yan

Nankai University, **Nankai**, China

C. Yu, X. Zhang, W. Zhu

Novosibirsk State University, **Novosibirsk**, Russia

A. E. Blinov, S. Kononov, E. A. Kravchenko

Budker Institute of Nuclear Physics, **Novosibirsk**, Russia

E. Antokhin, A. Yu. Barnyakov, K. Beloborodov, V. E. Blinov, I. A. Kuyanov, S. Pivovarov, E. Pyata,
Y. Tikhonov

Institut de Physique Nucléaire, CNRS-IN2P3, Univ. Paris-Sud, Université Paris-Saclay, 91406, **Orsay**
cedex, France

R. Kunne, B. Ramstein

University of Wisconsin Oshkosh, **Oshkosh**, U.S.A.

G. Hunter, M. Lattery, H. Pace

Dipartimento di Fisica, Università di Pavia, INFN Sezione di Pavia, **Pavia**, Italy

G. Boca

University of West Bohemia, **Pilsen**, Czech

D. Duda

Charles University, Faculty of Mathematics and Physics, **Prague**, Czech Republic

M. Finger, M. Finger, Jr., A. Kveton, M. Pesek, M. Peskova, I. Prochazka, M. Slunceka, V. Slunceckova

Czech Technical University, Faculty of Nuclear Sciences and Physical Engineering, **Prague**, Czech
Republic

P. Gallus, V. Jary, O. Korchak, M. Marcisovsky, G. Neue, J. Novy, L. Tomasek, M. Tomasek, M. Virius,

V. Vrba

Institute for High Energy Physics, **Protvino**, Russia

V. Abramov, S. Bukreeva, S. Chernichenko, A. Derevschikov, V. Ferapontov, Y. Goncharenko, A. Levin, E. Maslova, Y. Melnik, A. Meschanin, N. Minaev, V. Mochalov, V. Moiseev, D. Morozov, L. Nogach, S. Poslavskiy, A. Ryazantsev, S. Ryzhikov, P. Semenov, I. Shein, A. Uzunian, A. Vasiliev, A. Yakutin

Sikaha-Bhavana, Visva-Bharati, WB, **Santiniketan**, India

U. Roy

University of Sidney, School of Physics, **Sidney**, Australia

B. Yabsley

National Research Centre "Kurchatov Institute" B. P. Konstantinov Petersburg Nuclear Physics Institute, Gatchina, **St. Petersburg**, Russia

S. Belostotski, G. Fedotov, G. Gavrilov, A. Izotov, S. Manaenkov, O. Miklukho, A. Zhdanov

Kungliga Tekniska Högskolan, **Stockholm**, Sweden

A. Atac, T. Bäck, B. Cederwall

Stockholms Universitet, **Stockholm**, Sweden

K. Makonyi, M. Preston, P.E. Tegner, D. Wölbing

Sardar Vallabhbai National Institute of Technology, Applied Physics Department, **Surat**, India

K. Gandhi, A. K. Rai

Veer Narmad South Gujarat University, Department of Physics, **Surat**, India

S. Godre

Florida State University, **Tallahassee**, U.S.A.

V. Crede, S. Dobbs, P. Eugenio, D. Lersch

Università di Torino and INFN Sezione di Torino, **Torino**, Italy

M. P. Bussa, S. Spataro

INFN Sezione di Torino, **Torino**, Italy

D. Calvo, P. De Remigis, A. Filippi, G. Mazza, A. Rivetti, R. Wheadon

Politecnico di Torino and INFN Sezione di Torino, **Torino**, Italy

F. Iazzi, A. Lavagno

Università di Trieste and INFN Sezione di Trieste, **Trieste**, Italy

A. Martin

Uppsala Universitet, Institutionen för fysik och astronomi, **Uppsala**, Sweden

A. Akram, H. Calen, W. Ikegami Andersson, T. Johansson, A. Kupsc, P. Marciniewski, M. Papenbrock, J. Regina, K. Schönning, M. Wolke

Instituto de Física Corpuscular, Universidad de Valencia-CSIC, **Valencia**, Spain

J. Diaz

Sardar Patel University, Physics Department, **Vallabh Vidynagar**, India

V. Pothodi Chackara

National Centre for Nuclear Research, **Warsaw**, Poland

A. Chlopik, G. Kesik, D. Melnychuk, A. Trzcinski, M. Wojciechowski, S. Wronka, B. Zwieglinski

Österreichische Akademie der Wissenschaften, Stefan Meyer Institut für Subatomare Physik, **Wien**, Austria

C. Amsler, P. Bühler, N. Kratochwil, J. Marton, W. Nalti, D. Steinschaden, K. Suzuki, E. Widmann, S. Zimmermann, J. Zmeskal

Karlsruhe Institute of Technology, Institute for Data Processing and Electronics, **Karlsruhe**, Germany

M. Caselle, S. Chilingaryan, A. Kopmann

Editors: Myroslav Kavatsyuk Email: m.kavatsyuk@rug.nl
 Wolfgang Kühn Email: w.kuehn@physik.uni-giessen.de
 Klaus Götzen Email: k.goetzen@gsi.de
 Tobias Stockmanns Email: t.stockmanns@fz-juelich.de
 Johan Messchendorp Email: j.g.messchendorp@rug.nl

Technical Coordinator: Lars Schmitt Email: l.schmitt@gsi.de
Deputy: Anastasios Belias Email: a.belias@gsi.de
Spokesperson: Klaus Peters Email: klaus.peters@gsi.de
Deputy: Tord Johansson Email: tord.johansson@tsl.uu.se

Preface

This document describes the technical layout and the expected performance of the Data Acquisition system for the \bar{P} ANDA experiment (Phase 1).

The use of registered names, trademarks, *etc.* in this publication does not imply, even in the absence of specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

Contents

Preface	vii	3.3 On-line storage and requirements for event filtering	23
1 Executive Summary	1	3.4 DAQ partitioning and running modes	24
2 The $\bar{\text{P}}\text{ANDA}$ Experiment	3	4 System Architecture	25
2.1 The $\bar{\text{P}}\text{ANDA}$ Experiment	3	4.1 Basic building blocks of the system .	25
2.1.1 A worldwide unique facility for antiprotons	3	4.1.1 The Data Concentrator module .	25
2.1.2 Hadron spectroscopy	3	4.1.2 FPGA based Compute Nodes . .	25
2.1.3 Hyperon physics	4	4.1.3 FPGA based HPC interface . . .	28
2.1.4 Proton structure	4	4.2 Protocols, data formats, interfaces and data flow	28
2.1.5 Hadrons in nuclei	5	4.2.1 Synchronization protocol SO-DANET	29
2.1.6 High-Energy Storage Ring	5	4.2.2 Data Flow	37
2.1.7 Targets	5	4.2.3 Burst Building Network	40
2.1.8 Luminosity Considerations	6	4.2.4 DAQ partitioning, implementation	41
2.2 The $\bar{\text{P}}\text{ANDA}$ Detector	7	4.2.5 High-speed ROCE links and embedded GPU/AI computing nodes	41
2.2.1 Target Spectrometer	7	4.3 Event filtering and partitioning of algorithms	42
2.2.2 Forward Spectrometer	7	4.3.1 FPGA based online tracking algorithm for the STT	42
2.2.3 The Particle Identification System	8	4.3.2 FPGA based online clustering algorithm for the EMC	47
2.2.4 Infrastructure	8	4.3.3 EMC clustering. Performance Assessment by Hardware Simulation	54
3 DAQ requirements	13	4.4 Run Control, error handling and data quality monitoring	56
3.1 Event rates and pile-up situation for Phase 1 Physics	13	5 Performance	59
3.1.1 Event Rates	13	5.1 Simulations	59
3.1.2 Event Overlap	13	5.1.1 Framework	60
3.1.3 Online Event Building	14	5.1.2 Results	60
3.2 Readout of PANDA subsystems	15	5.2 Measurement with DAQ prototype .	63
3.2.1 Micro Vertex Detector	15	6 Project managements and resources	75
3.2.2 Straw-Tube Tracker	16	6.1 Work-packages and responsibilities .	75
3.2.3 Barrel DIRC	18	6.2 Schedule and Milestones	77
3.2.4 Endcap Disc DIRC	18	6.3 Cost	78
3.2.5 Barrel TOF	18		
3.2.6 Muon detectors	18		
3.2.7 Electromagnetic Calorimeter	19		
3.2.8 Forward Tracker	20		
3.2.9 Luminosity Detector	22		
3.2.10 Readout summary	23		

1 Executive Summary

The $\bar{\text{P}}\text{ANDA}$ Experiment

$\bar{\text{P}}\text{ANDA}$ [1] will be one of the four flagship experiments at the new international accelerator complex FAIR (Facility for Antiproton and Ion Research) in Darmstadt, Germany. With the $\bar{\text{P}}\text{ANDA}$ detector unique experiments will be performed using the high-quality antiproton beam within a momentum range from 1.5 GeV/c to 15 GeV/c, stored in the HESR (High Energy Storage Ring) [2]. The exploration of fundamental questions of hadron physics in the charm and multi-strange hadron sectors will deliver essential contributions to many open questions of QCD. The scientific program of $\bar{\text{P}}\text{ANDA}$ [3] includes hadron spectroscopy, properties of hadrons in matter, nucleon structure, hypernuclei and much more like the exotic bound quark states with or without gluonic degrees of freedom. The cooled antiproton beam colliding with a fixed proton or nuclear target will allow hadron production and formation experiments with a luminosity of up to $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ in the fully completed version of the facility. In the Modular Start Version (MSV) the luminosity will be $1 \times 10^{31} \text{ cm}^{-2}\text{s}^{-1}$.

[3] PANDA Collaboration. Physics Performance Report for PANDA: Strong Interaction Studies with Antiprotons. *arxiv:0903.3905*, 2009.

$\bar{\text{P}}\text{ANDA}$ Overview: Data Acquisition

The $\bar{\text{P}}\text{ANDA}$ experiment adopts a free-running data acquisition concept in order to allow as much flexibility as possible which the complex and diverse physics objectives of the experiment require, and also to fully exploit the high interaction rate of up to 2×10^7 events/s. Each sub-detector system runs autonomously in a self-triggering mode, yet synchronised with a high-precision time distribution system, SODANET. Zero-suppressed and physically relevant signals are transmitted to a high-bandwidth computing network implementing a software trigger. Without a selection of data the whole data rate could be as high as 200 GB/s. The data acquisition system aims for an online data reduction of factor 100 – 1000.

Bibliography

- [1] PANDA Collaboration. Technical Progress Report, FAIR-ESAC/Pbar. 2005.
- [2] D. Prasuhn. Status HESR. PANDA meeting, June 2014.

2 The $\bar{\text{P}}\text{ANDA}$ Experiment

2.1 The $\bar{\text{P}}\text{ANDA}$ Experiment

2.1.1 A worldwide unique facility for antiprotons

The ambition of $\bar{\text{P}}\text{ANDA}$ (anti-Proton ANnihilation at DArmstadt) collaboration [1] is to provide an exceptional and versatile subatomic physics laboratory complementary to the existing facilities. The unique feature of $\bar{\text{P}}\text{ANDA}$ is to exploit antiprotons as a hadronic probe to study the dynamics of the strong interaction. This in contrast to the currently operating hadron-physics facilities that primarily use electromagnetic probes or beams of protons and secondary pions or kaons.

In the past decades, major progress has been made in the field of strong QCD. In particular, recent discoveries of new forms of hadronic matter, e.g. tetraquark candidates in the heavy-quark sector (also known as XYZ-states), have given the field an enormous boost forward. Hadron physics has thereby a viable future with new challenges to overcome and questions to address in the upcoming decades. $\bar{\text{P}}\text{ANDA}$ has the potential to provide new insights in the field since it aims to harvest complementary and unique data exploiting a beam of antiprotons in the momentum range from 1.5 to 15 GeV/c annihilating with protons and nuclei.

Antiprotons in the $\bar{\text{P}}\text{ANDA}$ momentum range will probe both perturbative and non-perturbative aspects of QCD. This can be achieved by studying the structure and interaction dynamics of various hadrons with gluon-rich and light-, strange-, and charm degrees-of-freedom. The complementarity with respect to other probes is many-fold.

First of all, antiproton annihilations are associated with large production rates, in contrast to the use of electromagnetic probes. In particular, hadrons composed of (multiple) strange quarks and gluon degrees-of-freedom are abundantly produced as demonstrated at the predecessor LEAR.

Another unique feature of antiprotons is the potential to populate directly a large spectrum of spin-parity states. It allows to measure, with an unprecedented resolution, the line shape of (narrow) states and to access very high-spin states. Experiments that exploit electromagnetic probes, such as the ones using electron-positron collisions, and heavy

meson decays, i.e. B decays, cannot produce high-spin states. It is of outmost importance to provide complementary information of the recently discovered XYZ-states to pinpoint their true nature and to search for new forms of hadrons that cannot be seen by other experiments.

The challenge of $\bar{\text{P}}\text{ANDA}$ lies in identifying reaction channels that largely vary in production cross section, ranging from microbarns to picobarns. Later in this document, we present a study of a few benchmark channels that are exemplary of the rich physics program with the emphasis on channels with a large variation in cross section. Below, we briefly discuss the present the main physics pillars that we defined for the first two phases of $\bar{\text{P}}\text{ANDA}$.

2.1.2 Hadron spectroscopy

The first pillar is devoted to provide precision data for hadron spectroscopy with light to charm constituent quarks, and gluons. Concerning the light-quark and gluon sector, we plan to map out the glueball spectrum and to search for exotic forms of hybrids and meson-like or molecular states. Final states best suited for that include one or multiple ϕ , ω , η , J/ψ as well as D and D_s mesons since their production is OZI suppressed due to the large flavor component, and thus those channels have reduced conventional matter content.

The data we foresee to harvest are complementary to studies conducted at BESIII, COMPASS, and GLUE-X due to various reasons like the probe, the analysis technique and the accessible quantum numbers. The LEAR facility has demonstrated the strong advantage of using antiprotons for gluon-rich matter such as various scalar states that were first discovered using antiproton annihilations including the candidate for the glueball ground state, i.e. the $f_0(1500)$. $\bar{\text{P}}\text{ANDA}$ will extend these measurements by probing a sufficiently larger energy range and with a detector capable to perform a fully exclusive study of practically all final states. As demonstrated by LEAR, the cross sections are spectacularly high in contrast to reactions with electromagnetic probes. Already with lower luminosities, this gives excellent prospects for the search for massive glueballs with $J^{PC} = 2^{++}$ and 0^{-+} , hybrid states with gluonic degrees of freedom and exotic quantum numbers as well as meson-like states with light and

strange quarks. The analysis technique at \bar{P} ANDA is based on exclusive final states and constraints on the initial state for the ease of reliable quantum number assignments that are crucial for the interpretation of the data. This is superior to scattering experiment in terms of systematics and backgrounds, and can systematically be extended to the hidden and open-charm sector. This will e.g. shed new light on the recently discovered charmonium-like XYZ-spectrum with precision lineshape scans (with a resolution of about $50 \text{ keV}/c^2$ per point). In particular, it will allow for searches in the high spin segment (up to $J = 6$) or at larger mass (up to $5.5 \text{ GeV}/c^2$) to complement the existing findings such as $X(3872)$, $Z_c(3900)$ and $Z_c(4020)$ (to name a few) for which most of the decay channels are not found yet and counter parts for an understanding of the spectrum are missing.

One of the highlights within the XYZ program of \bar{P} ANDA is a rigorous study of the lineshape parameters of narrow resonances, such as the $X(3872)$, a $J^{PC} = 1^{++}$ hidden-charm state of exotic, but yet unknown, nature. Via a resonance scan exploiting the excellent momentum resolution of the antiproton beam, \bar{P} ANDA will be able to probe already at phase-one its narrow natural width, presently known to be less than 1.2 MeV , with a sensitivity better than 100 keV . Such a measurement will provide crucial information that will eventually help to shed light on the true nature of this state. One of the typical final states that will be used for the detection of the $X(3872)$ is its decay to the vector J/ψ charmonium state via the emission of a $\pi^+\pi^-$ pair. The J/ψ can be identified via its $\ell^+\ell^-$ decay. The production cross section of the $X(3872)$ in $\bar{p}p$ interactions will be less than 50 nb .

2.1.3 Hyperon physics

The second physics pillar of \bar{P} ANDA is the formation of mesons and baryons with open-strangeness and charmness in $\bar{p}p$ annihilations. This is truly unique in the world and will provide insight in the underlying mechanisms that play a role in the creation of strange and charm quark pairs. For the phase-one and phase-two programs of \bar{P} ANDA, we foresee to exploit the pair production of $|S|=1, 2$, and 3 and $|C|=1$ baryons and mesons near their production thresholds. It delivers a rigorous test for the validity of few-body models at various mass and energy scales. The self-analyzing feature of weakly-decaying hyperons gives experimental access to spin observables such as polarization and spin correlations, which are sensitive to the production mecha-

nism. A multidimensional analysis of the particle-antiparticle symmetric final state provides a model-independent test of CP violation in baryon decays. This has to be seen in the context of the long-standing puzzle of the matter-antimatter asymmetry of the Universe.

The associated production also will provide the opportunity to extensively investigate the spectrum of the charm and strange sector. This allows for a tagging method, an ideal tool for spectroscopy purposes near production thresholds where the missing-mass resolution is optimal and the background is small. Limitations in the luminosity are compensated by the huge cross sections (up to microbarns) that can be expected because of past experience, in particular for $|S|=1, 2$ hyperons. In general, the production cross sections of $|S|=2$ states with antiprotons are two- or three-orders in magnitude larger than that of photons. \bar{P} ANDA will be able to study the $|S|=3$ sector, where only a very few experiments can contribute measurements. Furthermore, the antiproton probe makes \bar{P} ANDA sensitive to states that do not couple to photons or kaons.

A typical example process that will be accessible in \bar{P} ANDA already in the first phases of the experiment, is the $\bar{p}p \rightarrow \Lambda\bar{\Lambda}$ hyperon-antihyperon channels its subsequent decay of the Λ ($\bar{\Lambda}$) into $p\pi^-$ ($\bar{p}\pi^+$). The production cross section at an antiproton beam momentum of $1.64 \text{ GeV}/c$ will be about $64 \mu\text{b}$ leading to 7 million reconstructed $\Lambda\bar{\Lambda}$ pairs per day at the phase-one luminosity. With such high rates, \bar{P} ANDA will be able to provide high precision cross section and polarization data to study the underlying production dynamics.

2.1.4 Proton structure

Proton structure experiments with electromagnetic final-states in antiproton-proton collisions are the third physics pillar of \bar{P} ANDA. In the first phases, the electric and magnetic form factors, $|G_E|$ and $|G_M|$, and their ratios at various q^2 values in $\bar{p}p \rightarrow \ell^+\ell^-$ will be measured. These time-like electromagnetic form-factor studies are a necessary complement to the space-like counterpart in lepton-scattering experiments. The foreseen measurements will extend the time-like data of electron-positron annihilation experiments, such as the ones conducted at BESIII, substantially in terms of energy range as well as accuracy. Moreover, \bar{P} ANDA will collect data that are sensitive to probe the unphysical regime, $4m_e^2 < q^2 < 4m_p^2$ by exploiting a final-state with an additional pion that brings the (anti)proton offshell.

With a $\bar{p}p \rightarrow \ell^+\ell^-$ cross section of 0.64 nb at a beam momentum of 1.5 GeV/c, the present day accuracy on the ratio $R=G_E/G_M$ will be improved from 20% at low q^2 values to 3%. Combining data from (yet unpublished) BESIII and \bar{P} ANDA, the analytical nature of the form factors can be studied. Besides these unique measurements, \bar{P} ANDA will be the first experiment capable to measure the proton time-like form factors in both e^+e^- and $\mu^+\mu^-$ final states. It might offer new insights in the radius puzzle of the proton by testing the lepton universality in these reactions.

2.1.5 Hadrons in nuclei

The fourth pillar of the phase-one and two program is dedicated to study the properties of hadrons in a nuclear medium. Firstly, we propose studies of the antihyperon-nucleus potential, accessible via the associated production of hyperon pairs close to the production threshold, which cannot be studied at any other laboratory in the world. It offers essential information for the interpretation of data from heavy-ion collision experiments. Furthermore, the foreseen program includes measurements of the basic (mass, width) parameters of hidden-charm states at nuclear densities. This might give complementary information that helps to shed light on the formation of hadrons and measures signatures that could point to the (partial) restoration of chiral symmetry. Also, in the context of color transparency we can deliver significant contributions that exceed those of other experiments, e.g. due to the access to other final states due to the $\bar{p}N$ entrance channel.

Furthermore antiprotons can be used to implant hyperon pairs into a nucleus. \bar{P} ANDA offers also the unique possibility to search for X-ray transitions from very heavy hyperatoms as e.g. Ξ^- - ^{208}Pb which is not possible at other labs. This will complement experiments at J-PARC which attempt to measure X-rays in medium-heavy nuclei. The measurement at \bar{P} ANDA will for the first time allow to constrain the interaction of Ξ^- -hyperons in the neutron skin. In a later stage, \bar{P} ANDA will extend the studies on double hypernuclei by performing for the first time high resolution γ -spectroscopy of these nuclei.

2.1.6 High-Energy Storage Ring

The preparation of an anti-proton beam for the \bar{P} ANDA experiment needs a complex sequence of different accelerators: First protons have to be ac-

celerated up to 29 GeV in a sequence of three different accelerators: first a proton linac, than SIS18 and finally SIS100. The 29 GeV protons are then shot onto an anti-proton production target where, among other particles, antiprotons are produced. These are collected and pre-cooled in the Collector Ring (CR) for about 10 s.

The High-Energy Storage Ring (HESR) and the \bar{P} ANDA experiment at the FAIR facility are shown in Fig. 2.1. The experiment aims at both high reaction rates and high resolution in order to study rare production processes and small branching ratios. With a design value of 10^{11} stored antiprotons for beam momenta from 1.5 GeV/c to 15 GeV/c and high density internal targets the anticipated antiproton production rate of $2 \cdot 10^7 \text{ s}^{-1}$ governs the experiment interaction rate in the order of cycle-averaged $1 \cdot 10^7 \text{ s}^{-1}$. The stored antiprotons are freely coasting except for a 10% to 20% bunch structure allocated to a barrier bucket for compensation of energy losses.

Two complementary operating modes are planned, named *high luminosity* and *high resolution* mode, respectively. The high luminosity mode with $\Delta p/p = 10^{-4}$, stochastic cooling and a target thickness of $4 \cdot 10^{15} \text{ cm}^{-2}$ will have an average luminosity of up to $L = 2 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. For the high resolution mode $\Delta p/p = 5 \cdot 10^{-5}$ will be achieved with electron cooling for momenta up to $p = 8.9 \text{ GeV}/c$. Operation will mainly be in conjunction with a cluster jet target which will not impose a time structure onto the event rate. The cycle-averaged luminosity is expected to be $L = 2 \cdot 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$.

In phase one of the FAIR facility every 10 s 10^8 anti-protons are transferred to HESR for 1000 s until 10^{10} anti-protons are accumulated in the accelerator. In the HESR they are further cooled down to a momentum resolution better than $5 \cdot 10^{-5}$ and accelerated to the wanted beam momentum in a range of 1.5 to 15.0 GeV/c.

2.1.7 Targets

The \bar{P} ANDA Target Spectrometer is designed to allow the installation of different targets. For hydrogen as target material both a Cluster-Jet target and a Pellet target are being prepared. One technical challenge is the distance of 2.1 m between the injection nozzle and the Interaction point, plus the same distance until the target particles are dumped in an efficient catcher keeping the whole target line under high vacuum.

The Cluster-Jet target is homogenous in space and




Figure 2.1: Schematic of the future FAIR layout incorporating the current GSI facilities on the left; on the right the future installations, the SIS 100 synchrotron the storage and cooler ring complex including CR and HESR and the Super FRS experiment being some of the new parts. \bar{P} ANDA is positioned right in the center of the image inside the HESR.

time whereas a Pellet target with average inter-pellet spacing of 3 mm exhibits large density variations on the 10–100 μ s timescale.

An extension of the targets to heavier gases such as deuterium, nitrogen, or argon is planned for complementary studies with nuclear targets.

2.1.8 Luminosity Considerations

The luminosity is directly linked to the number of stored antiprotons. The maximum luminosity depends on the antiproton production rate. The cycle-averaged antiproton production rate and reaction rate must be equal in the consumption limit. Due to injection losses and possible dumping of beam particles at the end of a cycle the time-averaged reaction rate will be lower. In Figure 2.2 the beam preparation periods with target off and data taking periods with target on are drawn. The red curve showing the luminosity at constant target thickness is proportional to the decreasing number of antiprotons during data taking. In order to provide a constant luminosity, compensation by adjusting the target thickness is studied.

In the case of a Pellet target, variations of the in-




Figure 2.2: [2] Time dependent macroscopic luminosity profile $L(t)$ in one operation cycle for constant (solid red) and increasing (green dotted) target density ρ_{target} . Different measures for beam preparation are indicated. Pre-cooling is performed at 3.8 GeV/c. A maximum ramp of 25 mT/s is specified for acceleration and deceleration of the beam.

stantaneous luminosity will occur. These are depending on the antiproton beam profile, pellet size, pellet trajectories and the spacing between pellets. In the case of an uncontrolled pellet sequence target thickness fluctuations with up to 2–3 pellets in beam do occur during timescales of the pellet transit time which is 10–100 μ s. The pellet high lumi-

osity mode (PHL mode) features smaller droplet size, lower spread in pellet relative velocity and smaller average pellet distance. The latter being much smaller than the beam size, hence the thickness fluctuations would be much reduced. However, this pellet target mode is currently being developed.

2.2 The \bar{P} ANDA Detector

Figure 2.3 shows the \bar{P} ANDA detector viewed with partial cut-outs. As a fixed target experiment, it is asymmetric having two parts, the Target Spectrometer (TS) and the Forward Spectrometer (FS). The antiproton beam is scattered off a pellet or cluster-jet target (left side in Fig. 2.3). \bar{P} ANDA will measure $\bar{p}p$ reactions comprehensively and exclusively, which requires simultaneous measurements of leptons and photons as well as charged and neutral hadrons, with high multiplicities.

The physics requirements for the detectors are:

- to cover the full solid angle of the final state particles,
- to measure energy and momenta of the reaction products, and
- to identify particle types over the full range of momenta of the reaction products.

2.2.1 Target Spectrometer

Figure 2.4 shows a side view of the \bar{P} ANDA target spectrometer. The TS, which is almost hermetically sealed to avoid solid angle gaps and which provides little spare space inside, consists of a superconducting solenoid magnet with a field of 2 T and a set of detectors for the energy determination of neutral and charged particles as well as for the tracking and PID for charged tracks housed within the magnet. The silicon microvertex detector (MVD) closely abuts the beam pipe surrounding the target area and provides secondary vertex sensitivity for particles with decay lengths on the order of 100 μm .

The main tracker is a straw tube tracker (STT). There will be several gas electron multiplier (GEM) tracking stations in the forward direction. The tracking detectors like MVD and STT also provide information on the specific energy loss in their data stream.

Two Internally Reflected Cherenkov light (DIRC) detectors are to be located within the TS. Compared to other types of Ring Imaging Cherenkov

(RICH) counters the possibility of using thin radiators and placing the readout elements outside the acceptance favors the use of DIRC designs as Cherenkov imaging detectors for PID. The Barrel DIRC covers the polar angles θ from 22° to 140° inside the \bar{P} ANDA TS. The Endcap Disc DIRC (EDD) covers the polar angles θ from 10° to 22° in the horizontal plane and 5° to 22° in the vertical plane. For the analysis of the DIRC data the tracking information is needed, as the Cherenkov angle is measured between the Cherenkov photon direction and the momentum vector of the radiating particle.

The Barrel TOF serves as precise (< 100 ps) timing detector cylindrically surrounding the target. It consists of small scintillator tiles read out by Silicon Photomultipliers (SiPMs) and is attached to the support frame outside of the Barrel DIRC providing reasonable π -K separation below 1 GeV/c.

The lead tungstate (PWO) crystals of the electromagnetic calorimeters (EMC) are read out with Avalanche Photo Diodes (APD) or vacuum pentodes. Both, the light output and the APD performance improve with lower temperature. Thus the plan is, to operate the EMC detectors at $T = -25^\circ\text{C}$. The EMC is subdivided into backward endcap, barrel and forward endcap, all housed within the solenoid magnet return yoke.

Besides the detection of photons, the EMC is also the most powerful detector for the identification of electrons. The identification and measurement of this particle species will play an essential role for the physics program of \bar{P} ANDA.

The return yoke for the solenoid magnet in the \bar{P} ANDA TS is laminated to accommodate layers of drift tubes (Iarocci-type detectors) for the muon detection. They form a range stack, with the inner muon layer being able to detect low energy muons and the cumulated iron layer thickness in front of the outer layers providing enough hadronic material to stop the high energy pions produced in \bar{P} ANDA. A similar lamination and instrumentation of the iron is foreseen in the downstream door of the yoke augmented by the addition of a muon filter located in between the TS and the FS.

2.2.2 Forward Spectrometer

Figure 2.5 shows a side view of the \bar{P} ANDA forward spectrometer. The FS angular acceptance has an ellipsoidal form with a maximum angular acceptance of ± 10 degrees horizontally and ± 5 degrees vertically w.r.t. the beam direction.

The tracking section of the FS is incorporated into

the large gap of a dipole magnet providing bending power of 2 Tm with a B-field perpendicular to the forward tracks. The other parts are placed further downstream outside the dipole magnet.

- 5 An aerogel RICH detector will be located right behind the dipole magnet followed by a the Forward Time-of-Flight wall (FTOF) which also covers the detection of slow particles below the Cherenkov light threshold. The energy is measured in the
 10 Shashlyk type electromagnetic calorimeter consisting of 1404 modules of $55 \times 55 \text{ mm}^2$ cell size covering $2.97 \times 1.43 \text{ m}^2$.

For the determination of the luminosity a detector based on four layers of monolithic active pixel
 15 sensors will be installed close to the beam pipe detecting scattered antiprotons under small angles.

2.2.3 The Particle Identification System

The charged particle identification (PID) will combine the information from the time-of-flight, tracking, dE/dx and calorimetry information with the output from the Cherenkov detectors, with their focus on positive identification of kaons.

The individual $\bar{\text{PANDA}}$ subsystems contributing to the PID and the combination of their data into a global PID information have been reviewed in the PID-TAG-report [3], a performance plot regarding the π -K separation power is shown in chapter ??.

2.2.4 Infrastructure

30 The $\bar{\text{PANDA}}$ detector is located below ground in an experimental hall, encased in smaller tunnel-like concrete structure, partially fixed, partially made of removable blocks. Most subsystems connect their FEE-components via cables and tubes placed
 35 in movable cable ducts to the installations in the counting house, where three levels are foreseen to accommodate cooling, gas supplies, power supplies, electronics, and work space.

Bibliography

- 40 [1] PANDA Collaboration. Technical Progress Report, FAIR-ESAC/Pbar. 2005.
- [2] PANDA Collaboration. Straw Tube Tracker Technical Design Report. 2012.
- 45 [3] G. Schepers et al. Particle Identification at PANDA. *Report of the PID TAG*, March 2009.




Figure 2.3: Aerial view of PANDA with the Target Spectrometer (TS) on the left side, and the Forward Spectrometer (FS) starting with the dipole magnet on the right. The antiproton beam enters from the left.




Figure 2.4: Side view of \bar{P} ANDA with the Target Spectrometer (TS). The antiproton beam enters from the left.




Figure 2.5: Side view of \bar{P} ANDA forward Spectrometer (FS). The antiproton beam enters from the left.

3 DAQ requirements

Section 3.1 describes the expected event rates and discusses the pileup situation for the first stage of operations, the so called phase 1 physics. Section 3.2 describes readout of the PANDA subsystems. Section 3.3 discusses the available capacity for on-line storage and the resulting requirements on event filtering. Section 3.4 discusses the partitioning of DAQ and DAQ running modes.

3.1 Event rates and pile-up situation for Phase 1 Physics

3.1.1 Event Rates

In the phase-one operation mode, see 2.1.6, the maximum luminosity is limited to $2 \cdot 10^{31} \text{cm}^{-2} \text{s}^{-1}$. In the final version of FAIR an additional accumulator ring is foreseen (RESR) which will increase the achievable luminosity by a factor of 10 (phase three).

The cross section for the reaction of an anti-proton beam with a proton target is momentum dependent and drops with rising momentum in the energy range of the PANDA experiment (see Fig. 3.1). It lasts from about 100 mbarn at the lowest beam momentum of 1.5 GeV/c and drops to about 60 mbarn at the highest beam momentum of 15 GeV/c. For the calculation of the interaction rate the highest value of 100 mbarn was chosen. Multiplied with the maximum luminosity of $2 \cdot 10^{31} \text{cm}^{-2} \text{s}^{-1}$ one gets a mean interaction rate of 2 MHz or a mean time between events of 500 ns. This interaction rate we use as design requirement for the current TDR (phases one and two).

3.1.2 Event Overlap

The beam structure inside HESR has a very unique pattern. Four fifth of the accelerator ring are filled with anti-protons while one fifth is empty. With a duration of about $2 \mu\text{s}$ for one revolution one gets $1.6 \mu\text{s}$ of quasi continuous beam followed by a time gap of 400 ns at the target station of PANDA. The interaction probability of the anti-proton beam with the target is described by a Poisson distribution leading to a probability distribution for the time between two consecutive interactions described by




Figure 3.1: Measured antiproton-proton total (black) and elastic (red) cross section and calculated as the difference between the two the inelastic (blue) cross section in the momentum range of the PANDA experiment




Figure 3.2: Simulated time between two consecutive anti-proton target interactions with a mean time between events of 500 ns, a continuous beam with 1600 ns and a gap of 400 ns

an exponential drop with a high probability for very short time gaps and a long tale for very long time gaps (see Fig. 3.2).

Even though the mean time between two events is with 500 ns quite large the high probability for shorter times between events increases the probability that two or more events overlap in the data stream of a sub-detector. How big this effect is depends on the time resolution of the detector for the initial interaction time (t_0). The time resolution of the t_0 time depends on the uncertainty of the flight time of the particle until it hits the sub-detector




Figure 3.3: Time difference between two consecutive hits in the MVD for 1000 simulated background events




Figure 3.4: Time difference between the last hit of an event and the first hit of the next event in the MVD

and the time resolution of the detector itself.

If one takes a relatively fast detector close to the interaction point like the MVD with an internal time resolution of better than 10 ns 2 % of the events are overlapping with a neighboring event. The STT on the other hand has an overlap of about 18 % due to its additional uncertainty coming from the drift time of 250 ns in the drift tubes.

3.1.3 Online Event Building

A fast way to perform initial event building is to look at the time-sorted data stream coming from the sub-detectors. If the time between two events is larger than the time difference between hits inside an event it is possible to separate events just by grouping those hits which are close together in time. To test how good this would work 1000 background events have been simulated with the FTF background simulator and studied in the Panda-Root simulation framework. Figure 3.3 shows the time between two hits of the Micro Vertex Detector (MVD) which are in the same event. In this logarithmic plot one can see that most hits have a time difference less than 1 ns and none has a time difference longer than 19 ns. If the time gap between two consecutive hits is longer than 20 ns the probability is high that they are not from the same event.

Figure 3.4 shows the time between the last hit of one event and the first hit of the next event. This time gap can be several μ s. The integral of this distribution normalized to the total number of events is shown in Figure 3.5. With a time gap of 20 ns one would have about 5 % of events which are merged together, while 95 % could be separated from each other with this simple method.




Figure 3.5: Integral of the time difference between the last hit of one event and the first hit of the next event as percentage value for all events.

For detectors with worse t_0 time resolution like the STT with its drift time up to 250 ns, one still gets 60 % clean events while the other 40 % contain more than one MC event.

The coarse event candidates generated by looking at the time gaps between two consecutive hits will then be used as input for the event reconstruction algorithms like tracking and pid. By combining the information of different sub detectors a t_0 time resolution much better than 1 ns can be achieved which is sufficient to disentangle overlapping events and separate events for the later physics analysis.

With the low luminosity at the beginning of the PANDA experiment the event overlap is quite small and does not effect the performance of the PANDA detector. With the higher luminosity once RESR is available this situation changes and the reconstruction algorithms have to take event overlap into account to separate single events from each other.

This TDR focuses on the phase one of the PANDA experiment with the maximum event-rate of 2 MHz, while readout capable of higher rates of the phase three is subject to a latter upgrade.

3.2 Readout of PANDA subsystems

This section provides an overview of all PANDA subsystems in terms of number of readout channels and output data rates. All PANDA subsystems and corresponding front-end electronics are designed for operation with the highest possible interaction rate of 20 MHz. Therefore, in this section data-rate and other characteristics are presented for the high interaction-rate which is relevant for the phase three. Numbers relevant for the phase one are derived by down-scaling by factor 10 corresponding values.

3.2.1 Micro Vertex Detector

Subsystem structure

The Micro Vertex Detector (MVD) consists of two silicon devices: hybrid silicon pixels and double sided silicon micro-strips. Four barrels, two innermost composed of hybrid pixels and two outermost composed of micro-strips, are located around the interaction point. Six disks based on hybrid pixels and two micro-strips rings compose the forward part of the MVD.

The readout electronics for both pixels and micro-strips is custom made due to the challenging request of a trigger-less readout.

The readout of the pixels is based on the chip named ToPix, bump bonded to the sensor. A ToPix chip implements a 12760 pixel readout channel matrix and includes two drivers which transmit data at 320 Mb/s each. The pixel module layout is based on a basic unit corresponding to a readout chip size and modules of several sizes are obtained by tiling from two to six units. In total, 176 sensors and 810 readout chips equip the pixel detector (10335600 readout channels): 64 sensors and 338 chips in the barrels (4312880 readout channels), 112 sensors and 472 chips (6022720) in the disks.

The readout ASIC of the strip part is named ToASt and implements 64 readout channels. Two 160 Mb/s links will transmit the data. The strip part foresees 3 different geometries: square sensor, read out by 8 ASICs each, and rectangular ones, read

out by 11 ASICs each, equip the two barrels and trapezoidal sensors, read out by 12 ASICs each, compose the two rings. In total, 296 sensors and 3112 chips equip the micro-strips detector (199168 readout channels): 64 square and 184 rectangular sensors and 2536 chips in the barrels (162304 readout channels), 48 trapezoidal sensors and 576 chips in the rings (36864 channels).

Readout architecture and data transmission for the pixel detector

A dedicated radiation hard optoelectronic board, composed of a GBT chip-set and a Versatile link (generically indicated as GBT), both developed at CERN for the LHC projects, receives the electrical data transmission lines and transmits the combined data via fast optical links. The 320 Mb/s serial links, based on the SLVS (Scalable Low Voltage Signaling) differential standard, in order to be compatible with the 1.2 V voltage supply typical deep sub-micron technologies, interface the GBT to the front-end electronics. 68 GBT for the barrels and 95 GBT for the disks will be arranged upstream the MVD.

Readout architecture and data transmission for the micro-strips detector

The serial output links of the ToASt front-end ASICs for the readout of one sensor are multiplexed by one Module Data Concentrator (MDC). A total number of 296 MDCs is foreseen, and each one provides one link (foreseen at 320 Mb/s) to transmit data via differential SLVS lanes to the eLink interface port of the GBT chip-set. A number of 30 GBTs thus will route the entire strip system and 30 optical links will connect the off-detector electronics.

Expected hit rates

The expected particle distribution in antiproton-proton collisions obtained with 10^7 Dual Parton Model (DPM) events at two different beam momenta (p_{beam}) are reported in figure 3.6. This distribution is mainly peaked in the forward direction, where the pixel disks and the strip rings are positioned.

The expected particle distribution in antiproton-nucleus collisions obtained at maximum beam momenta (p_{beam}) are mainly distributed in the angular range covered by the barrel part of the MVD as presented in figure 3.7.

The rates over all individual channels of the micro-




Figure 3.6: In the top of the figure, particle distributions are reported as a function of particle momentum and polar angle. In the bottom part of the figure the 1D profile is projected onto the polar angle. The red line results from the elastic scattering process that is implemented in the DPM generator (MVD TDR).

strips and the pixel front-end ASICs were studied and they are reported in figure 3.8.

Considering the charge sharing effect and the output format of the hit data, a maximum raw data rate of roughly 40 Gbit/s is expected at highest momentum with the nominal interaction rate of 2×10^7 s⁻¹, 10 times higher than in phase one, at which the contribution of the pixel and strip part are in the same order of magnitude. The expected data rate for the phase one is 4 Gbit/s.

3.2.2 Straw-Tube Tracker

System Characterization

The STT system is mechanically split into two independent semi-barrels. One semi-barrel is mounted on the left and the other on the right side of the Central Systems Frame. The electronic readout chain from the detector front-end electronics, cable routing and readout system back-end in front of the PANDA TS is divided into two corresponding sub-systems.

The electronic readout records straw signal time (leading edge) and signal time-over-threshold (TOT). The TOT measurement gives information about the particle specific energy-loss (dE/dx)




Figure 3.7: In the top of the figure, particle distributions, obtained with 4×10^4 UrQMD events, are reported as a function of particle momentum and polar angle for antiproton-N and antiproton-Au collisions. In the bottom part of the figure, the 1D profile is projected onto the polar angle, obtained with a statistical sample of 2×10^6 (MVD TDR).

and is used for particle identification (PID) in the momentum range below about 1 GeV/c.

A PANDA-straw specific ASIC, named PAST-TREC (= Panda STTracker REadout Chip), has been developed for the signal time and TOT readout. The 8-channel chip features a software controlled range of electronic parameters for each channel, like amplifier gain, peaking time, baseline level, threshold and further signal shaping parameters which can be selected by DAC. The 16-channel front-end boards are equipped with two PAST-TREC ASICs per board.

Readout overview

The STT subsystem has the following structure:

- 4224 straws in 21-27 radial layers, arranged in 2 subsystems (2x 2112 straws)
- 270x front-end boards (FEB)




Figure 3.8: Extracted average count rates for antiproton-proton reactions obtained with the DPM generator. The distribution over all individual channels of the micro-strips and the pixel front-end with the highest occupancy are shown in the frames on the left at top and bottom respectively. Arrows indicate the position inside the MVD. Results for all readout chips (FE) in the different detector layers are shown on the right (MVD TDR).

- few FEBs with not all channels connected to straws
- 16 readout channels per FEB
- Front-end PASTTREC ASIC
 - Straw signal LE-time and TE-time measurement
 - ASIC with amplifier, shaper and discriminator circuitry
 - 8x channel per ASIC, 2x ASIC per FEB
- SW control of selectable DAC range per ASIC (gain, peaktime, threshold, shaping parameters), BL level selectable per channel
- LVDS signal output of discriminated signal leading and trailing edge
- Micro-twisted-pair cable (2x20 lines) for LVDS straw signals and ASIC control

The Monte Carlo simulations were used to estimate expected hit rates and corresponding data rates, see 10




Figure 3.9: Simulation of P-bar p reactions giving the number of hits per event and per cm along the straw tubes in the innermost layer of the PANDA STT. The target position is at $z=0$ cm. (Figure from PANDA STT TDR).




Figure 3.10: The readout structure of the barrel DIRC subsystem using TRB3 boards.

figure 3.9. Deduced hit rates for the full luminosity mode ($2 \times 10^7 \text{ s}^{-1}$ p-bar p, PANDA phase three) are:

- ca. $8 \times 10^5 \text{ s}^{-1}$ per straw in innermost layer (max. intensities);
- ca. $5 \times 10^5 \text{ s}^{-1}$ per straw on average all layers.

Taking into account that the data word size per hit is 10 bytes the total expected raw data rate is about 160 Gbit/s.

3.2.3 Barrel DIRC

The subsystem consists of 16 units with 8 photon detectors each. Each photon detector has 64 pixels. In total we have $16 \times 8 \times 64 = 8192$ pixels = 8192 readout channels, see figure 3.10. At 20 MHz interaction rate, phase 3, it is expected to register on average two charged particles in the barrel. For each

registered charged particle there are about 50 detected Cherenkov photons, therefore, the expected hit rate is about 2×10^9 hits/s. 96 bytes are required to record one hit, namely 32 bit for the leading edge, 32 bit – for the trailing edge, and 32 bit overhead. Therefore, a raw data rate of 200 Gbit/s is expected for the PANDA phase 3.

3.2.4 Endcap Disc DIRC

The Endcap Disk DIRC (EDD) consists of four independent quadrants and each contains 24 Read-Out Modules (ROMs). Each ROM consists of one MCP-PMT with 300 (PHOTONIS) or 384 readout channels (Hamamatsu) and corresponding optical components. Each ROM will be readout by 6 TOF-PET ASICs (64 channels per ASIC). In total the EDD will have 7200 (9216) channels per quadrant in case of PHOTONIS (Hamamatsu) MCP-PMTs or 28800 (36864) channels for the complete EDD. Assuming the high luminosity mode of operation for the PANDA experiment the maximum expected hit rate is about 100 kHz/channel (for PHOTONIS), which results in about 30 MHz/ROM (for PHOTONIS or Hamamatsu). All in all the total hit rate for the high luminosity corresponds to 7.2×10^8 per quadrant and 2.9×10^9 for the complete detector. Taking into account that one hit corresponds to 32 bits the maximum expected data rate is about 88 Gbit/s for the phase 3.

3.2.5 Barrel TOF

The barrel TOF detector consists of small scintillator tiles read out by Silicon Photomultipliers (SiPM). The whole detector is sectioned to 16 Supermodules (SM) of $2 \times 60 = 120$ tiles per each SM. Each tile is read out by 8 or 12 SiPMs combined to two readout channels yielding $120 \times 2 = 240$ readout channels per SM and $240 \times 16 = 3840$ readout channels in total. According to Monte Carlo simulations average hit rate at 20MHz interaction rate (high luminosity) is 27 kHz per readout channel. Therefore, expected data rate for the high luminosity is 6.9 Gbit/s assuming that each hit produces 64 bit of data.

3.2.6 Muon detectors

The laminated yoke of the solenoid magnet acts as a range system for the detection of muons. There are 13 sensitive layers, each 3 cm thick (layer “zero” is a double-layer). They alternate with 3 cm thick iron absorber layers (first and last iron layers are 6

cm thick), introducing enough material for the absorption of pions in the PANDA momentum range and angles. In the forward End Cap more material is needed due to the higher momenta of the occurring particles. Therefore, six detection layers will be placed around five iron layers of 6 cm each within the downstream door of the return yoke, and a removable Muon Filter with additional four layers of 6 cm iron and corresponding detection layers will be moved in the space between the solenoid and the dipole.

As sensors between the absorber layers, rectangular aluminium Mini Drift Tubes (MDT) are foreseen. Basically, these are drift tubes with additional capacitive coupled strips. All together, the laminated yoke of the solenoid magnet and the additional Muon Filter will be instrumented with 2751 MDTs and 424 MDTs, respectively. The MDTs are read out by the Trans-Resistance Amplifier ASIC Ampl-8.3 coupled with the Comparator ASIC Disc-8.3 and corresponding digital readout boards. Assuming the high-luminosity mode (phase 3) the muon detector will produce 19 Gbps of raw data transmitted over 62 optical fibres to the DC level.

3.2.7 Electromagnetic Calorimeter

Subsystem structure

The EMC consists of a target and a forward calorimeters. The target EMC is a homogeneous calorimeter composed of three sub-detectors: backward end-cap (Bw end-cap), barrel and forward end-cap (Fw end-cap), and is made of PWO crystals:

- Bw end-cap:
 - 524 PWO crystals read out by two LAAPDs each, one readout channel per LAAPD
 - $524 \times 2 = 1048$ readout channels
- Barrel:
 - 11360 PWO crystals read out by two LAAPDs each, two readout channels per LAAPD (high and low amplification gain)
 - $11360 \times 4 = 45440$ readout channels
- Fw end-cap:
 - 3856 PWO crystals read out by LAAPDs (two readout channel per LAAPD), and VPTT (one readout channel per VPTT)




Figure 3.11: Hit rate in the target EMC from the DPM background generator at $p\bar{p} = 14$ GeV/c. Top panel – barrel EMC, bottom panel – forward end-cap EMC. [EMC TDR, p 36.]

$$- 768 \text{ (VPTT)} + 6176 \text{ (LAAPD)} \times 2 = 13888 \text{ readout channels}$$

The forward EMC (FSC) is a sampling calorimeter read out by about 1500 photo-sensors with one readout channel each.

Expected hit-rate

Since the PANDA experiment employs fixed-target geometry the hit-rate distribution is non-uniform and peaking in the forward direction. The average hit-rate per crystal for the highest beam energy and luminosity is shown in fig. 3.11. In order to simplify the estimation of the data flow we divide the Fw end-cap into two regions, namely, the central high-rate region read out by VPTTs (768 crystals) with a hit-rate of 500 kHz, and the periphery region with a hit-rate of 100 kHz (3088 crystals). The barrel EMC is split into three regions as shown in tab. 3.1. The hit-rate for the Bw end-cap we estimate as 10 kHz (524 crystals).

Front-end electronics

To read out photo-sensors two different front-end types, readout units (RU), will be employed, namely digitizers based on a sampling ADC (SADC) and hit-detection ASIC. The SADC-based digitizer reads out 32 photo-sensors with dual-gain amplification in order to increase dynamic range and will be employed for the readout of the Fw end-cap. The digitizer has one radiation-hard 4.8 Gbps optical interface for communication with a data concentrator module (DC). There is the possibility to combine data from multiple digitizers using the backplane of the crate. However, due to the required redundancy and fail-save operation in the radiation environment the data from different digitizers will not be combined, even though the bandwidth of the link will be not completely employed. Therefore, number of optical interfaces to the DC layer is completely defined by the number of photo-sensors in the Fw end-cap, namely 768 VPTTs read out by 24 digitizers and 6176 LAAPDs – by 193 digitizers. The FSC is read out by 24 64-channel digitizers. The number of optical interfaces together with corresponding data-rates is summarized in the tab. 3.2. The data-rate estimation is based on the maximum hit-rates per sub-detector parts and assumes data size of 16 bytes per hit, which is defined by the communication protocol between the RU and DC.

The barrel and Bw end-cap EMC are employing RU based on a hit-detection ASIC. The ASIC reads out four LAAPDs and has four digital links to the RU. Depending on the hit-rate the hit-detection ASIC uses either two or four links to the RU. Four links are employed only in the most-forward region of the barrel EMC (1120 crystals, see tab. 3.1). The RU has 106 links to communicate with hit-detection ASICs and one radiation-hard 4.8 Gbps optical interface to DC. Therefore, for the most forward part of the barrel EMC one RU will read out 26 crystals (52 LAAPDs) and in all other cases 53 crystals (106 LAAPDs). The number of optical interfaces together with corresponding data-rates is summarized in the tab. 3.2. Even though for some RU

Table 3.1: Number of crystals with expected hit-rate for the barrel EMC.

Hit rate	# of crystals
< 10 kHz	$28 \times 160 = 4480$
< 50 kHz	$24 \times 160 = 3840$
< 90 kHz	$19 \times 160 = 3040$

estimated data rate is much lower than available bandwidth of output link data of different RU will not be combined at the detector to avoid substantial data loss in case of malfunction of a RU.

As can be seen from tab. 3.2 the total number of optical links towards DC level is 526 with a combined maximum data-rate of $(254 + 8)$ Gbit/s for target and FSC calorimeters. One has to mention that at the latter stage amount of the data for the target calorimeter will be reduced by the factor two due to combining data from two photo sensors reading out the same EMC crystal. Therefore, the total maximum data-rate for the phase three to the DAQ is $254/2 + 8 = 135$ Gbit/s.

3.2.8 Forward Tracker

High Resolution mode data rates

The estimated data rate is presented taking into account the existing firmware for the TDC and the Data Board (DB) that define the data format, buffering and throughput capabilities of the TRBv3 platform. One registered hit in the TDC (leading and trailing edge of a signal for TOT) is represented by two 32-bit words, which include the channel number and a value containing encoded time measured in the range up to $10.24 \mu\text{s}$. Additionally, after each $10.24 \mu\text{s}$ period, another 32-bit word containing epoch counter is added on individual channels. In a free-running measurement system (time counters are not initialized at the beginning of the measurement window), it is assumed that the epoch counter words will appear on each active channel with a frequency close to 100 kHz.

The buffering is performed on several data processing stages. Each TDC channel can store up to 40 hits between consecutive readouts. Output buffer from one TDC module can store up to 2048 hits from all the channels. Assuming uniform hit distribution over all channels, the TRBv3 is capable to process up to 55 kHits/s on all 192 channels. When needed, one can limit the amount of hits stored per channel in order to gain buffer space for additional epoch counters (for lower readout frequencies) or higher channel density at the cost of lost hits.


In the High Resolution mode the average load on channels is below 55 kHits/s (see Table 3.3). Straws close to the beam pipe can have a much higher hit rate, reaching 140 kHits/s. Such effect can be compensated by mixing straws from various detector regions on one TRBv3 and therefore achieve balanced, aggregated data volumes per DB. A setup with 69 TRBv3 boards with 192 channels each is

Table 3.2: Summary of the EMC readout.

	photo sensors	# of RU	hit-rate/RU hits/sec.	max data rate/RU Gbps
Fw end-cap	768 VPTTs	24	$1.6 \cdot 10^7$	1.9
	6176 LAAPDs	193	$3.2 \cdot 10^6$	0.4
Barrel	3040×2 LAAPDs	117	$5.2 \cdot 10^6$	0.6
	3840×2 LAAPDs	73	$5.3 \cdot 10^6$	0.6
	4480×2 LAAPDs	85	$1.1 \cdot 10^6$	0.1
Bw end-cap	524×2 LAAPDs	10	$1.1 \cdot 10^6$	0.1
FSC	1512	24	$1.1 \cdot 10^6$	0.3

Table 3.3: Estimated data rates per tracking station in High Resolution mode.

Tracking station	Number of channels	Average hit rate	Maximum hit rate	Number of TRBv3s	Bandwidth	Total bandwidth
FT 1,2	2304	35 kHits/s	140 kHits/s	12	81 MB/s	972 MB/s
FT 3,4	3328	31 kHits/s	90 kHits/s	18	69 MB/s	1242 MB/s
FT 5,6	6592	9 kHits/s	39 kHits/s	35	20 MB/s	700 MB/s

**Figure 3.12:** Readout architecture for the forward tracker in the High Resolution mode.

needed to read all 12224 channels of the FT (see figure 3.12). 65 TRBv3 provide enough channels for the time measurement and are called slaves. Four additional boards are required to construct a tree-like architecture. One module, called master is a root and controls the readout process of the entire system. It also acts as a gateway for Slow-Control commands. The remaining 3 modules provide connectivity and synchronization of the slaves. The setup requires 13 SFP Addons, which are used to connect the boards between themselves.

The expected hit rate differs between FT stations, but so does the number of straws. Stations FT1 and FT2 are constructed out of 1152 straws each and expect 35 kHits/s in the High Resolution mode. For such number of channels 12 TRBv3s are required, each generating 81 MBps summing up to total of 0.927 GBps. Stations FT3 and FT4 have

3328 channels with which a setup of 18 TRBv3s has to deal with, generating 1.242 GBps data stream. Stations FT5 and FT6 have higher channel number (6592) but significantly lower hit rate. 35 TRBv3s are needed and a 0.7 GBps data stream will be generated. All those values are presented in table 3.3 and sum up to a total of 2.914 GBps out of the system in the High Resolution mode.


High Luminosity mode data rates

In the High Luminosity mode, the expected hit rates per channel are ten times higher than in the High Resolution mode. The readout system described above is not capable to cope with such data volumes and, therefore, development of a new hardware platform is planned. The concept is to develop a new version of the TRB board (conventionally named TRB5G), equipped with more advanced and up-to-date FPGA devices, mountable in a crate system that includes backplane with power supply and communication features. It is an evolution of the already existing TRB3SC system based on ECP3 FPGA, the same as in TRBv3 platform. Resources available in the proposed ECP5G-45 FPGA devices allow to implement up to 64 low resolution (below 0.5 ns) TDC channels sufficient for readout of 4 FEE modules.

A new implementation of the time measurement component allows also to redefine its output data format. One 32 bit word is enough to combine channel number, leading time (relative to the burst

Table 3.4: Estimated data rates per tracking station in High Luminosity mode.

Tracking station	Number of channels	Average hit rate	Maximum hit rate	Number of TRB5Gs	Bandwidth	Total bandwidth
FT 1,2	2304	350 kHits/s	1400 kHits/s	36	90 MB/s	3240 MB/s
FT 3,4	3328	310 kHits/s	900 kHits/s	52	79 MB/s	4108 MB/s
FT 5,6	6592	90 kHits/s	390 kHits/s	103	23 MB/s	2369 MB/s

**Figure 3.13:** Readout architecture for the forward tracker in the High Luminosity mode.

update signal) and its width for the TOT measurement. It is a significant (more than factor 2) reduction of the data volume as compared to the system for the High Resolution mode.

One crate can house up to 9 DB modules. One master module (DCB) will interface with the time-distribution protocol and propagate the timing information to the remaining 9 slave modules (see figure 3.13). In order to reduce the amount of physical links, the master will also concentrate the measurement data from the slaves, through the backplane and transfer collected data further to the CN for event building. The master module has to be equipped with a different FPGA device (e.g. Artix or Kintex) because ECP5G has not enough high-speed transceivers to accommodate 9 data streams from slaves and handle output links.

In the High Luminosity mode, the rate of 350 kHits/s is expected in stations FT1 and FT2. Under such conditions, one DB module with 64 channels will generate 90 MBps (0.72 Gbps). Chosen ECP5G FPGAs are equipped with 5 Gbps transceivers, both for backplane connections and for data transmissions out of the system. Fully-loaded crate with 9 slave and one master modules, will generate 6.5 Gbps data stream to the CN farm. The proposed solution will be constructed out of 191

new DB slave modules. In order to house them, 22 crates will be needed and the same amount of master modules. Each crate can be connected directly to the SODANET for time synchronization and to the CN farm for event building and further processing. In table 3.4 are presented the values of expected hit rates, required hardware and generated data volumes. They can be compared with those quoted in the previous section. All stations summed up will produce 9.6 GBps in the High Luminosity mode.

3.2.9 Luminosity Detector

Subsystem structure

The luminosity detector (LMD) consists of 4 layers of high voltage monolithic active pixel sensors (HV-MAPS). Each layer is composed of 80 HV-MAPS resulting in a total of 320 HV-MAPS. Hits are already zero-suppressed and digitized on the sensor itself. Then they are serialized into an 8b10b encoded data stream which is sent to front end FPGA boards using an LVDS link. Therefore a maximum of 320 data channels are to be expected from the LMD.

Expected hit-rate

Due to the boosted event topology of PANDA, the hit rate is not distributed uniformly over the LMD. Instead it is peaked forward strongly as can be seen in figure 3.14. The highest hit rate for the innermost sensors is around 130 kHz and falls off to rates below 20 kHz with increasing distance from the beam pipe. The above numbers were obtained using simulations with a beam momentum of 1.5 GeV/c and a luminosity of $1 \times 10^{32} \text{ cm}^{-1}\text{s}^{-2}$. The total hit rate from all sensors is given in table 3.5 where we assume a peak luminosity of $2 \times 10^{32} \text{ cm}^{-1}\text{s}^{-2}$. As can be seen from the table the highest data rates are reached at 1.5 GeV/c and drop down by a factor of 5 at 15 GeV/c. To compute the expected data rate, we make the following two assumption:


- At 95% hit detection efficiency a noise rate of

Table 3.5: Total hit rate for 320 HV-MAP sensors at $L = 2 \times 10^{32} \text{ cm}^{-1} \text{ s}^{-2}$.

p_{beam} [GeV/c]	$\text{rate}_{\text{total}}$ [MHz]	$\text{rate}_{\text{elastic}}$ [MHz]	$\text{rate}_{\text{inelastic}}$ [MHz]
1.5	18.82	16.56	2.26
15	3.70	2.20	1.54

Table 3.6: Data rate at $L = 2 \times 10^{32} \text{ cm}^{-1} \text{ s}^{-2}$ and 95% hit detection efficiency. The numbers do not include SODANET headers or other overhead.

	Physics Rate [Mb/s]	Noise Rate [Mb/s]	Total Rate [Mb/s]
Total Detector	677.52	37454	38131.92
Average per Sensor	2.1	117.1	119.2
Inner Sensor	10.0	117.1	132.1

**Figure 3.14:** Event rates at sensors of the first detector plane at 1.5 GeV/c for a average luminosity of $1 \times 10^{32} \text{ cm}^{-1} \text{ s}^{-2}$.

50 Hz is to be expected [1]. This value is likely to decrease with improved sensor prototypes.

- each hit produces 36 bit of data (row and column address, timestamp and deposited charge)

The calculated data rates can be found in table 3.6. The data rate is clearly dominated by the pixel noise and the difference in terms of data rate between an inner and outer sensor can be neglected in the following considerations.

10 Front-end Electronics

As front-end FPGAs the Hades Trigger and Readout Board (TRBv3) will be used. Each TRB has four peripheral FPGAs and one central FPGA. One peripheral FPGA will be used to readout the data from 4 HV-MAPS. This means deserializing and decoding the 8b10b data stream, transforming the chip row and column address into the physical address, clustering, and changing the timestamps from a gray counting scheme to a binary counter. The data coming from the peripheral FPGAs is then merged in the central FPGA and send to the data

concentrator using optical signal transmission. This optical link will have to transmit $4 \cdot 4 \cdot 130 \text{ Mb/s} \approx 2.0 \text{ Gb/s}$. As the optical link speed of the TRB is limited to 2 Gb/s while using SODANET, two optical links will be used per TRB. The total number of links from the LMD is: $(\# \text{HV-MAP sensor}) / (\# \text{Peripheral FPGA} \cdot \# \text{HV-MAP sensor per FPGA}) \times (\# \text{Links per TRB}) = 2 \cdot \frac{320}{4 \cdot 4} = 40$. The front-end FPGAs will also handle the slow control of the sensors and the communication with the SODANET source.

3.2.10 Readout summary

Subsystems presented in this section will produce most of the data which will be streamed to the data acquisition. In summary, for the high-luminosity mode the total amount of data is 5 (MVD) + 20 (STT) + 25 (Barrel DIRC) + 11 (EDD) + 1 (Barel TOF) + 3 (Muon) + 17 (EMC) + 10 (Forward Tracker) + 4 (luminosity) = 96 GB/s. All other subsystems, not included into this overview will produce not more than few GBps of the raw data. Since in the phase one PANDA will operate at about 10 times lower luminosity the expected data rate streamed to the DAQ is about 10 GB/s. Therefore, this TDR is based on the latter data rate.

3.3 On-line storage and requirements for event filtering

The required online storage depends on the size of the data set one wants to have access to at the same time. In the case of well separated events it should be sufficient to access just a single event with an

event size of a few hundred kByte of detector raw data depending on the number of particles inside an event. With the quasi continuous data stream of PANDA the events are not well separated in time which makes an event separation on an early processing stage difficult. Therefore a more suitable packaging would be on bursts of about $2 \mu\text{s}$, matching the convolution time of HESR. Within one burst three quarters of the ring are filled with antiprotons while one quarter is empty. Within this time on average 3-4 antiproton proton reaction take place which would lead to a data size per burst of about 1 MByte. During the processing of the raw data the data size increases by a factor of 5.

To transmit small packages is very inefficient due to the protocol overhead which is more or less independent of the package size. Therefore it is better to increase the size of a data package by grouping 36 bursts into a super-burst which would have the size of 180 MByte. The typical size of 1-2 GByte of RAM per processing node should therefore be fully sufficient to hold one super-burst of data plus the needed calibration parameters from a parameter data base.

3.4 DAQ partitioning and running modes

The PANDA DAQ should have modular design which would allow to run in different configurations. Different modes of operation include:

- complete PANDA detector at FAIR facility (production mode);
- one or few PANDA subsystems at FAIR facility;
- parts of several subsystems in a stand-alone mode;
- part of one subsystem in stand-alone mode.

In all listed above modes it should be possible to:

- enable/disable data-processing algorithms in case some components are not present/used, e.g. track finding/fitting in case the magnet is off;
- acquire debugging information from any sub-detector, e.g. digitized waveforms for the EMC (in this case taking data from complete detector is not possible, however, it should be possible to run another instance of partial DAQ for other subsystems);

- DAQ to interface with experiment operations system for production data taking ;
- possibly DAQ to interface with DCS.

It is required that all data-processing before the data-reduction step is not-destructive. Namely, all raw data after zero-suppression should be kept. All data-processing algorithms should re-structure incoming data-stream according to the physics features and add new headers to newly-formed packets with relevant description of found features, e.g. find hits which belong to the same cluster, put these raw data into one packet and specify in the header position, time end energy information of the found cluster.

Bibliography

4 System Architecture

The trigger-less readout of the PANDA detector is realized with self-triggered intelligent front-end electronics (FEE), a very precise time-distribution system and on-line data processing algorithms. Figure 4.1 shows the functional diagram of the trigger-less readout scheme for the PANDA experiment. An intelligent front-end detects particle hits in the detectors and processes the raw data. The resulting data are sent to the data concentrator (DC). The DC modules assign incoming data to certain time periods, called bursts, and send data through a burst-building network for on-line processing to compute nodes (CN). The burst-building network guarantees that the data from all PANDA sub-detectors, which belong to the same burst, arrive at a single CN. At the CN hits are combined to clusters and tracks and after identification of particles primary annihilation events are reconstructed. Since at this stage all incoming data has to be processed, the applied tracking and reconstruction algorithms should be fast and use fewer resources, therefore, yielding lower resolution. Depending on the reconstructed secondary particles events will be selected for further precise processing on a computing farm where the final event filtering will take place and the resulting data will be stored on disk.

During phase 1, the experiment will be operating with high interaction rates up to average of 2 MHz in a freely streaming mode. Without any hardware trigger, up to 10 GBps of raw data are streaming into the data acquisition (DAQ) system. Before mass storage, the data rate has to be reduced by at least two orders of magnitude. This will be achieved by filtering the events based on their physics content using information from reconstruction algorithms implemented on Field Programmable Gate Arrays (FPGA) as a first level and on a farm of servers with attached GPUs as a second level. In the following, we discuss two FPGA based algorithms, a fast track finder for the STT and a cluster finder for the EMC.

Section 4.1 describes the hardware building blocks of the system as data concentrators, FPGA based compute nodes, CPU/GPU farm. Section 4.2 describes synchronization and data-transport protocols, data formats, interfaces and data flow. Section 4.3 describes the event filtering system and its partitioning into FPGA based and CPU/GPU based algorithms. Section 4.4 discusses run control (RC), error handling and data quality monitoring (DQM).

4.1 Basic building blocks of the system

4.1.1 The Data Concentrator module

The Data concentrator for the PANDA is a designed to collect data from the PANDA front end electronics, as well as to distribute detector synchronization signals. The same module can be used to implement burst-building network, see figure 4.1. The module is housed in a single AMC board with full size width. The module features 60 optical transceivers and 16 backplane multi-gigabit links, whereas it is planned to reach at least 12 Gbit/s on all optical and electrical links, see figure 4.2.

In order to maintain high optical link density, the module is equipped with 12-channel FireFly transmitters and receivers from Samtec. Since the FireFlies are simplex devices, it is necessary to implement one 12-channel FireFly receiver and a 12-channel FireFly transmitter to achieve every 12 full duplex channels. Optical fibre pigtailed connect FireFlies with front panel MPO-type feed-through, see figure 4.3. There are five 12-channel feed-throughs. The data are collected, buffered and processed by an on-board Kintex Ultrascale FPGA.

The FPGA has a general LVDS connection for "slow" control (up to 2 GB/s) over the AMC bus. In order to support defined latency systems, a stable clock with low phase noise is necessary. The Data Concentrator module comprises a dual-PLL jitter cleaner and clock distributor LMK04610 from National Semiconductor (currently Texas Instruments).

The module is supplied from 12 V MicroTCA crate and all necessary internal voltages are generated by on-board DCDC converters. The estimated power consumption amounts to 49 W, hence the module needs efficient cooling. The voltage turn-on and off sequencing are controlled by an Atmel microcontroller. For full power control, key voltages and currents are monitored by a System Monitor located in the on-board FPGA.

4.1.2 FPGA based Compute Nodes

The FPGA based compute nodes with multi-Gbit/s bandwidth capability are implemented using the




Figure 4.1: The overview of the PANDA readout system.




Figure 4.2: Block diagram of the PANDA Data Concentrator.




Figure 4.3: A preliminary layout of the PANDA Data Concentrator.

ATCA architecture and are designed to handle tasks such as event building, feature extraction and high level trigger processing. Several incarnations of the CN were developed in close collaboration between IHEP Beijing and JLU Giessen. The original version consists of a single full-size ATVA board hosting 5 Virtex 4 FPGAs. The ATCA shelf is equipped with a full-mesh backplane allowing for high speed serial data transfers between any of the slots. This module has been used for the development of the FPGA nased trigger algorithm for the STT. The present version of the CN follows a mod-

ular concept consisting of a full-size ATCA carrier board and up to 4 AMC cards ("xFP").

The FPGA on the xFP card is a commercial-grade Xilinx Virtex-5 FX70T with speed grade -2 in an FFG1136 package. The FPGA contains 11200 slices with 44800 6-input LUTs and 44800 flip-flops. There are 148 block RAM units with a total capacity of 5328 kbit. I/O is provided by 640 general purpose pins and 16 GTX transceivers supporting line rates up to 6.5 Gb/s. A PowerPC 440 hard core CPU and 4 Ethernet MACs are included on the chip.

The xFP cards feature two soDIM sockets supporting up to 4 GB DDR2 SDRAM modules. The large storage capacity is an essential feature, allowing




Figure 4.4: ATCA compliant FPGA based compute node. (top): carrier board; (bottom) xFP daughter board. The architecture is based on XILINX[®] Virtex FPGAs. A single CN can support up to 16 optical links and a total of 18 GB DDR2-RAM. The FPGAs are equipped with embedded PowerPC - CPUs, running Linux operating systems for slow control functions.

for complex algorithms with large latency. An attached Flash memory chip provides 64 MB of non-volatile storage. The configuration bitstream can be stored in a 4 MB Xilinx Platform Flash chip. Two SFP+ cages are connected to the high-speed serial transceivers of the FPGA. They can be used either with an optical transceiver or with a 1000BASE-T-Ethernet transceiver allowing GbE connections using standard copper network cables. Six more MGTs as well as 12 bi-directional LVDS links serve as connections to the carrier board. A GbE PHY chip with RJ45 connector provide an additional GbE connection. For debugging purposes, there is an UART interface.

The carrier board hosts a commercial-grade Xilinx Virtex-4 FX60 FPGA with 25280 slices. Two PowerPC 405 CPUs and four Ethernet MACs are included as embedded cores on the chip. There is 64 MB of flash memory attached but only a single SO-DIMM socket supporting up to 2 GB of RAM. The FPGA is connected to each of the four AMC slots via 4 bi-directional LVDS links. A GbE PHY connects the FPGA to the first ATCA backplane base channel. All 16 MGTs of the FPGA are routed to fabric channels. In a full mesh ATCA shelf equipped with CNs, it is possible to route data packets from

the FPGA in any xFP card to any other one in the shelf, the FPGA on the carrier board serving as a switch.

Three AMC ports of each AMC bay are used to create a threefold full-mesh interconnection between the inserted cards in the carrier board. Five more ports of each AMC are routed to the RTM connector. The switch FPGA also has 16 LVDS links and an additional Ethernet interface via a second PHY connected to the RTM.

The carrier board can host up to 4 xFP cards which are equipped with Virtex5 FX70T FPGAs, up to four optical links, 4 GB DDR2 - RAM and GBit ethernet. The system is scalable and can be optimized for pipelined and parallel architectures. The FPGAs feature embedded PowerPCs running Linux operating systems for slow control functionality. The large on-board RAM per FPGA is essential, allowing to store data for complex algorithms with large latency.

We are currently developing an upgrade based on Kintex 7 Ultrascale, following the same architecture but with significantly increased resources. Two prototypes of the new carrier board are already in operation. They new design is backward-compatible with the existing Virtex-5 based xFP cards.

4.1.3 FPGA based HPC interface

After pre-processing in the FPGA-based compute nodes data are streamed to High-Performance Computer (HPC) cluster for final event reconstruction and filtering. The data transfer is done using optical interface. New generation of FPGA-based compute nodes as well as DC hardware will be equipped with FireFly optical transmitters which employs custom communication protocol while the HPC cluster at FAIR uses InfiniBand network-interface. Therefore, a dedicated PC is used to match both interfaces. The data are received by a High-Flex readout card [1], developed at KIT, Karlsruhe, Germany, see figure 4.5. The current prototype of the up-



Figure 4.5: PCIe High-Flex readout card [1].

dates High-Flex card is equipped with two FireFly optical interfaces with 12 lanes at 28 Gb/s (full-duplex) which provides 330 Gb/s bandwidth and additional FireFly modules can be connected using FMC-FireFly card (commercially available). Received via optical inputs data can be processed by Xilinx ZYNQ UltraScale FPGA and sent through 16 Gen 3 and 4 PCIe lines with total bandwidth of 240 Gb/s to GPUs installed in the host PC or to HPC cluster via InfiniBand interface card of the host PC, see figure 4.6. Unique feature of the




Figure 4.6: The basic concept of direct FPGA-GPU data-link tunnelled through InfiniBand network [1].

High-Flex card is capability of direct data transfer from the on-board FPGA to GPU memory with

low latency bypassing the system memory, see figure 4.7. This opens-up possibility for online data




Figure 4.7: Illustration of the data-flow of the traditional approach (dashed line) and direct FPGA-GPU communication (top), and round-trip latency measurements. (Figure taken from [1].)

processing on installed locally GPUs. In addition the High-Flex board is capable to write data directly into GPU memory through InfiniBand link, see figure 4.6, for remote GPUs installed in the HPC cluster.

In the PANDA DAQ the High-Flex card is employed as an interface between the FPGA-based layer of the data acquisition and CPU/GPU-based one. This interface card will be used with local GPUs/CPU for stand-alone setups (outside FAIR facility) while InfiniBand interface will be employed for the PANDA experiment only.

4.2 Protocols, data formats, interfaces and data flow

As can be seen from figure 4.1 and section 4.1 hardware layer of the PANDA DAQ below computing farm is composed from custom FPGA-based hardware. In order to control, synchronize and read-out these hardware a dedicated protocol SODANET has been developed. This section describes the SODANET protocol and data flow within the DAQ.

4.2.1 Synchronization protocol SODANET

4.2.1.1 Architecture

The time-distribution system is of key importance for the trigger-less readout, since all front-end modules acquire data independently, and precise time-stamps are needed to correlate hits from different sub-detectors. A single clock and synchronization source is used for the complete PANDA detector. The precise clock and synchronization commands are distributed using optical-fiber connections. After successful tests of the stand-alone synchronization protocol [2, 3], it has been combined with the slow-control and data-transfer protocol TRB-NET [4], used in the HADES experiment [5]. The combined protocol is named SODANET. The topology of the SODANET network for the PANDA experiment is shown in figure 4.8. A single SO-




Figure 4.8: The topology of the SODANET network for the PANDA experiment.

DANET source supplies via hubs the system clock and synchronization commands of fixed latency. In addition, the network is used to distribute slow-control (e.g. the front-end configuration, system status) data. The transmission of slow-control data can be interrupted at any time by a synchronization command. The synchronization command, a 32 bit long word, together with four dedicated "K-characters" (with hexadecimal value FB) is embedded into the currently transmitted data. At the receiving end the incoming data are continuously analyzed for the presence of the dedicated four K-characters. Once the K-characters are detected, the synchronization command is extracted from the incoming data stream and the receiving end takes the synchronization action. The remaining slow-control data are not distorted. All SODANET

K	Data	K	Data	K	Data	K	Data
fb	bits	fb	bits	fb	bits	fb	bits
	31-24		23-16		15-8		7-0

Figure 4.9: Structure of the SODANET packet as it is sent through a serial link. Each block corresponds to 1 byte

links in the downstream direction, i.e. from the source to the FEE, are operated in synchronous mode: all Serializers-Deserializers (SERDES) have fixed phase of the serial and parallel clocks [2] and no buffers in the data-path, which might change the latency of the data transfer. The FEE uses the recovered parallel clock as the system clock. The upstream link, i.e. from the FEE to the SODANET source, is not synchronous. The upstream link is employed to transfer slow-control data from the FEE and to measure the time interval of signal propagation through the link with a precision of about 10 ns. This precision is enough for an initial calibration of the readout system during the start-up phase. The larger uncertainty arises due to the lack of a synchronous upstream link. The operation of both the down- and up-stream links in synchronous mode would require a very long start-up time, since all transmitters and receivers should lock on a same phase and this would make the operation of larger systems impractical.

4.2.1.2 Protocol description

Interface: The SODANET uses serial link with 8b/10b encoding. SODANET commands are 32-bit long words. In order to make SODANET compatible with the TRB protocol, SODANET package has the following structure as shown in figure 4.9, where the total length of the sent word is 64 bits. The K-characters *FB* are used to recognize synchronous data embedded into the data stream. Serial clock of the link is derivable from the base 40 MHz clock and speed of the link is determined by capability of the hardware to operate in synchronous mode.

Clocking: The base clock for the SODANET is 40 MHz. All other clocks should be derivable from the base clock. The serial clock of a link is used to transfer the system clock. At the receiving side SERDES transforms serial clock into a parallel without introducing random phase shift which might occur during standard initialization of a SERDES. Therefore, a special configuration of the SODANET, or synchronous, SERDES is required. After decoding, the parallel clock is cleaned with

a jitter cleaner circuit which guarantees the same phase of the cleaned clock, so called zero-phase jitter cleaner.

Commands: SODANET packet with command data can be issued at any time. Each SODANET command consists of one 32-bit word. Each command has the following structure:

- bit 31: type of the command defined by the value:
 - 1: super-burst number (this command is used for synchronization and is named "SB-start");
 - 0: DAQ command;
- bits 30 – 8: super-burst number or SODANET command;
- bits 7 – 0: CRC checksum (CRC8-CCITT).

The SODANET protocol foresees the following commands which should be recognised by all front-end electronics:

- bit 31: 0 (SODANET command);
- bit 30: 1 (rough propagation-time calibration of the SODANET links);
- bit 29: 1 (start DAQ, start data taking including current super-burst);
- bit 28: 1 (stop DAQ, stop data taking from the next following super-burst);
- bit 27: 1 (Reset FEE);
- bits 26 – 8: reserved;
- bits 7 – 0: CRC checksum (CRC8-CCITT).

In order to implement continuous monitoring of the readout each received SODANET packet is acknowledged. The feedback packet is 16-bit long:

- Reply to the packet "super-burst number": [K-character (FB)] [bits 7-0 of the burst number];
- Reply to the DAQ command: [K-character (FB)] [CRC checksum].

Each reply is being checked at the source side and in case of wrong or missing feedback the error is reported to the slow-control subsystem by changing values of the dedicated registers. The monitoring is performed at each physical serial link. In general, the monitoring protocol for the readout can be described using figure 4.8. The monitoring protocol is defined as following:

- DCs have a report-register and a status-register with n-bits for n-clients/FEEs;
- status-register holds '1' for each known failing-client; '0' for others;
- at soda-command transmit, status-register is copied to report-register; watchdog started;
- client soda-reply causes corresponding bit to be set in report register;
- at watchdog timeout '0's in report-register cause error-report to HUB; all '1' = success;
- HUB handles DC-reports in a similar way;
- slow-control handles error-reports, tracing through HUB- and DC- report-registers;
- slow-control can set/reset status-bits on HUB and DCs.

Synchronization: Since all FEE modules share the same clock, see section 4.2.1.1, timing of detected hits is done locally at the FEE side and the local timing circuits are periodically synchronized. The PANDA experiment will employ anti-proton beam which has a "burst" structure: time of the beam on target and a time-gap before the following beam-on period. Length of the burst depends on the momentum of anti-protons and is changing in a range of 1.4 – 2.0 μ s. At the beginning of each 16-th burst the SODANET source, see figure 4.8, issue SODANET command "start of a super-burst" (**SB-start**), which consists of the number of the starting super-burst period. All timing circuits at the FEE side are regularly reset to zero by this synchronization command and in this way are synchronized. The SODANET receivers check if received SB-number is sequential. In the case of error (not sequential number) the receiver uses number distributed by the SODANET, sets special error bit in the output data, and informs the slow-control system. If some part of a SODANET message is missing, receiver uses burst number from the local counter and reports problem to the slow-control system.

The start of the super-burst signals are issued periodically with the period which is the closest to duration of 16 accelerator bursts and is a multiple of 25 ns. Once the difference between issued SB-start signal and the start of accelerator burst becomes larger than 25 ns the duration of the current burst is corrected by one 40 MHz period.

Propagation-time: A dedicated SODANET command is foreseen for calibration of the signal-propagation time (see SODANET Commands in the




Figure 4.10: Measured propagation time (T_d) of a synchronization command from the source to each FEE is used to correct local timestamps.

31	16	15	0
last-packet flag, packet number		data size in bytes	
Not used (same as HADES)		Not used (same as HADES)	
Status and error		System ID	
Super-burst number			
Data			

Figure 4.11: Header structure of a SODANET data-packet.

text above). Once the command is received a reply is send to the transmitter side, and original message is forwarded further through the SODANET network. At the transmitter side propagation time is calculated and stored in a register. The register values are read out by a slow control system. These data are used to pre-calculate signal-propagation delays T_d with precision of about 10 ns and are used to correct local time-stamps produced by FEE, see figure 4.10:

- each FEE module has own timing;
- local "time zero" is reset with each synchronization SODANET command "start of a super-burst";
- each hit time-stamp is corrected with the T_d value;
- after correction the time-stamp the hit data, including current super-burst number, are sent to DC module;
- at the DC module decision is taken to which super-burst the hit belongs.

Once rough calibration is performed more precise determination of T_d values is performed using special detector data, e.g. simultaneous injected light-pulses. These precise T_d values for all FEE are valid for whole data-taking period unless there are some hardware changes in the system.

Output data format: SODANET protocol specifies structure of the output data as following, see figure 4.11, at the level of data concentrators. DC starts transmitting FEE data once it is available, without waiting till the end of a super-burst. Each data-packet should contain:

- super-burst number;
- system identifier (System ID) (EMC, MVD, etc.);
- packet number within super-burst data;
- a flag "last packet", for the last packed with data associated to the current super-burst.

The SODANET data-header is compatible with HADES sub-event structure [4]. The adoption of the existing data structure allows to reuse major parts of code: GbE packet builder in FPGA, "Event-builder" server software. In order to keep compatibility, the maximum packet size is 62 kB (maximum UDP packet size). However, arbitrary number of packets per super-burst can be transmitted by a data concentrator.

4.2.1.3 Hardware implementation

For the development of the SODANET protocol the Lattice ECP3 FPGA hardware is employed. At present, the network operates at a speed of 2.0 Gb/s while the maximum achievable synchronous speed of a link is 2.4 Gb/s. The synchronization commands are periodically issued with a period chosen as a multiple of 25 ns (40 MHz clock). Such a combination of link speed and synchronization frequency allows to synthesize, at the FEE side, clocks of 40, 80, 120, 160, 200 MHz which are always in phase with the single SODANET source. At present the SODANET source, hub and endpoint are implemented and perform up to specifications.

As mentioned in the text above, the synchronous optical-link connection is a crucial component of the trigger-less readout. In order to send data via optical link the parallel data is encoded byte-wise into serial data. In the current prototype a standard 8b/10b encoding (mapping 8-bit symbols to 10-bit symbols) was used. On the receiving side, serial data is decoded into parallel data. The standard implementation of the serial optical data-link does not guarantee a stable phase of a recovered parallel clock signal with respect to the beginning of a serial data word. For different power or reset cycles the phase can take 10 different values. This problem can be avoided by using a specially designed communication protocol and a suitable modification of the hardware related to the optical-link transceivers/receivers [3]. In this work was found that for the Xilinx FPGA it is possible to establish a synchronous serial optical-link connection by applying special settings to the FPGA serializers/deserializers (SerDes) and building a specialised




Figure 4.12: The photograph (upper panel) and a sketch (lower panel) of the set-up used to test the synchronous serial optical-link connection.

state-machine which is establishing and monitoring a connection. At least the GTP and the GTX SerDes of the different Xilinx FPGA families can be configured to guarantee a stable phase, however, the SerDes will lock on a signal only once in about 10 trials. Therefore, a special state machine was designed, which resets the SerDes until it is locked on a signal.

4.2.1.4 Stability verification of the SODANET

A special set-up, shown in figure 4.12, was constructed in order to test the reliability and performance of the synchronous optical-link connection. The set-up consists of the Data-MultipleXer (DMUX) module and two FPGA development boards, serving as digitizers (front-end modules). For the tests, the DMUX module internally generated a clock signal and SODA time-synchronisation commands. The SODA commands were generated with a frequency of 1 kHz and sent to the front-end modules. In addition, for every




Figure 4.13: Stability of the recovered phase. During the measurement the front-end module was periodically reset. The insert shows a zoom of the region around 30 ns.

SODA command a short NIM pulse was generated by the DMUX module. This pulse is further referred to as T_0 . On the front-end side decoded SODA commands were triggering TTL pulses T_1 and T_2 for both front-ends, respectively. For properly functioning synchronous serial connections the time differences $\Delta t_1 = T_1 - T_0$, $\Delta t_2 = T_2 - T_0$ and $\Delta t_J = T_1 - T_2$ should be constant. In order to estimate the stability of these time differences, all timing signals were measured by a CAEN V775 TDC [6]. Figure 4.13 shows the time-difference spectrum for $\Delta t_1 = T_1 - T_0$. During the measurement, the front-end module was periodically reset by a hardware signal and by disconnecting the optical fibre. The measured single narrow peak, without satellites, proves the phase stability of the recovered parallel clock signal with respect to the beginning of the serial data-word.

The measurement of the width of the $\Delta t_J = T_1 - T_2$ distribution can provide information on the time jitter of the synchronisation commands. Even though the used hardware did not allow to make a precise measurement of the Δt_J value, it was possible to estimate an upper limit of the jitter. Figure 4.14 shows the $\Delta t_J = T_1 - T_2$ distribution. The RMS value of the distribution is measured to be about 40 ps. However, a separate measurement of the jitter induced by the set-up, e.g. by level adapters or the precision of the TDC, yielded a value of about 22 ps for the same time difference. The measurement was performed using the same set-up, but electrically one front-end board was triggering T_1 and T_2 TDC channels. The resulting Δt_J distribution is shown in figure 4.15. The shape of the Δt_J distributions in figures 4.14 and 4.15 is not symmetric. This asymmetry in both distributions indicates a small systematic instability of the time-measurement set-up. Nevertheless, it is safe to con-




Figure 4.14: Difference of arrival times of synchronization commands at different front-ends.




Figure 4.15: Measurement of the jitter induced by the measurement set-up.

clude that the jitter per front-end channel is less or about $\sqrt{(40 \text{ ps})^2 - (22 \text{ ps})^2}/\sqrt{2} \approx 23 \text{ ps}$. This value of the jitter is sufficiently low for the PANDA EMC readout.

The long-term stability of synchronization provided by SODANET was tested using set-up shown in figure 4.16. Three peripheral FPGAs of a TRBv3 board were used as a SODANET source and two EMC data concentrators. The SODANET synchronous link was connected from the source to both DC FPGAs. Both DC FPGAs were connected with two FEE modules: EMC digitizer and Xilinx Kintex-7 development board, which was pro-




Figure 4.16: Scheme of the setup used for the long-term stability test of the SODANET.

grammed as EMC digitizer. Firmware of digitizers was modified to provide a signal which corresponds to the signal "start of super-burst". In case of EMC digitizer hardware two LVDS signals were generated – one using recovered clock directly after receiving SERDES and another with a recovered clock which was stabilized by hardware jitter cleaner. The Kintex-7 development board provided only "start of super-burst" LVDS signal based on recovered clock without jitter cleaner. All LVDS signals were converted to ECL standard and fed into CAEN V775 TDC [6]: signal from the development board as a "start" input and all other signals as TDC stop. One LVDS signal from the EMC digitizer, generated using cleaned clock, was split into two channels before the LVDS-ECL level adapter. These two signals were used to determine performance of the level adapter and TDC input stages.

Time-data were recorded with the rate between 100 Hz and 2 kHz, depending on the duration of the experiment. Collected data were split into the time-bins of one minute. For each time-bin histograms for measured time-differences were filled and analysed in terms of the mean value. Changes of this measured mean values in different time-bins were used to conclude on stability of the synchronization. The following time-differences were measured:

- $\Delta T = T_1$: time difference between receives "super-burst start" signals – measure of stability of the SODANET system;
- $\Delta T_{j.clean} = T_2$: time difference between receives "super-burst start" signals with jitter cleaner in one branch – stability of the SODANET system with jitter cleaner;
- $T_{j.clean} = T_1 - T_2$: difference between signals "start of super-burst" generated with recovered clock and the one after jitter cleaner – effect of jitter cleaner;
- $T_{TDC} = T_3 - T_2$: difference between the same signal from digitizer – stability of the level adapter and TDC input stages;

where T_1 , T_2 and T_3 are referred to the outputs of the FEE as shown in figure 4.16.

Typical histogram of measured T_{TDC} for few time-bins is shown in figure 4.17. AS can be seen the standard deviation for the time-difference is about 20 ps which determines the maximum accuracy of the measurements. Histogram for ΔT for a single time-bin is shown in figure 4.18. The standard deviation of the distribution is about 45 ps. Taking into




Figure 4.17: Time-difference between the same signal measured by two TDC channels (T_{TDC}).




Figure 4.18: SODANET jitter (ΔT).

account width of the distribution for the same signal, see figure 4.17, we conclude that the standard deviation induced only by signal differences from different FEE boards is $\sqrt{45^2 - 20^2}/\sqrt{2} \approx 28$ ps per channel which agrees with the previous measurements, see text above.

The long-term stability measurements are shown in figure 4.19. As can be seen from the figure (top panel) at the beginning of the measurements there is a jump in all measured time-differences. At the beginning of the measurement the EMC digitizer was cooled to the room temperature while all other components were warmed-up. The warming-up process influenced the measured values. After all components of the readout reached their working temperature the performance of the SODANET system was stable within 30 ps, see figure 4.19 (lower panel), which is sufficient precision for the PANDA experiment.

In order to estimate effect of temperature changes on the synchronization accuracy additional measurement was performed with warming-up of the




Figure 4.19: Stability of SODANET synchronization. The lower panel shows stability of the system once the hardware warmed up to the working temperature.




Figure 4.20: Effect of the warming-up process of one FEE on the synchronization stability (upper panel) and the temperature of the FEE FPGA (lower panel). Curves with different colours represent two independent measurements.

EMC digitizer. The warming-up measurements were performed two times to check if the results can be reproduced. The ΔT measurements as a function of time are shown in figure 4.20 in correlation with the temperature of the FPGA of the EMC digitizer. During these measurements cooling of the digitizer was off in order to achieve extreme




Figure 4.21: Hardware set-up used for long-term stability-tests of the SODANET synchronization protocol.

temperatures. As can be seen from the measurements the measurements are reproducible and can be employed for active delay correction. However, assuming that temperature in the experiment environment would not change more than 10 degrees it is possible to conclude that the possible shift in synchronization should be within 30 ps, which is acceptable for the PANDA experiment.

Precise long-term SODANET stability-tests have been performed with a setup composed out of three TRBv3 boards: one serving as SODANET source (TRB0) and two serving as both DC and FEE, designed for PANDA Forward Tracker subsystem (TRB1 and TRB2), see figure 4.21. TRB0 provided the SODANET synchronization packets, which were received by the central FPGAs on the endpoint TRBs (they can be treated as Data Concentrators) and the superburst update pulse was delivered to peripheral FPGAs configured with TDC firmware (they can be treated as Front-End modules). The TDC FPGAs are clocked with local oscillators and therefore are independent of SODANET clocking. The superburst update signal is used for both: common reference time for the TDC and as input to couple of TDC channels. In this way, it is possible to synchronize multiple TDCs in the system, because they all use the signal from the same origin, but locally recovered as time reference.

For each SODANET synchronization packet, a superburst update signal is generated and its time is registered by all TDCs in a form of 32-bit data word and stored in a buffer. The update signal also triggers buffers readout and construction of UDP packets, which are transmitted to the analysis computer. The data is stored with DABC software [7] and analyzed with GO4 framework [8]. First the times are normalized by subtracting the value registered as common reference time and then the time difference between channels of two separate TDC FPGAs can be calculated.

Described above setup and procedure have been




Figure 4.22: Long-term stability of SODANET synchronization at constant room temperature. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.

used for collecting and analyzing data in order to determine the SODANET stability over long run-times and temperature changes. Temperatures are collected at fixed time intervals from dedicated temperature sensors located nearby each FPGA in the system. The RMS of superburst update signal is calculated as the RMS of the time difference between two TDC channels of the TRB1 and TRB2 on statistics collected over the each defined time interval.

On figure 4.22, result of such scan is presented for constant, room temperature. The FPGAs get slowly, not significantly warmer over time (blue, green, yellow, red and black curves, right axis). The RMS of the time difference (pink curve, left axis) remain stable on the level of 19 ps with 10% jitter which also comes from the TDC time measurement resolution.

Similar results have been achieved for the setup where an addition SODANET hub has been introduced, see figure 4.23. TRB1 has remained connected directly to SODANET source but TRB2 has been connected through SODANET hub.

Further investigations involved forced temperature changes on selected system components. On figure 4.24 a significant change of the temperature on SODANET source (yellow curve) has been enforced. The RMS does not reflect the temperature shift. The same behaviour is present in case the temperature on one of the endpoints is drastically changed (blue curve on figure 4.25).

The measurement presented above, show stable operation of SODANET synchronization sys-




Figure 4.23: Long-term stability of SODANET synchronization at constant room temperature measured with additional SODANET hub in the network. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.




Figure 4.24: Long-term stability of SODANET synchronization at forced temperature change of the SODANET source. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.




Figure 4.25: Long-term stability of SODANET synchronization at forced temperature change of one of the SODANET end-points (TDC). Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.

tem which is resistant for significant temperature changes and long run times. One can notice distortions in the time difference RMS, but their levels remain within 10%.

4.2.1.5 Protocols summary

5

The SODANET protocol is employed to configure and monitor all front-end electronics including data concentrators and compute nodes. The SODANET foresees two types of communication messages: synchronous, with fixed latency (SODA) and asynchronous (TRBNET). Synchronous messages are sent only to the front-end electronics, namely to the DC-level and below, while TRBNET is used to configure and monitor DAQ till the HPC interface. In summary information distributed by the SODANET includes:

10

15

- fixed-latency messages:
 - synchronization messages;
 - control messages for data-taking as start/stop of data taking (these messages are distributed not only at the beginning/end of a run, but during a run in order to implement data-taking throttling);
 - calibration of a signal-propagation time (these commands are triggered by Detector Control System (DCS) at the start-up phase of the DAQ or to monitor changes in the signal-propagation time);

20

25

- asynchronous messages:

30

- FEE configuration (DC level and below);
- topology of the readout which defines topology of the burst-building network, CN configuration and overall data-flow (DC level and above);
- monitoring readout including status of the FEE
 - * DCS information as temperatures/voltages;
 - * relevant performance parameters monitored by FEE;
 - * status of error registers from all components of DAQ (these registers store information on any detected fault as missing data from particular FEE/DC, buffers overflow).

SODANET defines format of data packages after the DC level, see section 4.2.1.2. This format of the data-packets is kept till the CN level, where data from complete detector for every super-burst is combined. Format for the combined data as well as physical protocol between DC, burst-building, CN and HPC interface will be defined once final specifications of all hardware are defined.

4.2.2 Data Flow

The overview of the PANDA readout system is shown in figure 4.1. Front-end electronics pushes data to the Data Concentrator (DC) modules where it is time-ordered, possibly pre-processed and pushed to the Burst-Building Network (BBN). The BBN is ordering incoming data according to the SB-number (see section 4.2.1), possibly pre-process it and distribute to Compute Nodes (CN) for processing by high-level algorithms, e.g. tracking, event pre-selection. After the CN level data is streamed to a high-performance computing cluster for the final event selection and storage.

A different picture of the data flow model of PANDA is given in figure 4.26. It shows, starting from the different sub-detectors, the sequence of the various processing stages and the required data at each stage for the barrel part of the detector. The sequence for the forward part is identical. The gray boxes to the left resemble the local processing stages for each sub-detector. All have to do time-sorting of the raw data based on the time-stamps assigned to each detector hit. After that the MVD, STT and EMC detector use their calibration parameters from a database to calculate the deposited energy in each sensor element followed by a clustering algorithm which is also done for the GEM detector.

After that the hit position for the tracking detectors or the bump position for the EMC is calculated based on the alignment parameters of the detector. To do tracking in the barrel part the data of all tracking detector is needed. They first perform a rough event building before a primary and secondary track finder are operating on the combined data. At this stage in addition the information of the field maps and material maps are needed. The resulting tracks are then combined with the data from the PID detectors and the EMC to do the construction of charged particle candidates. Bumps in the EMC which could not be associated to a particle track are then used to create neutral particle candidates. In the last stage of the processing the different track candidates with the now much better timing information are combined into events which are then the input to the event selection.

This section describes in details how the PANDA readout can be implemented starting with FEE interfaces for crucial PANDA subsystems.

4.2.2.1 Micro Vertex Detector

Readout of the MVD subsystem, see section 3.2.1, is done by GBT optical links. In total there are 68+95 GBT links from pixel part (barrel and pixel parts, respectively) and 10 from micro-strips. PANDA DC module, see section 4.1.1, can be used as the MVD Multiplexer Board (MMB) – the off-detector electronic that will manage data from up to 48 GBTs and send them to the BBN of the DAQ system, using 12 optical link operated at 12 Gbit/s. In addition, it receives the clock signal from the time distribution system of the experiment (SODANET). **Four DC boards** with $4 \times 12 = 48$ outgoing links to the BBN are enough to readout complete MVD detector. In the phase-1 of the PANDA experiment the interaction rate will be an 10 times lower than the maximum design value. Therefore, the total required bandwidth from DC to BBN level is 10 times lower as well and number of output links can be reduced to **four (links to BBN)**. If required, the load on the output links can be implemented using backplane communication between four DC boards. One should note that the total amount of data produced by the MVD is much less than the total bandwidth of the data-lines coming from the detector, see section 3.2.1. Therefore, one link per DC board has sufficient bandwidth for the phase 1.




Figure 4.26: Schematic view of the data flow to the different processing stages within the PANDA experiment

4.2.2.2 Straw-Tube Tracker

The STT front-end boards, the PASTTREC3s 3.2.2, may be readout by different TDC back-ends. For the PANDA phase 1 and phase 2 it is foreseen to use TRBv3 boards as back-ends. The discriminated straw signals are fed via micro-twisted pair cables to the TRBv3 boards. In addition to the time and TOT readout, the central FPGA in each TRBv3 board is used to set and control parameters of the connected ASICs. The TRBv3 readout system fulfils the data bandwidth requirements for the PANDA phase 1 and phase 2 only. For the PANDA full luminosity mode, the use of a new, next generation TDC system with higher bandwidth limits is foreseen.

In summary, 2×11 TRBv3 boards are required to readout complete STT subsystem. Two more TRBv3 boards are required to distribute the SO-DANET data. All in all, 24 optical data-links, operated at 1.2 Gbps each, are required to readout all TRBv3 boards. The TRBv3 boards can be readout by one DC module, see section 4.1.1. This allows to collect complete STT data in one FPGA and perform the following data pre-processing:

- receive and collect SuperBurst fragments;
- time-order hits-data on channel by channel basis;

- parse and decode hit data;
- add calibration constants, apply first raw hit filter (e.g. noise, pickup hits);
- find hit-track candidates;
- search for hit-cluster in multi-dim space (channel, LE-time, TE-time);
- associate hits LE-times to 250ns time slots (drift time range), $32 \mu\text{s}$ SB/128;
- associate hits TE-times to 50ns time slots;
- associate channels to sector channel slot, sector \pm neighbour sector;
- associate channels to neighbour straw/layer mapping T_0 determination;
- determine T_0 (rough) from track candidate hits;
- track pre-fit if possible from STT data alone, calculate track pre-fit parameters.

The following data-processing is performed at the CN level, where the STT data is combined with other subsystems:

- track processing;
- track building, fitting, calculate track parameters;

- determination of PID information from TOT truncated mean;
- association of tracks to events.

Taking into account that the maximum data-rate produced by STT is 160 Gbit/s, see section 3.2.2 and the fact that in phase 1 interaction rate is an order of magnitude lower all collected STT data can be pushed from STT DC to the BBN level via **two optical links**.

4.2.2.3 Barrel DIRC

Current readout scheme of the Barrel DIRC, see section 3.2.3, is based on the TRBv3 boards used as front-end. In future, the TRBv3 boards will be replaced by the DIRICH system. However, internally readout of the DIRICH system is the same as the TRBv3. Therefore, change to DIRICH readout boards will not change significantly readout scheme described in this TDR.

48 TRBv3 boards are required to readout complete subsystem, see figure 3.10. Each TRBv3 board pushes data to the DC via two optical links operated at 2 Gbit/s. Two DC modules, see section 4.1.1, are able to readout complete subsystem provided that each DC is connected to the BBN level with 10 optical links. However, for the PANDA phase 1 with reduced luminosity it is enough to readout each TRBv3 board with one optical link. Therefore, number of DC modules can be reduced to one with **two optical links to the BBN level**. In addition to the data routing the DC module will clean up of data by setting timing conditions on fired channel and time-order hits.

4.2.2.4 Endcap Disc DIRC

Front-end electronics of the Endcap Disc DIRC is briefly described in section 3.2.4. The six TOFPET ASICs of each ROM will be connected to a modified TOFPET D-Board (TOFPET ASICs and the D-Board are developed by PETsys electronics) which basically consists of an FPGA and is able to handle up to $1 \cdot 10^8$ hits/s. Eight FEE-Board (one side of the EDD dedecagon) are connected to a DAQ Board (also developed by PETsys electronics) which can handle up to $2.5 \cdot 10^8$ hits/s. In total 8 DAQ Boards are needed. Each of the DAQ boards will push data to the BBN via one optical link operated at 12.5 Gbit/s. At the phase 1 due to the reduced luminosity the occupancy of the output links will be about 1/10 of the bandwidth. This is taken into

account during the calculations of the requirements for the BBN.

4.2.2.5 Barrel TOF

The Barrel TOF subsystem, see section 3.2.5, is made of 16 Supermodules. Each of the module is connected with an optical fibre with the DC unit. Since the total data-rate, produced by the subsystem is only about 7 Gbit/s in the high-luminosity mode, the same DC unit might be employed to readout other subsystems as well, sharing the same output link to the BBN.

4.2.2.6 Electromagnetic Calorimeter

The EMC subsystem, see section 3.2.7, for readout uses DC modules described in section 4.1.1. EMC DC-unit has the following EMC-specific functionality:

- energy calibration of incoming hits;
- combine hit-data from two LAAPDs reading out the same EMC crystal (this yields data-reduction by factor two);
- combing incoming data into one time-ordered stream;
- pre-clustering and burst building for incoming data;
- slow control of the FEE, communication with the SODANET source;
- data monitoring, error detection.

The DC uses Xiling Kintex Utrscale architecture and is equipped with 60 optical links. Resources of the FPGA are enough to process data from 48 EMC RUs. Depending on the input data-rate multiple outputs towards a burst building network (BBN) level, up to 12, might be used with a speed of 12.5 Gbps. To process all EMC data 11 DC units are required with four output links toward the BBN level. For the phase 1 a single output link per DC towards the BBN is sufficient due to the reduced luminosity.

4.2.2.7 Forward Tracker

The Forward Tracker (see section 3.2.8) readout system, composed out of TRB modules (actual TRBv3 or new version described in the previous

section 3.2.8) as Data Concentrator Boards is a subsystem that has to be integrated with the PANDA data acquisition and synchronization system. In the High Resolution mode, the readout system based on TRBv3 modules requires in total 65 boards. Each TRBv3 has a direct output link, which has to be connected point-to-point to the BBN and Compute Node for event building and further processing. Since the data-rate per each link is rather low, see table 3.3 all these data can be combined into two links using DC modules, see section 4.1.1. Since the DC hardware has FPGA resources, in parallel to data combining, on-line tracking algorithms can be executed on specific data parts. Especially, tracklets from station pairs FT1, 2, FT3, 4 and FT5, 6 can be reconstructed during the event building and the calculated results replacing or appended to the raw data fragments. Finally, when all parts arrive to a single device, tracklets can be assembled into a fully reconstructed track.

In summary, only two optical fibres are used to connect the subsystem to the BBN.

4.2.2.8 Muon and other systems

The Muon detectors (see section 3.2.6) are read out via 62 optical fibres which will require two DC modules to receive and time-order the raw data. Since the incoming data-rate for the phase one is very small (2 Gbps) with respect to the available bandwidth of the employed DC modules it will be possible to employ all free inputs of the DC modules to read out other smaller subsystems as FTOF. All other subsystems as GEM will be included in a similar fashion as the ones described above.

4.2.3 Burst Building Network

From the DC level measured data have to be delivered to data processing farm of Compute Nodes via the BBN. Function of BBN is to route packets to different compute nodes according to the Super Burst number, see section 4.2.1, in a way that all fragments from a particular Super Burst end up on a particular CN module. Since no packet routing is required for connections between BBN and CN, the low-level 8b/10b encoding with defined K characters as data delimiters is sufficient in order to achieve maximum throughput over 12 Gbit/s. The BBN should have over 30 optical links as inputs from the DC level and about 8 outputs to the CN level, see figure 4.27. The number of outputs in this case are determined by the required number of CN modules required for online data processing,

and not by the total bandwidth of the incoming links since the expected total data-rate is 80 Gbit/s, see section 3.2.10. In case if number of required CN modules is less than 8 it is possible to connect each CN with BBN with few links. The BBN can be implemented as a virtual network switch in one FPGA. As hardware platform a DC module, see section 4.1.1, can be used since it has enough resources. Moreover, using the same hardware for DC and BBN tasks eases DAQ hardware development and maintenance. Functionally, for each incoming packet of data the BBN switch assigns destination CN-module and transfers this packet accordingly. Therefore, next to the network-switch functionality the BBN module contains supervisory functionality that monitors status of the system and changes the CN assignment, in case a CN-module stops properly function.

The Compute Node farm performs event building and the first level of event filtering. As mentioned in the text above, the CN module is reassembling data packets with the same Super Burst number coming from different data concentrators. It might be that latency of data-processing within the DC level is different for different subsystems (due to different complexity of data pre-processing). Therefore, the CN module should have sufficient memory to buffer the incoming data until all packets from all subsystems are collected. Once complete Super-Burst data is collected it is being processed by CN. Number of required CN modules is determined by complexity of the data-processing algorithms. One of the most resources demanding algorithm is track finding and track fitting procedures discussed in section 4.3.1. Assuming that during the PANDA experiment an upgraded version of CN modules, based on Xilinx Kintex-7 Ultrascale FPGA family, will be employed the complete track finding and fitting can be done by five full ATCA CN-blades (20 AMC cards, see conclusion of the section 4.3.1) for the high-luminosity mode. Since during the phase 1 luminosity is lower by a order of magnitude we assume that five ATCA CN-blades should be sufficient for implementation of all required algorithms.

The processed data from all CN modules is being send using point-to-point optical links to the High-flex readout boards, see section 4.1.3. The readout board transfers data via InfiniBand to HPC computing for further data processing.




Figure 4.27: The overview of the PANDA readout system starting at the Data Concentrator (DC) level.

4.2.4 DAQ partitioning, implementation

As mentioned in section 3.4 it should be possible to run DAQ in different modes. Different modes of operation are defined by topology of connected hardware and/or configuration of the burst-building network and online data-processing. As example of different modes of operation one can consider data-taking run at FAIR facility. In this case DAQ has to collect all data-packets with the same super-burst number in one compute node for event building. Topology of connected hardware is represented in figure 4.27. With the data-taking run data-packets with the same super-burst number are sent to the same CN number $\text{mod} (SB/N_{CN})$, where SB is the current super-burst number and N_{CN} – number of running CN modules. In compute nodes data is pre-processed and sent to HPC cluster. However, the same hardware configuration can be used to test separate detector subsystems. In such case data-combining algorithms in compute nodes is modified in order to combine data of separate subsystems in a separate packets and after required pre-processing send data to HPC. In case of stand-alone DAQ data from DC modules are sent directly to high-flex readout card and stored/processed on a local PC. In case data-preprocessing in CN is necessary DC modules can be directly connected to CN modules (if number

of physical data links is sufficient) and functionality of the burst-building network is implemented in the CN FPGA. Such modular design of the DAQ and usage same family of FPGAs at the DC, burst-building, CN and high-flex modules allows to build DAQ for different experiments, from test experiment of a single detector to complete PANDA experiment, using same hardware and library of pre-compiled data-processing ip-cores.

4.2.5 High-speed ROCE links and embedded GPU/AI computing nodes

Future experiments require a new concept of data acquisition (DAQ) systems to handle the largely increased detector data rates. In general, fast standard communication protocols and off-the-shelf devices should be used where possible and need to be integrated with the custom detector readout. The use of standard networking protocols, i.e. TCP/IP, RDMA, RoCE, iWARP and commercially available PCIe-based data receiving and processing cards optimizes cost and effort and ensures seamless incremental upgrades of individual subsystems with technology progress. Moreover, the recent accelerator cards can be programmed in high-level languages, what facilitates migration of analysis algo-




Figure 4.28: The overview of the PANDA readout system employing commercial hardware at the burst-building level.

rithms implementations from simulations to actual data processing and improves source code compatibility and maintenance. It is especially important for projects with limited human resources. The key feature of this DAQ framework is the escape from custom electronics and protocols at the stage of subsystems data concentrators and proceed with the networking, Burst Building and data processing using commercially available hardware, firmware and software. One of such methods is RDMA (Remote Direct Memory Access) that realizes the transfer of data from the memory of one host or device into that of another one without any CPU intervention, allowing high-throughput. We are exploring the advantage offered by recent Ethernet standards which are nowadays available RoCE (RDMA over Converged Ethernet) [9] and iWARP [10]. In particular the RoCE standards offers several advantages: it is a non-complex UDP/IP protocol based on the Infiniband, higher RDMA performance over 100G, direct implementation on modern FPGA by the Xilinx Embedded Target RDMA enabled NIC (ETRNIC) IP-Core [11]. It will enable a direct connection to the high-performance computing (HPC) system for fast data processing including edge AI technology. The Burst Building might be performed on the network packets switching level. Complete Super-Burst data, in a form of a number of data units coming from the PANDA subsystems will be delivered through the network to the HPC node, which are equipped with GPU and FPGA accelerator cards

(i.e. NVIDIA EGX [12] GPU and Xilinx Alveo). The node will delegate particular data analysis algorithm implementation to a suitable hardware resource in order to achieve highest possible burst processing throughput, see figure 4.28.

4.3 Event filtering and partitioning of algorithms

One of the main tasks of the DAQ is data reduction by filtering events of interest from the background. In case of phase one/two experiments the required reduction factor is 100. In section 5.1 it is demonstrated that the required reduction factor can be achieved by event filtering based only on online tracking and calorimetry with reduced resolution. Therefore these two online algorithms are considered crucial components of the DAQ and are described in this section.

4.3.1 FPGA based online tracking algorithm for the STT

FPGAs as a platform to run tracking algorithms have been employed in many experiments. In the D0 experiment at Fermilab, the central track trigger logic is implemented in FPGAs which compare hit inputs to a predefined list of track equations. In the CDF experiment, the online tracking sys-

tem performs pattern recognition with Associative Memories and a linear track fitter implemented in FPGAs. In the CMS experiment, the resistive plate chambers based trigger system implements a tracking algorithm in FPGAs, detecting cosmic muon tracks. In the ATLAS experiment, the Fast Tracker system as a part of the ATLAS trigger upgrade programme performs global tracking at 100 kHz. Here, FPGA devices are employed for track fitting calculations.

Employing the full PANDA STT detector with 4636 straw tubes, an FPGA based tracking algorithm has been designed with a complex track finder, followed by a track fitter. Due to the high event rate of 1 MHz, some overlapping of events might occur and has to be treated properly. The algorithm features a high tracking efficiency, low latency and sufficient momentum resolution for event filtering in realtime.

Operating in a freely streaming mode, the quasi-continuous beam delivered to PANDA is organised in bursts of 2000 ns duration, with a 400 ns gap between two bursts. Assuming that the event start time T_0 is provided by a fast external detector (e.g. inner detector as MVD 3.2.1, or outer as Barrel TOF 3.2.5), a collection of hits from T_0 to T_0+200 ns in a burst is packed as one STT burst event, potentially containing some overlapping proton-antiproton interactions. The drift time of each hit in this burst event is approximated by the difference of the measured arrival time of the hit and T_0 , with the flight time (a few ns) of the particle from the interaction point to the tube being ignored.

At 1 MHz interaction rate, an average of 2 events are expected in one burst. Due to the large electron drift time (up to 200 ns), the STT hits from the events in one burst might have some overlapping with respect to the arrival time. One obvious feature of the overlapping of events is the increase of the number of hits in a burst event, which then introduces a number of fake tracks due to the overlapping events. Due to the 400 ns gap between two bursts, there are no overlapping of events in two bursts.

With the drift time information, the drift circles of hits in the burst event are calculated and fed to the tracking algorithm which is divided into two stages, the track finder and the track fitter.

4.3.1.1 Track finding

The procedure of track finding is depicted in Figure 4.29. In this figure, the black solid circles repre-

sent straw tubes. The red solid curve indicates the trajectory of a charged particle. In the tubes on the trajectory, the drift circles calculated using T_0 are shown as red dashed circles. The track finder starts from hits in the innermost layer, marked as H_0 . The hits in H_0 are considered as a seed and are used to create a tracklet. Moving to the next layer, two adjacent tubes, tube 1 and 2 in the second layer are searched and the active tube is attached to the tracklet. If both adjacent tubes are not active, the next-to-adjacent tubes, 3 and 4, will be searched. With the attached hit in this layer, referring as H_1 , the tracklet is updated and extrapolated to the next layer, till the last layer of the inner axial tubes.

The track finder starts to enter the stereo layers. At the transition from axial to stereo layer, a wider search window of ± 6 tubes is applied. Inside the stereo layer, a normal search and attach process will be performed with 2 adjacent and 2 next-to-adjacent tubes.

After finishing the stereo layers, the track finder enters the outer axial layers. As in the previous step, a wider search window of ± 6 tubes is applied at the transition from stereo layer to axial layer. Using this track finder, tracklet candidates consisting of a list of axial and stereo straws are formed. The track finder is tolerant towards a missing layer. However, in the case that no hit is found in two consecutive layers, the track finder quits the current track finding. A track candidate is only delivered to the next step when there are at least 3 axial hits and 2 skewed hits. The neighbouring information is kept in the identification number of a tube. For each layer, the maximum tube number is saved in one array. A tube, with a number larger than 0 and smaller than the maximum indicates that it is not located at the border of the segment, for which the neighbouring tubes can be located by -1 or $+1$ operation to the tube number. When a tube is at the border, the -1 or $+1$ operation will jump to the neighbouring tube at the neighbouring segment. This is fulfilled by using a map of neighbouring tube at the segment border. With this, tracks passing through different STT segments are found correctly. Two crossover tracks will produce 2 good tracklets and 2 combinational tracklets. These combinational tracklets can be suppressed later with informations of track qualities.

4.3.1.2 Track fitting

With a group of hits in one tracklet obtained from the track finder, a track fitter is designed to extract the helix parameters of the track. Here, a least




Figure 4.29: Illustration of the track finding algorithm. The black circles represent the individual straw tubes. The red solid curve indicates the trajectory of a charged particle. The red dashed circles inside tubes represent the isochrones.

squares fit is used based on minimising the mean square distance from the fitting curve to the data points. The axial hits are used for the measurement of the helix curvature related to the transverse momentum. In the XY plane, a helix is projected to a circle. Rather than the coordinates of the hit in three dimensions, each measurement of a straw tube provides a drift circle. A fit fully considering drift circles by minimising the mean square distance from the fitting curve to drift circles would make the calculation too complicated and not suitable for implementation in VHDL. Thus, a simplification by using the central position of each drift circle with the reversed drift radius as a weight is performed. By this method, a 6 % transverse momentum resolution is achieved at 1 GeV/c.

To improve the resolution, a second fit iteration is applied. Using the extracted momentum from the previous fit, the relation between each drift circle and the track can be determined. The center of a drift circle is either inside or outside of the helix circle in equation 1. With this information, a point on each drift circle is chosen as the best conjecture of the trajectory. Using these points selected at drift circles, the second fit iteration improves the resolu-

tion to 3 % at 1 GeV/c. No significant improvement of resolution was observed with further iterations.

After this stage, three of the five parameters of the track helix are known (the radius R, the position of the helix center in the X-Y plane). The remaining two parameters of the helix are determined by using the hits of the skewed straws.

4.3.1.3 Implementation in VHDL

To test the tracking algorithm in the FPGA chip, an experimental test bench was set up, with one PC connected to an Advanced Telecom Computing Architecture (ATCA) based compute node (CN) [13] via a GBit Ethernet link. In Fig. 4.31, a first generation CN prototype is shown, which was used for the tests described below. The PC is emulating the detector by sending straw tube hit information from Monte Carlo simulations to the tracking algorithm. The CN which is designed as a generic solution for trigger and data acquisition was equipped with 5 Virtex 4 FX60 FPGA chips. It is plugged into an ATCA shelf, capable of supporting up to 14 CNs. The features of CN can be found in [13].

The algorithm has been implemented in VHDL. In




Figure 4.30: Example of a single track displayed in X-Y view (left) and Φ -Z view (right). In the Φ -Z view, the blue solid line represents the fit result from the first iteration using central position of each hit. The red dashed line shows the fit result after the second iteration.




Figure 4.31: The Compute Node used in this study, equipped with 5 Virtex 4 FX60 FPGA chips.

number and 6 bits for the tube number.

The simplified block diagram of the implementation is shown in Fig. 4.32. The data sources for the algorithm include a collection of STT hits in a burst and a list of event start times T_0 .

5




Figure 4.32: Simplified block-diagram of the algorithm in VHDL.

the code, fixed-point numbers are used. The hit information, such as the position x y z of a wire, is expressed as 24 bit fixed-point numbers, with the first bit as a sign bit the following 7 bits as the integer part and the last 16 bits as the fraction.

Each tube has its own identification number, with 3 bits for the segment number, 5 bits for the layer

5

Following the data flow, the algorithm is comprised of:

- 1) The Input Interface module taking charge of reading the STT hits. Hits in the window of T_0 to $T_0 + 200$ ns are packed as one burst event which is copied to a ring buffer.

10

2) The ring buffer is a circular FIFO with a depth of 1024 and a width of 96 bits, composed of 72 bits for position x y and z , 14 bits for the ID of the tube and 10 bits for the arrival time of the hit in the burst. Typically, a burst event has 100 hits. The ring buffer is designed to buffer one event for a certain time. After the hits of one event are filled into the buffer, the input interface stops sending further hits, and the indices of the first hit and the last hit in the ring buffer will be kept. When the tracking of this event is finished, the hits of next event will be sent to the buffer, with the indices of first and last hits updated.

3) The 2D mapping module contains a dual port block RAM, with a depth of 16384 and width of 16 bits. Each tube has one correspondence in the 2D map. The 14 bit ID of a tube is taken as the address line of the block RAM. Reading one hit, its corresponding bin in the 2D map is marked as "Occupied". At the last hit, the 2D mapping task is finished and a start signal is passed to the next module, the Track Finder.

4) Once a start signal is received, the Track Finder starts to work. Hits from the innermost layer are saved in one array and treated as seed. Starting from each seed, the Track Finder searches its neighbours in next layer, which is fulfilled by sending the address line (ID of a tube) backward to the 2D map. If the unit in 2D map is "Occupied", the corresponding hit is attached to the tracklet. When one tracklet is completed, a start signal is sent to the next module, the Pt Calc.

5) The Pt Calc module reads the hit index from each tracklet and calculates the transverse momentum. There are two input data streams, one from the Track Finder with only the index of hits in a tracklet, and another from the ring buffer containing the full information of corresponding hits. 11 Xilinx IP (Intellectual property) cores for 32 bit multiplication are used, which have a latency of 6 clock cycles. At the end of the first iteration, the module updates the position of each hit by choosing one point around the drift circle and starts the second iteration. When the second iteration is finished, a signal is sent to the Pz Calc module.

6) The Pz Calc is designed to calculate the Pz information in two iterations. The transverse momentum from the previous module is read in to calculate the intersection with skewed straws.

These modules are optimised to operate in a pipelined mode. The first 3 modules, the input interface, the ring buffer and 2D mapping modules work hit-based, with one clock cycle required for one hit. In contrast, the last 3 modules, the track finder, the Pt Calc and Pz Calc modules work tracklet-based. Once a tracklet is found by the Track Finder, the following Pt Calc module starts immediately while the Track Finder continues with the search for the next tracklet. This reduces the execution time. About 700 clock cycles ($7 \mu\text{s}$) are required for one event with 6 charged track (about 100 hits). The device utilisation summary table for the design is shown in Fig. 4.33.

Device Utilization Summary			
Logic Utilization	Used	Available	Utilization
Number of Slice Flip Flops	25,022	50,560	49%
DCM autocalibration logic	14	25,022	1%
Number of 4 input LUTs	33,120	50,560	65%
DCM autocalibration logic	8	33,120	1%
Number of occupied Slices	21,563	25,280	85%
Number of FIFO16/RAMB16s	148	232	63%
Number used as RAMB16s	148		
Number of DSP48s	124	128	96%

Figure 4.33: Device utilization summary for the design.

4.3.1.4 Performance Results

In the investigation of the performance of the algorithm, we first run the algorithm with well separated events, corresponding to very low event rate. Then tests at different event rates are performed. An example of a single event is shown in Fig. 4.34.




Figure 4.34: Example of a single event. The red circles represent the tracks found by the algorithm.

At sufficiently low event rate, events are well separated with no overlapping occurring. Event-based

simulation, which does not consider overlapping events is used to generate Monte Carlo (MC) samples in the framework of PandaRoot [14].




Figure 4.35: The tracking efficiency and the resolution as a function of the transverse momentum.

tracking efficiency is defined as $\epsilon = N_{rec}/N_{gen}$, where N_{rec} is the number of tracks reconstructed and N_{gen} is the number of generated tracks which has at least 3 hits in the STT. A reconstructed track is required to be within 6 standard deviations with respect to the momentum of the simulated one.

The tracking efficiency is studied at four transverse momentum bins, as shown in Fig.4.35. At high transverse momentum (with $p_t \geq 2$ GeV/c), a drop of efficiency is observed. This is due to the reduced bending of tracks at high transverse momentum.

A track, appearing like a straight line originating from the interaction point has a larger probability to trigger tubes which are at the same side of its trajectory, which makes the momentum extraction to be more difficult. The resolution $\sigma(p_t)/p_t$ at different momentum bins is shown in Fig.4.35. The resolution $\sigma(p_z)/p_z$ depends on p_t because p_z is calculated with the input of p_t . In the case of $p_t = 1$ GeV/c, $\sigma(p_z)/p_z$ at 0.5 GeV/c (1.0 GeV/c, 2.0

GeV/c) is 3.1% (3.8%, 4.2%).

The hit information from the prepared MC samples is transmitted into the connected FPGA chip, in which the track finder and track fitter are run accordingly as explained in the previous section. The resulting momentum information from the algorithm is send back to the PC. The top plot in Fig.4.34 displays one example event that was analysed by the FPGA.

To summarize this section, an FPGA-based tracking algorithm has been designed for the straw tube tracker and tested on a Virtex 4 FX60 FPGA. The algorithm features high tracking efficiency, sufficient momentum resolution (3% (4%) for p_t (p_z) at 1GeV/c) and low latency (7 μ s to process one event with 6 tracklets). Comparing to a CPU chip (single core Intel Xeon E5520 at 2.26 GHz), where the algorithm was implemented in C++, about 40 ms are required to process one event. An improvement of more than a factor of 5000 for the specific FPGA implementation compared to the CPU implementation is obtained. The expected performance boost using the CN with Kintex 7 Ultrascale would be approximately 10 which results in a requirement of 20 FPGAs for the tracking task.

4.3.2 FPGA based online clustering algorithm for the EMC

This section is a part of PhD thesis [15]

Cluster Finding - Input

As an input, the cluster-finding algorithm takes a stream of EMC digis grouped in time (see the start of the flow diagram in Figure 4.38). 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 4.36. 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.




Figure 4.36: 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.




Figure 4.37: 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.

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.

Distributed Cluster-Finding

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 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 4.37), 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.

Cluster Finding - Algorithm


The art of cluster finding is not new. The algorithm that is currently implemented in the Panda-ROOT 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 cluster-finding process is depicted graphically in Figure 4.38. 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, skips¹ one or more crystals, and starts a new particle shower in a neighbouring crystal, or
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 are discussed in [15].

1. Here, 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.



Set cluster properties

Assign hits to clusters

Hit	Cluster
1	1
2	1
n	...

$E_{clus} = \sum E_{hit}$
 $t_{clus} = t_{hit(max E)}$
 ...

1	2	3	4

Figure 4.38: 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.

Algorithm Implementation

The algorithm described above takes the input time-bunch and loops over all member digis to identify clusters, as shown in Figure 4.38. For distributed cluster-finding, there are two options:

- In the single-pass case, the algorithm takes the input time-bunch, and loops per (virtual² or real) 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, 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.

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 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. All algorithms described above follow the flow chart in Figure 4.38, 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.

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 of the PandaRoot simulation package, as the forward-boosted particle production forces most decays in that direction. The 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

² A 'virtual' Data Concentrator (DC) is a subsection of the EMC in the simulation, that would correspond to one DC.




Figure 4.39: 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. The right figure shows only the results of the time-based simulation.

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 4.39 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 4.39. 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 4.40. 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 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.

For a more advanced test, the four cluster-finding algorithms are pitted against each other for the two selected decay channels, namely $p\bar{p} \rightarrow \gamma\gamma$ and $h_c \rightarrow \eta_c\gamma \rightarrow \pi^0\pi^0\eta\gamma \rightarrow 7\gamma$ and five interaction rates 2 kHz, 200 kHz, 2 MHz, and 20 MHz. The




Figure 4.40: 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.

selected decay channels were chosen because all decay products a gamma rays and reconstruction of these channels solely depends on the performance of the EMC subsystem. Since 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 $h_c \rightarrow \eta_c\gamma \rightarrow \pi^0\pi^0\eta\gamma \rightarrow 7\gamma$ 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 4.41 and Figure 4.42 for the two decay channels. The cluster-finding algorithms are referred to by ‘DEF’ for the currently implemented in the PandaRoot default algorithm, ‘ONL’ for the online algorithm, ‘DIST’ for the single-pass distributed method, and ‘2STEP’ for the double-pass distributed 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 4.41 and 4.42, 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 4.43. 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 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.

5

10

15

20

25

30

35

40

45

50

55




Figure 4.41: 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. 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 4.42: Same as Figure 4.41, 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 the text.




Figure 4.43: 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.

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 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.

4.3.3 EMC clustering. Performance Assessment by Hardware Simulation

4.3.3.1 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




Figure 4.44: 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.

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 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 4.44. First, digitised hit (`digi`) data were simulated for the $h_c \rightarrow \dots \rightarrow 7\gamma$ decay channel using PandaROOT. As mentioned above, this channel was selected to test solely performance of the EMC subsystem. The simulated, digitised data were reformatted to mimic the data coming out of the SADCs. These data were mapped onto virtual Data Concentrators, and `digi` data from four adjacent DCs were written in the correct data format (see section 4.2.1) using a LabView program. Using the TRB protocol [4], 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 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




Figure 4.45: 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.

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, Beijing, China.

The result from the VHDL implementation is shown in Figure 4.45 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 performance can be optimised using simulated data. In addition, the DCs appear to have sufficient resources to run the precluster forming algorithm.

4.3.3.2 Simulation of the Burst Building Network

Next, a virtual version of the Burst Building Network 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 this work was considered that the burst building network is constructed with hardware modules




Figure 4.46: 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.

equipped with 24 optical links. Therefore, the complete burst building network employed multiple modules arranged into two levels L1 – collecting data from the data concentrators, and L2 – transferring data to the compute nodes. Each module from the L2 layer has a direct link with all modules from the L1 level. 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 equipped 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 4.46. 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 4.46, which draw different data from the same eight inputs of the module, based on their superburst³ 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 cores, 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 L2

³. 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 in this case to 32 μ s.

nodes) \times (2 modules) \times (4 outputs)=(64 CNs) were simulated. Data are collected in the L2 nodes, based on their superbunch number, and each data packet with a given superbunch number is sent sequentially to one of the 128 Compute Nodes. Data with superbunch number 1 are sent to the first CN, data with superbunch number 2 to the second, and so on. Hence, the time available to a CN to process the data, until the next superbunch data packet is offered, is $128 \times 32 \mu\text{s} = 4.1 \text{ ms}$. In the current simulation, half of the superbunches 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 [15]. 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 superbunch 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 4.1 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 inspires confidence that the envisaged network will be able to handle the data stream. In later stages, it is 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.

4.4 Run Control, error handling and data quality monitoring

All DAQ actions as preparation, start and stop of data taking, are triggered by the Experimental Control System (ECS) via the Run Control System (RCS). After receiving triggers the DAQ performs all required operations in a close cooperation with the Detector Control System (DCS). Such cooperation can be implemented according to two possible principles [16]:

- DAQ receives a list of commands from DCS via dedicated EPICS IOC. This list is executed by the DAQ, e.g. configured all FEE and DAQ hardware, check hardware status. After successful execution of all commands DAQ reports back to the DCS and RCS. In this case all FEE and DAQ parameters are handled by the DCS.
- DCS provides an access to a database where all required parameters are stored and the DAQ retrieves all required parameters from the database and applies them to the system.

This can be implemented as:

- DAQ provides a shared library which handles all DAQ internals and is being used by DCS;
- DAQ provides a control daemon which communicates with DCS and RCS.

In both cases DAQ internally uses SODANET protocol, see section 4.2.1.

Hardware of the PANDA DAQ is designed with access bandwidth and buffers which guarantees fault-free operation under all conditions. Problems as overflow buffers or missing data occur only under condition of faulty hardware. Therefore, in case of fault DAQ has to be reconfigured to omit faulty hardware which requires action of DCS and run control. Because of the slow nature of such operations PANDA DAQ does not foresee error reporting with fixed latency and, therefore, error reporting and handling is being performed by DCS. As described in section 4.2.1.5 all DAQ components have dedicated registers where all errors are recorded. Value of these registers is periodically pulled by DAQ/DCS and appropriate action is taken in case of detected error. In case the DAQ is not able to process all incoming data online, e.g. due to disables online data-filtering (taking unbiased data) or

major malfunction of several CN modules it is possible to reduce amount of incoming data by enabling throttling. Throttling is implemented by periodically enabling and disabling data taking on sequential super-bursts by sending stop-DAQ and start-DAQ SODANET commands. In this way it is possible to skip required amount of data.

Data-quality monitoring is performed on two levels: raw-data and reconstructed data. The Same registers which record all occurred errors store information on overall status of the DAQ, namely continuously monitoring if all hardware, which is listed by DCS as working, is delivering data for every super-burst. These registers allow continuously monitor quality of delivered raw data. Next to these monitoring parameters online computing provides to RCS high-level reconstructed quality-monitoring data, which summarize performance of all enabled subsystems.

Another aspect of error handling is related to the hardware. If necessary, during the maintenance periods firmware of used electronics is upgraded. Any fault during upgrade of the firmware will render electronics to be unusable. In case a bad image is flashed into the FPGAs, the boards have to be accessed and re-programmed manually, which is a tedious and time-consuming work that should be avoided. Therefore, all used electronic modules employ one of the following mechanism which guaranty fail-safe firmware upgrade:

- The FPGA operates under an external supervisor managing the communication with the SODANET network, e.g. an interface FPGA or CPLD. In case the main FPGA does not boot correctly, the communication FPGA takes control over the flash memory and re-program a working firmware image. If the firmware of the communication FPGA or CPLD stays untouched, it is possible to access the flash memory over the network even if the main FPGA does not boot.
- An external watchdog circuit switches the main FPGA to an alternative "golden" image, in case it does not boot and correctly initialize after a given timeout.

Bibliography

- [1] M. Caselle et al. *Jinst*, 12:C03015, 2017.
- [2] M. Kavatsyuk et al. Trigger-less readout electronics for the PANDA Electromagnetic Calorimeter. *Conference Record, 2011 IEEE Nuclear Science Symposium, Valencia, Spain, October 23 - 29, 2011*, DOI 10.1109/NSSMIC.2011.6154360.
- [3] I. Konorov et al. Trigger-less readout electronics for the PANDA Electromagnetic Calorimeter. *Conference Record, 2009 IEEE Nuclear Science Symposium, Orlando, Florida, USA, October 25-31 2009, N25-249*, DOI 10.1109/NSSMIC.2009.5402172.
- [4] J. Michel. Ph.D. Thesis, Johann Wolfgang Goethe-Universität, Frankfurt am Main. <https://hds.hebis.de/ubffm/Record/HEB314568271>, 2012.
- [5] HADES experiment. <http://www-hades.gsi.de>.
- [6] CAEN V775 32 Channel Multievent TDC (35 - 300 ps). <http://www.caen.it/csite/CaenProd.jsp?parent=11&idmod=>
- [7] J. Adamczewski, H. G. Essel, N. Kurz, S. Linev. Data acquisition backbone core DABC. *IEEE Transactions on Nuclear Science*, 55(1):251–255, 2008.
- [8] J. Adamczewski, et. al. Go4 on-line monitoring. *IEEE Transactions on Nuclear Science*, 51(3):565–570, 2004.
- [9] Mellanox Technologies, "RDMA/ROCE solutions", [Online] Available: <https://community.mellanox.com/docs/DOC-2283>.
- [10] Intel Corp., "Understanding iWARP: Delivering low latency Ethernet", [Online] Available: <https://www.intel.fr/content/www/fr/fr/-ethernet-products/gigabitserver-adapters/iwarp-brief.html>.
- [11] Xilinx RDMA IP-core, https://www.xilinx.com/support/-documentation/ip_documentation/-etnric/v1_0/pg294-etnric.pdf.
- [12] <https://www.nvidia.com/en-us/data-center/products/egx-a100>.
- [13] H. Xu et al. *Physics Procedia*, 37:1849, 2012.
- [14] S. Spataro et al. *J. Phys. Conf. Ser.*, 331:032031, 2011.
- [15] M. Tiemens. Ph.D. Thesis, University of Groningen, Groningen, The Netherlands. 2017.
- [16] The PANDA collaboration. Technical Design Report for the: PANDA Detector Control System. 2019.

5 Performance

5.1 Simulations

The PANDA online event filter aims to select interesting signal events with a high efficiency while reducing the total data rate to a feasible amount. It works without conventional hardware triggering but relies on prompt online reconstruction of particle and event candidates to apply complex and sophisticated filtering algorithms. In general it should be possible to trigger for multiple reactions simultaneously.

In order to demonstrate the principle feasibility of this approach foreseen for the PANDA data acquisition a dedicated set of benchmark channels has been studied. The channels are chosen to cover several aspects of the PANDA physics programme, namely hyperon physics at high production cross sections, charmonium physics a low cross sections, and form factor physics at very low cross sections. The primary goal in all cases is to retain a reasonable efficiency for the signal events while reducing the total data stream by at least two and three orders of magnitude for the high and low cross section cases, respectively.

Concretely the channels under consideration are:

- (1) $\bar{p}p \rightarrow \Lambda^0(\rightarrow p\pi^-)\bar{\Lambda}^0(\rightarrow \bar{p}\pi^+)$, at $p_{\bar{p}} = 1.641$ GeV/c, $E_{cm} = 2.304$ GeV. This channel is very important to study hyperon spin-observables allowing for probing QCD in the confinement domain. It has a quite high production cross section of $\mathcal{O}(100\mu\text{b})$ [1] with Λ having a large branching fraction $\mathcal{B}(\Lambda \rightarrow p\pi) \approx 64\%$ to the final state under study [2].
- (2) $\bar{p}p \rightarrow J/\psi(\rightarrow e^+e^-)\pi^+\pi^-$, at $p_{\bar{p}} = 6.988$ GeV/c, $E_{cm} = 3.872$ GeV. This channel is one of the key reactions to study and eventually understand the nature of the charmonium-exotic candidate $X(3872)$ that is discussed to be possible a multi-quark or molecule state [3, 4, 5]. It has a rather low cross-section of the order $\mathcal{O}(50\text{nb})$ [6] and thus is in combination with the smaller branching fractions $\mathcal{B}(J/\psi \rightarrow \ell^+\ell^-) \approx 12\%$ challenging to be isolated from multi-hadron backgrounds.
- (3) $\bar{p}p \rightarrow J/\psi(\rightarrow \mu^+\mu^-)\pi^+\pi^-$, at $p_{\bar{p}} = 6.988$ GeV/c, $E_{cm} = 3.872$ GeV. The same channel as (2) with the J/ψ decaying to the muonic final state.

(4) $\bar{p}p \rightarrow e^+e^-$, at $p_{\bar{p}} = 1.5$ GeV/c, $E_{cm} = 2.256$ GeV. This channel being complementary to elastic electron-proton scattering offers the possibility to study the electric G_E and magnetic G_M form factors of the proton in the region of $q^2 > 0$ (time-like region) [7]. In this region, the precision of the proton form factor measurements at the $\bar{p}p$ annihilation experiments and e^+e^- colliders has been limited [8]. Feasibility studies for the measurement at PANDA [9] show that $|G_E|$, $|G_M|$ and their ratio can be determined over a large q^2 range with a high precision. The cross section of this reaction is about $\mathcal{O}(1\text{nb})$, where the largest challenge is to separate it from $\bar{p}p \rightarrow \pi^+\pi^-$ background by mainly particle identification quantities.

(5) $\bar{p}p \rightarrow e^+e^-\pi^0(\rightarrow \gamma\gamma)$, at $p_{\bar{p}} = 1.5$ GeV/c, $E_{cm} = 2.256$ GeV. This channel, originating from $\bar{p}p \rightarrow \gamma^*\pi^0$ offers, in addition to the previous one, the unique opportunity to study the proton time-like form factors in the unphysical region below the kinematical threshold of the $\bar{p}p$ production of $(2M_p)^2$ [10, 11, 12]. This region has never been experimentally accessed. The cross-section is supposed to be larger than the previous one by about one order of magnitude [13, 14]. Due to the additional π^0 and the less particular kinematic situation it poses even higher challenge to a clean selection.

For the filtering algorithm we assume that that track finding, track fitting and neutral cluster finding and cluster/track matching was performed online to form event candidates, where also basic quantities for particle identification (e. g. dE/dx from the tracking system) are available being assigned to the reconstructed particle objects. The luminosity in the starting Phase-1 of up to $\mathcal{O}(2 \cdot 10^{31}/\text{cm}^2\text{s})$ suggests that event mixing can be neglected for now [REF].

In the following it is demonstrated that this goal is met for all cases individually, i. e. the background can be reduced by the required factor when filtering for a certain signature. The effect of simultaneous filtering for multiple channels will be discussed as well.

Table 5.1: Efficiencies for signal and background events after combinatorics.

Channel	ϵ_S [%]	ϵ_B [%]
$\Lambda \rightarrow p\pi^-$	60.6	58.0
$J/\psi \rightarrow e^+e^-$	42.2	1.1
$J/\psi \rightarrow \mu^+\mu^-$	62.3	0.9
$\bar{p}p \rightarrow e^+e^-$	61.1	2.1
$\bar{p}p \rightarrow e^+e^-\pi^0$	45.4	6.1

5.1.1 Framework

The simulations have been performed with the PandaRoot software framework [REF]. It is a Geant based microscopic simulation framework with detailed detector descriptions and a realistic reconstruction response. The used software is PandaRoot rev. #c757736, working with FairRoot v.17.10b and FairSoft oct17p1. The MC data have been produced with the complete offline simulation and reconstruction chain.

For each benchmark channel 10^5 signal events with the EvtGen generator [EvtGen] and 10^6 background event with the Dual Parton Model (DPM) generator [DPM] have been simulated and reconstructed, respectively. Since final algorithms for the online reconstruction were not available for tracks and neutral clusters found in the electromagnetic calorimeter at the time of writing the expected performance for the resolution figures were effectively simulated by smearing the track momenta and cluster energy deposits according to numbers from preliminary developments [REF]. Tracks have been smeared by a Gaussian distribution to roughly match $dp/p \approx 3\%$, the dE/E for clusters has been broadened by about a factor two.

5.1.2 Results

Based on events composed of reconstructed tracks, neutrals clusters and PID information available online, particle combinatorics has been performed for each channel individually. Either the main intermediate resonance ($\Lambda \rightarrow p\pi^- +$ charge conjugates, $J/\psi \rightarrow e^+e^-$, $J/\psi \rightarrow \mu^+\mu^-$) or the full exclusive final state (e^+e^- , $e^+e^-\pi^0$) has been reconstructed from the appropriate combination of charged track candidates and neutral particle candidates without usage of PID at this initial stage. The numbers of events S (not the numbers of entries in the histograms) contributing to the corresponding invariant mass distributions in a certain region of inter-

est were considered as primary observables. The efficiencies for signal events were then determined as $\epsilon_S = S/S_0$ with $S_0 = 10^5$ being the number of generated events. The background efficiencies were analogously determined as B/B_0 with B and $B_0 = 10^6$ defined in the same way as for signal events.

Figure 5.1 shows the invariant mass distributions in the regions of interest for the five channels after combinatorics. The black histograms are those of signal events, the red hatched areas the ones for background events. In all cases the signal show a clear peaking enhancement at the expected position. For channels (1) – (3) (Figs. 5.1 (a) – (c)) with the inclusively reconstructed resonances the background distribution does not show any peaking structure. For channels (4) and (5) ((Figs. 5.1 (d) and (e)) there is peaking background visible due to the exclusive reconstruction of the reactions. The predominant background reactions for the latter two channels are $\bar{p}p \rightarrow \pi^+\pi^-$ and $\bar{p}p \rightarrow \pi^+\pi^-\pi^0$ with the charged pions erroneously taken as the electron and positron candidates.

At this stage the signal efficiencies are with $\approx 45\%$ – 60% sufficiently large for all channels. The background efficiency without any further filtering are far below 10% for channels (2) – (5) as shown in Tab. 5.1. Channel 1. suffers from a massive background level of more than 50% due to the Λ resonance being close to the threshold. Many pairs of low momentum particle tracks have an invariant mass in that region if assigned this combination of particle hypotheses. However, to reach the target of a data stream reduction (background suppression) of two to three orders of magnitude, further filtering is necessary for all channels.

In order to ensure feasibility of the online filtering process in this proof-of-principle demonstration, mainly rather simple quantities were used.

The distributions of these quantities for signal (black histograms) and background events (red hatched histograms) as well as the applied selection (blue dashed lines) are visualized in Fig. 5.2. The resultant invariant mass distributions and efficiency numbers for the five reactions under study after application of these selection criteria are displayed and compiled in Fig. 5.3 and Tab. 5.2, respectively, and in detail discussed below.

- (1) $\Lambda \rightarrow p\pi^-$ (and c.c.) Due to the rather high initial background level, the filtering of this decay requires more restrictive selection criteria applied to multiple quantities with partially higher complexity. The quantities to select




Figure 5.1: Invariant mass spectra of the relevant signal resonances in the regions of interest after combinatorics for the five benchmark channels: (a) $\Lambda^0 \rightarrow p\pi^-$ (and charge conjugates), (b) $J/\psi \rightarrow e^+e^-$, (c) $J/\psi \rightarrow \mu^+\mu^-$, (d) $\bar{p}p \rightarrow e^+e^-$, (e) and $\bar{p}p \rightarrow e^+e^-\pi^0$. The black hollow histograms show the distributions for signal events, the red hatched histograms the corresponding distributions for background (DPM) events.

on were identified by a semi-automatic methode. Many kinematic and PID specific variables have been studied, those with the best selection performance were chosen.

5 The identified set of selection criteria are displayed in Fig. 5.2: (a₁) $FW_1 > 0.12$ (first Fox-Wolfram moment [15]); (a₂) $\theta_p < 0.55$ rad (polar angle of the proton in the laboratory system); (a₃) $dE/dx_{STT} > 8.5$ a.u. (specific energy loss of the proton candidate in the Straw-Tube-Tracker); (a₄) $\theta_{p\pi} < 0.43$ rad (polar angle of the Λ candidate in the laboratory system).

15 These criteria reduce the background by about a factor 80 while keeping the signal efficiency at 30% as shown in Tab. 5.2. The corresponding invariant mass distributions for signal and background events is displayed in Fig. 5.3 (a). The absolute background suppression factor $f_{\text{sup},(1)} = 134$ meets the requirement for this high cross section channel to reduced the data stream by two orders of magnitude.

(2,4,5) The channels comprising electrons and positrons in the final state can effectively by cleaned from background by a rather rough and simple electron identification. The ratio of the deposited shower energy in the electro-

magnetic calorimeter (EMC) in GeV and the momentum of the matched track in GeV/c is required to be larger than a certain threshold, since electrons/positrons usually have a velocity close to speed of light c_0 and deposite the complete kinematic energy in matter. The criterion for these three channels was (see Figs. 5.2 (b), (d), (e)) $E_{\text{EMC}}/p > 0.7c$, applied to both the electron and positron candidate individually. The quantity has a pronounced peak shaped distribution around $E_{\text{EMC}}/p \approx 1$ for signal events (black histograms) and a rather broad exponential-like shape (red hatched histograms) with just a tail leaking into the signal distribution. In all three cases the background levels are significantly reduced as visible in Figs. 5.3 (b), (d) and (e). Due to the different kinematic situations the absolute suppression factors of background events differ for the three kinds of reactions.

The requirement for a rather high invariant mass $m(e^+e^-)$ in addition to the PID filtering for the reaction $J/\psi \rightarrow e^+e^-$ (2) results in a background suppression to just one residual event out of 10^6 simulated, giving a nominal suppression factor of $f_{\text{sup}} = 10^6$. Due to the limited statistics we estimate the lower limit of the reduction factor corresponding to a

one standard deviation confidence interval to $f_{\text{sup,(2)}} > 303\,000$.

The suppression of the more specific peaking background in $\bar{p}p \rightarrow e^+e^-$ (4) reactions at this lower centre-of-mass energy with $f_{\text{sup,(4)}} > 70\,000$ (derived from 13 residual background events) is a bit less effective but still far above the required factor of $f_{\text{sup,req}} = 1000$.

The least suppression after applying this PID filtering is observed for reaction $\bar{p}p \rightarrow e^+e^-\pi^0$ (5). The originally higher background level introduced by the reconstruction of an additional $\pi^0 \rightarrow \gamma\gamma$ is still present after further filtering. It, however, leads to a background suppression factor of about $f_{\text{sup,(5)}} \approx 3700$ as it can be derived from Tab. 5.2, that still clearly exceeds the goal.

- (3) The decay channel $J/\psi \rightarrow \mu^+\mu^-$ mainly requires a muon identification for the two lepton candidates. For that purpose the penetration depth in the dense material of the muon detector (MDT) can be used. Due to the lower cross section of muons in nuclear matter compared to hadrons they typically cross more iron layers of the MDT. This is shown in Fig. 5.2 (c), where the number of layers is shown for signal (black) and background (red hatched) events. A simple requirement of $N_{\text{MDT}} > 5$ for both muon candidates of the reconstructed J/ψ effectively reduces the background by about a factor $f_{\text{sup,(3)}} \approx 34\,000$ as it can be derived from Tab. 5.2. This easily meets the requirement.

Concerning simultaneous selection of multiple channels, the effect can be roughly estimated from the information given in Tab. 5.2 for the set of channels (= filtering algorithms) under study. We can anticipate that the total accepted background level and thus data rate will increase when tagging multiple reaction types at the same time, since any of the corresponding filtering algorithms will accept a certain fraction of background events. This total fraction f_{tot} is at least as big as the largest fraction $\max_i f_i$ accepted by any of the simultaneously running filtering algorithms (if all other algorithms accept true subsets of the background events accepted by the maximum one), but at most the sum $\sum f_i$ of all fractions (if all algorithms accept pairwise disjoint sets of background events). The latter represents the worst case scenario.

In our case we have two occasions to tag simultaneously, namely where we have multiple reactions

Table 5.2: Efficiencies for signal and background events after combinatorics and additional filtering, and the background suppression factors. The numbers marked with * represent the lower border of the one standard deviation confidence interval.

No.	Channel	$\epsilon_S[\%]$	$\epsilon_B[\%]$	$f_{\text{sup}} [\times 1000]$
(1)	$\Lambda \rightarrow p\pi^-$	30.0	0.75	0.134
(2)	$J/\psi \rightarrow e^+e^-$	31.8	0.0001	> 303*
(3)	$J/\psi \rightarrow \mu^+\mu^-$	56.6	0.0029	> 28.2*
(4)	$\bar{p}p \rightarrow e^+e^-$	53.7	0.0013	> 56.5*
(5)	$\bar{p}p \rightarrow e^+e^-\pi^0$	36.2	0.0269	3.72

selected at the same centre-of-mass energy. These are the combination of channels (2) and (3) as well as the combination of channels (4) and (5). If we consider for both combinations the sum of the background efficiencies as the worst case scenario, we get $f_{\text{sup,(2)+(3)}} \approx 33\,000$ and $f_{\text{sup,(4)+(5)}} \approx 3\,500$ as suppression factors. Both numbers clearly meet the data stream reduction requirement.

Day 1 Scenario In addition to the expected performance for the agreed start setup of the PANDA detector this simulation study was repeated for the case that some of the sub-detectors, namely the barrel EMC and the GEM tracker, are not completely installed at start time. Concretely the EMC coverage was decreased to 75% of the crystals by omitting the top and bottom two of the total of 16 segments around the target pipe system. This has an impact not only on the acceptance of photons, but also on the electron identification required for three of our five reactions. The resultant efficiency numbers and suppression are given in Tab. 5.3. For channels (2), (4) and (5) we observe quite a significant reduction of the signal efficiencies. In particular the acceptance of reactions $\bar{p}p \rightarrow e^+e^-\pi^0$ is lowered by roughly a factor 2.5 to 16.5% due to the missing photon acceptance at the top and bottom of the detector in the barrel region with polar angles $22^\circ < \theta < 140^\circ$. However, these still acceptable signal efficiencies together with the mainly unaltered background levels prove the robustness and effectiveness of the online filtering system.




Figure 5.2: Selection criteria for $\Lambda \rightarrow p\pi^-$ (a₁) – (a₄), $J/\psi \rightarrow e^+e^-$ (b), and $J/\psi \rightarrow \mu^+\mu^-$ (c), $\bar{p}p \rightarrow e^+e^-$ (d), $\bar{p}p \rightarrow e^+e^-\pi^0$ (e₁), (e₂). The black histograms show the distributions for Monte Carlo truth matched signal candidates (the truth match is applied to emphasize the difference between signal and background shapes), the red hatched area the one for background combinations. The blue dashed vertical lines indicate the filtering criteria. Plot (e₁) shows the invariant $\gamma\gamma$ mass to form π^0 candidates, where the shown window size already represents a selection of about ± 3 standard deviations around the peak position.

Table 5.3: Efficiencies for signal and background events after combinatorics and additional filtering, and the background suppression factors. The numbers marked with * represent the lower border of the one standard deviation confidence interval. (DAY1 scenario with 75% EMC).

Channel	ϵ_S [%]	ϵ_B [%]	$f_{\text{sup}} [\times 1000]$
$\Lambda \rightarrow p\pi^-$	30.0	0.82	0.121
$J/\psi \rightarrow e^+e^-$	24.7	0.0001	> 303*
$J/\psi \rightarrow \mu^+\mu^-$	56.3	0.0038	> 22.2*
$\bar{p}p \rightarrow e^+e^-$	40.9	0.0009	> 76.3*
$\bar{p}p \rightarrow e^+e^-\pi^0$	16.5	0.0156	6.41

5.2 Measurement with DAQ prototype

The architecture of the triggerless readout system proposed in Chapter 4 is required to respond to demands of the PANDA experiment. One of the main requirements for the data acquisition concept mentioned in the Summary Chapter 1 and described in more depth in Chapter 3 is stable and synchronous data collection. The precise time-distribution system is an essential part of the data readout. This system provides the synchronization signal to coordinate the different branches of data flow with one reference time. As one of the constituents, the software protocol, SODANET, plays an important role in this process of time-sharing. As explained above in section 4.2.1.1, this protocol is designed to propagate synchronization commands within whole readout system with a certain precision. Through these commands a time-phase update is




Figure 5.3: Invariant mass spectra of the relevant signal resonances in the regions of interest after combinatorics and additional filtering for the five benchmark channels: (a) $\Lambda^0 \rightarrow p\pi^-$ (and charge conjugates), (b) $J/\psi \rightarrow e^+e^-$, (c) $J/\psi \rightarrow \mu^+\mu^-$, (d) $\bar{p}p \rightarrow e^+e^-$, (e) and $\bar{p}p \rightarrow e^+e^-\pi^0$. The black hollow histograms show the distributions for signal events, the red hatched histograms the corresponding distributions for background (DPM) events. Clearly the background levels are significantly reduced in all cases compared to Fig. 5.1.

applied for each FEE part simultaneously, which implies synchronization of the whole system. Thus, correlations between different sub-systems must be clearly seen.

synchronization, which meets our expectations.

5.2.0.1 Test-beam setup at COSY

Previous measurements showed that sigma can reach precision about of 20 ps in terms of RMS, as shown in Figure 4.17, if the same signal from two synchronized channels of the different FEE boards is compared. Even so, since the PANDA detector will have thousands of such channels, one cannot conclude that such precision is valid with complex readout setups. The reason for the above-mentioned measurements was to estimate the precision level of SODANET. Therefore, it was proposed to test a current DAQ prototype with a multi-component detector system, which is complex, in order to determine the precision of the time-phase synchronization. Additional goals were to investigate possible bottlenecks in data flow and to fix them. Beamtime at the COSY facility corresponded perfectly to our requirements for such an experiment, as COSY could provide multi-constituent setup and high energy of beam to produce relatively large amounts of data. These measurements, using the current DAQ prototype, helped to confirm a high precision of the time-phase

We had beamtime at COSY in February 2019. All detector prototypes connected to the DAQ system were installed on the support frames at a height of at least 2 m to allow crossing with the proton beam. The full length of the experimental setup was around 10 m. The beam had a momentum of 2.5 GeV/c and came from the right side as shown in Fig. 5.4. The beam spot had a size of about one square centimeter. The frequency of the beam was around 40 kHz. The readout chain of DAQ was conform to the diagram shown in Fig 4.1 which has push-only architecture. The PANDA EMC prototype, FT straw tubes and TOF plastic scintillator paddles were read by the DAQ prototype. Hence, they represented the FEE part of the readout chain.

The PANDA EMC prototype is called Proto9 subunit. This prototype was built in accordance with the design of the subunits which will be used for the PANDA Forward Endcap EMC. Earlier, the EMC prototype was used to test the data processing algorithms, which are implemented in the FEE to perform the triggerless data acquisition




Figure 5.4: Final detector's setup during the testbeam time at COSY.

tion. By itself Proto9 represents an assembly of nine PWO-II crystals wrapped in reflecting foil to increase light collection. The crystals were inserted into a metal alveole serving as a frame for sixteen crystals; the other places were filled by seven brass bars of the same size in order to stabilize the framework. One crystal has two Large Area Avalanche PhotoDiodes (LAAPDs) covering the back plane and one preamplifier connected to these photodiodes. Each LAAPD collects the scintillation light created in case of energy deposition. As a next step, the charge signal produced by the photodiode is amplified by the preamplifier and converted to the analog waveform by the shaper. In total, there are eighteen photosensors, nine preamplifiers and nine shapers for the nine crystals in this subunit. It was placed inside the detector box and attached to an aluminum mounting plate with metal pipes soldered on the other side. The detector box was thermally isolated and connected to the cooling system. Water cooling provided a stable temperature for the crystals during the measurements. Thus, this cooling system was sufficient to obtain the acceptable results from the EMC subunit in this experiment. For a reliability check of the readout parts, a light pulser was provided for the alveole. The light pulser was used as a reference: the light entered inside the subunit through an optical fiber, and, as result, all channels produced a signal corresponding to the certain

energy peak, observed in the final histogram.

The next constituent of the EMC readout chain was the Sampling Analogue-to-Digital Converter board (SADC), developed by Pawel Marciniewski of the University of Uppsala, Sweden, for the FwE EMC. Thirty-two input channels are available for connection with FEE, and one bi-directional interface link serves for reading out data or configuring the SADC. In our case, each board read 9 channels, one LAAPD of each crystal respectively, as can be seen in Fig. 5.5. During this sampling the waveform is doubled onto the low gain and high gain components. The former is used for representation of energy range from 1 GeV to 15 GeV while the latter is for energies from 3 MeV up to 1 GeV. As shown in Fig. 4.8, the slow control can access the FEE part, which is SADC in our case. This access allows setting the thresholds and filter parameters. The SADC board also includes two Field-Programmable Gate Arrays (FPGAs) which perform preprocessing of the waveforms. This preprocessing concerns the baseline subtraction, pulse detection, pile-up identification and recovery, and feature extraction (determination of timestamp and amplitude of pulse over the threshold). This board can also send raw data without preprocessing. Such features make this board an important part of the trigger-less readout.




Figure 5.5: Photograph of the SADC prototypes connected to the Proto9.

The output channel of the board is connected to the DC modules, introduced above in section 4.1.1. To concentrate the data during the beam-




Figure 5.6: Photograph of the TRBv3 board. Two FPGAs on the right side perform functions of the DC. Another two on the left side: the top one is the UDP server, and the bottom one is the SODANET source.

time, the Time Readout Board of the third version (TRBv3) [16] was used instead of the hardware platform proposed as the DC unit for PANDA experiment. This board has five ECP3-150 FPGAs, one master and four slaves respectively. This flexibility allows us to use different firmware, dependent on the field of application. In our case, the TRBv3 performed several functions, serving as DC module or SODANET source/SODANET hub. A photo of the connected TRBv3 can be seen in figure 5.6.

Slow control was hosted by a Windows machine which was connected to DC modules through a LAN. The slow control was realized as the Labview GUI program, which establishes links between all building blocks of the readout chain. As connection was established, the FEE was configured with default parameters. One can set the different parameters and initialize them for each FEE separately. After that, the user can begin data collection. The slow control commands came to the UDP server and were delivered to the central FPGA, which sent commands to the DC and the SODANET source. As the next step, the synchronization signal came from the SODANET source or SODANET hub, depending on the firmware, to the DC modules, which distribute this signal among connected FEEs. Each period of signal corresponds a certain superbunch number as explained in section 4.2.1.2

Further, the synchronized preprocessed data came out from DCs and was sent to a set of three Compute Nodes. These modules are described in subsection 4.1.2. Compute Nodes are expected to do event building and preprocessing, but in our case they functioned as a Burst Building Network (BBN), being its constituents. All three CNs were inserted into an ATCA crate, which provided connection between them. Two peripheral CNs received the data coming from the DC modules connected to the EMC and the FT straw tubes. The third central CN combined these streams of data into one. Finally, a combined stream of data was sent to the PC for storage.

5.2.0.2 Strategy of measurements

To estimate stability of the data acquisition and precision of the synchronization, a strategy of increasing complexity was chosen. This was done because of the previously mentioned requirement that the data acquisition system for PANDA experiment should demonstrate the same performance regardless of readout complexity. As the simplest case,




Figure 5.7: Diagrams of readout schemes which were tested during the beamtime at COSY. The top two, (a) and (b), were used for tests only with Proto9. The bottom diagrams represent the cases when other subdetectors were connected in addition to the EMC prototype.

the scheme containing one DC in the readout chain and two SADCs was selected, see Fig. 5.7(a). This scheme allowed investigating hit distributions in the EMC prototype. Further, the ADC channels that belonged to one crystal and had highest statistics were selected in order to determine the deposited energy in this crystal by combining the channel's output in case of time coincidence. For this reason, a time difference between hits from LAAPDs corresponding these ADC channels was an important quantity to determine synchronization stability. Another scheme 5.7(b) had two DCs, one DC per SADC board. The same channels as in case of the first scheme were measured. Hence, it was possible to investigate the impact of splitting the data stream into two separate DCs. Thus, the SODANET performance was estimated within the EMC prototype.

As the next step, more complex readout chains were composed which contain another two subdetectors - TOF and FT-STT. Diagrams of these schemes are shown in Fig 5.7 (c,d). For these cases, separate

TRBv3 functioned as the SODANET source. Synchronization signals from that board were sent directly or via SODANET hub to all the readout components. The main goal of scheme (c) was to investigate the correlations between EMC and TOF systems. The time difference between EMC and TOF hits was checked determine the event coincidence, when one proton made a hit in the TOF paddle and further in the EMC subunits, taking into account that clusters might occur. Therefore, one TOF hit corresponded to several hits of different EMC crystals. Thus, the synchronization precision was estimated by combining the data from the whole EMC subunit and the TOF paddle.

The final configuration of the readout chain included the FT subdetector. Consequently, coincidence events between FT and other two sub-systems were expected. The same analysis procedures were performed as in the case of EMC-TOF, but for FT-TOF or FT-EMC combined data. The main reasons for the last two schemes were assessment of the synchronization precision in the case of the multi-




Figure 5.8: Distribution of detected hits between ADC channels.




Figure 5.10: Discrepancy between ADC channels in case of proton peak position before alignment.




Figure 5.9: Energy distribution of hits coming from the ADC channel. Light pulser peak is located on the right side. Proton peak can be seen on the left side.




Figure 5.11: Divergence of time difference between EMC ADC channels and TOF system before alignment.

component system and evolution of this precision over time.

5.2.0.3 Calibration

Figure 5.8 represents hits distribution between ADC channels. Each hit has its own energy and timestamp respectively. An example of the hit energy distribution corresponding to a certain ADC channel is shown in Fig. 5.9. Ideally, all channels should have same positions of the proton peak and the light pulser, but, because of difference in the electronic pieces, there is a discrepancy in energy present between ADC channels. Figure 5.10 shows the discordance of SADC channels for the proton peak energy. To fix this problem, the calibration coefficient between a constant energy value and the proton peak was calculated for each ADC channel. After that, all channels were normalized in order to make them consistent between each other. Timestamps were also affected by the non-identical

readout electronics, so there was a discrepancy in time for different ADC channels in case of coincidence of events between ADC channel and other sub-systems. The shift constants, which were found beforehand for each ADC channel, were applied in order to fix the time discrepancy. The constants shown in Fig. 5.11 represent the time difference between the ADC channel and the TOF paddle. After applying the constant to the corresponding ADC channel, all time differences between sub-systems become distributed around 0.

5.2.0.4 Analysis results

The readout configuration with one DC shown in Fig. 5.7(a), had the hit distribution which is shown in Fig. 5.8. Due to the relatively small size of the beam spot and the position of the beam, some crystals had a higher interaction rate, therefore corresponding readout channels obtained more hits. A first group of channels below 20




Figure 5.12: Deposited energy distribution for a selected crystal. This distribution was obtained by combining the data from two SADC channels which correspond to the LAAPDs connected to this crystal.




Figure 5.13: Time difference between two LAAPDs of the same crystal.

belongs to the first ADC board, and the second group which starts from 64 to the second ADC. The distribution of hit energy is presented by Fig. 5.9. The peak of the light pulser can be seen on the left side. The continuous exponential shape with peak is caused by various proton reactions with scintillation material. The deposited energy, left via direct interaction of a proton with a single crystal, can be seen in Fig. 5.12. This energy was reconstructed by combining information from two LAAPDs of this crystal. This reconstruction was made in accordance with the requirement of their coincidence in time, which was determined by the time difference between these LAAPDs, see Fig. 5.13. Evolution of this time difference can be seen in Fig. 5.14. The time values along the X axis were obtained by converting the superburst numbers of hits. As the figure shows, the time difference was not changed over forty minutes that confirmed stability of the synchronization. Otherwise, a rapid




Figure 5.14: Time evolution of the time difference.




Figure 5.15: Time difference-deposited energy correlation for the single crystal.

change would be observed because of difference in crystal oscillators placed at the SADC. Figure 5.15 shows time difference - hit energy correlation for a single crystal. One can conclude that a significant number of coincidences between LAAPDs occurred at energy corresponding to the above-mentioned proton peak. As appropriate energy cuts were applied, the time difference distribution became similar to the Gauss distribution, see Fig. 5.16, similar to the proton energy distribution. The evolution of the time difference in applied energy cuts is represented by the time profile in Fig. 5.17. Each point in this profile corresponds to the mean value of the time difference distribution obtained over one minute of measurements. Error bars are defined by sigmas of these distributions divided by a square root of a total number of events per minute. The mean time difference varies within 100 ps, which determines the precision of time synchronization. Performing the same analysis for two DCs scheme, we obtained the time difference between two LAAPDs of one crystal depicted in Fig. 5.18. This time difference stayed within a 100 ps range for half an hour, see Fig. 5.19. One main




Figure 5.16: Time difference after application of the energy cuts.




Figure 5.18: Time difference in the case of two DCs. The time shift caused by wire length difference is observed.




Figure 5.17: Profile representation of the time difference evolution.




Figure 5.19: Time difference profile in the case of two DCs.

difference, comparing the previous readout scheme, was the peak position, which had an offset because of the difference in length of fibers connected to the DC. Both readout schemes, in which only the EMC was used, confirmed the high stability of synchronization, however, the precision was lower than expected.

The data obtained with the next scheme, when we had EMC+TOF, revealed some interesting observations about beam structure. Figure 5.20 shows a time correlation of TOF and EMC hits in logarithmic scale. The main peak from that plot corresponds to a hit coincidence between both systems, but the other bumps represent a random coincidence between different bunches of the beam pulse. These bumps cannot be seen in linear scale, see Fig. 5.21. This distribution was filled with time differences between all ADC channels and TOF paddle, not only channels corresponding to the specified crystal as before; all these differences were calibrated as explained above. Based on this time difference, the cluster energy distribution shown




Figure 5.20: Time difference between TOF and ADC channels in logarithmic scale.




Figure 5.21: Time difference between TOF and ADC channels.




Figure 5.22: Cluster energy distribution in the case of time cuts.




Figure 5.23: Time difference between TOF and ADC channels if energy cut was applied.




Figure 5.24: Time difference evolution (EMC+TOF).




Figure 5.25: Time difference between TOF and FT sub-systems.

in Fig. 5.22 was found. As in the case with a single crystal, the proton peak region was selected from this energy spectrum in order to estimate the synchronization. The time difference for these energy region is presented in Fig. 5.23. The time profile corresponding to this time difference shown in Fig. 5.24, shows that the deviation of the mean is about of 20 ps.

The last scheme in our measurements provided an opportunity to compare results with the previous scheme. There were two possible combinations used to look at the time correlations: TOF-FT and FT-EMC. In case of checking the TOF-FT combined data, we again observed the beam microstructure. The bumps related to the random coincidence between events can be seen in Fig. 5.25. The peak which is shown in Fig. 5.26 corresponds to the coincidence of TOF and FT for one event. The small value of sigma, around 0.9 ns, confirms high synchronization between these two systems. For the combined data from another pair of subdetectors,




Figure 5.26: Time difference between TOF and FT in linear scale.




Figure 5.28: Time difference between FT and all EMC channels in linear scale.




Figure 5.27: Time difference between FT and all EMC channels.




Figure 5.29: Cluster energy distribution in the case of FT-EMC event coincidence when time cut was applied.

EMC and FT, we repeated the same investigation procedure as for TOF-EMC in order to determine the synchronization precision. The pictures 5.27 and 5.28 show the time difference between their detected events in logarithmic and linear scales. Selecting the main coincidence regions yields the cluster energy distribution as shown in Fig. 5.29. Figure 5.30 shows the time difference for the proton peak region. The evolution of this time difference with duration of measurements is represented by the time profile in Fig.5.31. The synchronization precision obtained from this profile to be around 20 ps.

Therefore, we can conclude that, in case of the last two schemes, synchronization provided by the SODANET source was stable and had the same precision order, independently of the data source. The analysis indicated that precision values were comparable with results obtained earlier in less complex readout systems. The time deviation between the synchronized readout channels was not higher than 20 ps in case if the calibration was performed, and

statistics was sufficient.




Figure 5.30: EMC-FT Time difference in the case of cluster and energy cuts




Figure 5.31: EMC-FT Time difference evolution.

5.2.0.5 Conclusion

With the help of the beamtime at COSY, we estimated the performance of the current DAQ prototypes. The obtained results of these investigations showed that a high level of time-phase synchronization can be provided, using the SODANET protocol. The precision of this synchronization was comparable to the previous measurements described in 4.2.1.4. The standard error of the mean of the time difference distribution, when two systems are synchronized, did not exceed 20 ps. This result was independent of the complexity of the readout schemes used during the measurements.

Therefore, we can conclude that the proposed design of the triggerless readout approach suits the requirements of the PANDA experiment for the Day 1 scenario. This readout system provides stable synchronous data acquisition from all connected constituents of the multi-component detector setup. However, additional tests must be done with a final version of readout components such as Data Concentrators and CNs. In addition, each subdetector of PANDA should be capable of working in a stand-alone regime, and further software/firmware developments and tests are needed in this direction.

Bibliography

[1] V. Flaminio, W. G. Moorhead, D. R. O. Morrison, N. Rivoire. Compilation of cross sections III: p and \bar{p} induced reactions. *CERN HERA 84-01*, pages 177–178, 1984.

[2] M. Tanabashi *et al.* (Particle Data Group). Review of Particle Physics. *Phys. Rev. D*, (98):030001, 2018.

[3] E. Braaten *et al.* Effects of charged charm mesons on the line shapes of the $X(3872)$. *Phys. Rev. D*, (77):014029, 2008.

[4] C. Hanhart *et al.* Reconciling the $X(3872)$ with the near-threshold enhancement in the $D^0\bar{D}^{*0}$ final state. *Phys. Rev. D*, (76):034007, 2007.

[5] Yu. S. Kalashnikova *et al.* Quark and Meson Degrees of Freedom in the $X(3872)$ Charmonium. *Phys. Atom. Nucl.*, (73):1592–1611, 2010.

[6] K. Götzen *et al.* Precision resonance energy scans with the PANDA experiment at FAIR (draft for RN-QCD-2016-002). 2018.

[7] A. Zichichi, S. M. Berman, N. Cabibbo, and R. Gatto. *Nuovo Cim.*, (24):170–180, 1962.

[8] S. Pacetti, R. Baldini Ferroli, and E. Tomasi-Gustafsson. *Phys. Rept.*, (1):550–551, 2015.

[9] B. Singh *et al.* *Eur. Phys. J.*, (A52(10)):325, 2016.

[10] M. P. Rekaló. *Sov. J. Nucl. Phys.*, (1):760, 1965.

[11] A.Z. Dubnickova, S. Dubnicka, M.P. Rekaló. *Z. Phys. C*, (70):473–481, 1996.

[12] C. Adamuscin, E. A. Kuraev, E. Tomasi-Gustafsson, and F. E. Maas. *Phys. Rev. C*, (75):045205, 2007.

[13] J. Boucher. *Feasibility studies of the $\bar{p}p \rightarrow e^+e^-$ electromagnetic channel at PANDA*. PhD thesis, 2011.

[14] J. Guttmann and M. Vanderhaeghen. *Phys. Lett. B*, (719):136–142, 2013.

[15] G. C. Fox and S. Wolfram. *Phys. Rev. Lett.*, (41):1581, 1978.

[16] J. Michel. Ph.D. Thesis. *Johann Wolfgang Goethe-Universität, Frankfurt am Main*.

6 Project managements and resources

6.1 Work-packages and responsibilities

All DAQ-related work packages can be split into four categories:

- hardware;
- firmware;
- software;
- integration, support and maintenance.

Responsibility for each work package is taken either by a research group/institution or DAQ core team – team of experts coordinated directly by the DAQ system coordinator. In case the work package is lead by the DAQ core team the amount of required person-years (PY) is specified.

Hardware work packages. These work packages include research and development activities, design, testing, production and screening of PANDA-specific hardware.

Data-Concentrator module, see 4.1.1, is part of the DAQ core, used as data concentrator hardware for different subsystems and for the burst-building network as well.

- *Hardware design and prototyping:*
 - responsibility: Uppsala University;
 - milestones:
 - * design completed: 01.2020;
 - * first hardware prototype produced and screened: 08.2020;
 - * re-design completed (if necessary): 03.2021;
 - * hardware prototype of the final version produced and screened: 11.2021;
- *Hardware prototype testing/debugging:*
 - responsibility: DAQ core, 0.5 PY;
 - milestones:
 - * testing completed: 01.2021;
- *Hardware production and screening:*
 - responsibility: Uppsala University/DAQ core, 0.5 PY;

– milestones:

- * production and testing completed: 09.2023.

Compute Node modules, see 4.1.2 are used for initial tracking and event building.

- *Hardware design and prototyping:*
 - responsibility: IHEP, Beijing;
 - milestones:
 - * final design completed: 01.2021;
 - * hardware prototype produced and screened: 08.2021;
 - * re-design completed (if necessary): 03.2022;
- *Hardware prototype testing/debugging:*
 - responsibility: IHEP, Beijing;
 - milestones:
 - * testing completed: 01.2022;
- *Hardware production and screening:*
 - responsibility: IHEP, Beijing;
 - milestones:
 - * production and testing completed: 01.2023.

HPC interface, see section 4.1.3, receives data from compute nodes and transfers via Infiniband network to the HPC cluster.

- *Hardware design and prototyping:*
 - responsibility: KIT, Karlsruhe/DAQ core;
 - milestones:
 - * final design completed: 01.2021;
 - * hardware prototype produced and screened: 08.2021;
 - * re-design completed (if necessary): 03.2022;
- *Hardware prototype testing/debugging:*
 - responsibility: KIT, Karlsruhe/DAQ core;
 - milestones:

* testing completed: 01.2022;

- *Hardware production and screening:*

- responsibility: KIT, Karlsruhe/DAQ core;

- milestones:

- * production and testing completed: 01.2023.

Firmware work packages. These work packages include research and development activities, firmware design and testing.

SODANET, see section 4.2.1, is the synchronization and communication protocol between front-end electronics and the data acquisition. It includes several firmware cores as SODANET source, hub and receiving end-point. These cores are designed to run on DC hardware, see 4.1.1, and TRBv3 hardware platform. The same communication protocol without synchronisation capabilities is employed to configure and control the burst-building network and compute nodes.

- *Development and support:*

- responsibility: DAQ core, 0.5 PY;

- milestones:

- * test firmware on the final version of DC modules: 12.2020;

- *DAQ-accelerator interface:* develop a communication protocol between the White Rabbit hardware, used by accelerator facilities at FAIR, and SODANET. The main point of the development is to precisely receive timing of the beam bunch – target crossing.

- responsibility: DAQ core, 0.5 PY;

- milestones:

- * design completed: 12.2022;

- * testing completed: 12.2023;

Data-handling IP cores: This work package foresees development, testing and deployment of a framework for data-handling IP cores and communication protocols in the burst-building, compute-nodes and HPC-interface levels of the DAQ. All IP cores have common interfaces. This allows skip some data processing if required or place particular IP core on a different hardware, e.g. implement burst building network and compute nodes functionality in the HPC-interface module, see section 4.2.4 for details.

- *Development, testing and deployment:*

- responsibility: DAQ core, 2 PY;

- milestones:

- * definition of the framework and communication protocols, specifications of IP-core interfaces: 07.2020;

- * implementation of the framework and readout test with several subsystems: 03.2021;

- * implementation of all key data-handling IP-cores and start of commissioning of the complete readout: 01.2020;

Data-processing algorithms: These work packages include only key online algorithm which are running on the FPGA-based level of the DAQ and are crucial for the event building and filtering. Final event building and selection takes place in the HPC cluster and is part of the online-computing TDR. Zero-suppression algorithms as well as data packing/unpacking algorithms running on the data-concentrator level are provided by corresponding subsystems.

- *Event and burst-building:* includes package-addressing algorithms running on the data-concentrator level, routing functionality of the burst-building module and rough event building at the compute-node level;

- responsibility: DAQ core, 0.5 PY;

- milestones:

- * test firmware on the final version of DC modules: 05.2022;

- *Tracking (central tracker):* development, validation and implementation of online track finder and fitting algorithms running on compute nodes or GPUs for the central tracker.

- responsibility: Lanzhou;

- milestones:

- * design completed: 12.2022;

- *Tracking (forward tracker):* development, validation and implementation of online track finder and fitting algorithms running on compute nodes or GPUs for the forward tracker.

- responsibility: Krakow;

- milestones:

- * design completed: 12.2022;

- *EMC clustering (support):* support of the developed and tested online clustering for the EMC.

- responsibility: DAQ core, 0.5 PY;
 - milestones:
 - * deployment for complete detector: 12.2022.
- 5 **DAQ-software work packages.** These work packages include development of the DAQ software which controls, configures DAQ hardware and communicates with the experiment control system.
- 10 • *DAQ control:* This software configures and controls DAQ hardware using SODANET protocol.
 - responsibility: DAQ core, 1 PY;
 - milestones:
 - 15 * software for small-scale stand-alone DAQ: 12.2020;
 - * final design: 12.2022;
 - *DAQ interfacing:* Within this work package protocols for DAQ communications with DCS and ECS are developed.
 - 20 – responsibility: DAQ core, 1 PY;
 - milestones:
 - * protocols defined: 12.2020;
 - * protocols implemented and tested with large-scale system: 12.2022;
 - 25 • *Stand-alone DAQ:* This software includes on-line monitoring and data processing/storage with DAQ in minimal-required configuration of services, as database support, DCS interface.
 - responsibility: RUG/KIT/Bochum/Julich;
 - 30 – milestones:
 - * system ready: 12.2020.
- Integration, support and maintenance.** These work packages include developments required for integration of separately developed components as well as support of running DAQ infrastructure.
- 35 • *DCS – Run control integration:* This work package focuses on the efficiency of automation of error-handling and recovery of the DAQ subsystem. Since the error-handling higher-level recovery / automation is done by the ECS, which is not subject of this review, this work package foresees joint effort with the ECS group. This development is extension of the *Stand-alone DAQ* work package.
 - 40 – responsibility: DAQ core, ECS core;
- milestones:
 - * Error-handling software for small-scale stand-alone DAQ: 12.2021;
 - * final design: 12.2022;
 - 5 • *Integration of commercial components:* This work package aims to investigate possibility of employing commercial hardware and available communication protocols within the DAQ as described in section 4.2.5. This work package is strongly related to the *Data-handling IP cores* one.
 - responsibility: KIT, Krakow, DAQ core;
 - milestones:
 - 15 * decision on used communication protocols: 03.2021;
 - * decision on employed hardware: 12.2022.
 - *DAQ support and maintenance:* Within this work package support is provided to the subsystems to implement required functionality up to data-concentrator level, setup stand-alone DAQ instances, deploy PANDA DAQ at FAIR facility and maintain its operation.
 - 20 – responsibility: DAQ core, 2 fte.
- ## 6.2 Schedule and Milestones
- 25
- Planning milestones:*
- TDR approved (M3): 09.2020;
 - Funding established (ME1): 03.2021;
 - Contracts signed (M4): 09.2021;
 - Planning completed (M7): 09.2021.
- 30
- Construction milestones:*
- Tender completed (ME2): 06.2021;
 - All material is acquired (MX1): 09.2023;
 - Prototype/Pre-series testing complete, production readiness (M8): 09.2022;
 - 35 • Acceptance test completed (Factory Acceptance Test passed) (M9): 09.2023.
- 40
- Installation milestones:*
- Approval for installation (M10): 09.2023;
 - Ready for beam (M11): 03.2024;
 - 40 • Ready for operations (M12): 06.2025.

6.3 Cost

The DAQ core group next to the dedicated DAQ hardware will provide the DAQ server and a PC interfacing DAQ with HPC cluster. It is assumed that all data-base infrastructure will be provided by DCS and ECS subsystems. As a draft for the cost estimates, we currently assume the following numbers for the full PANDA setup (18 sub-detectors):

- DC hardware, 40 units (1 SODANET source, 33 DCs, 1 burst building network, spares): 40×8 kEuro = 320 kEuro;
- MicroTCA crates for DC modules (12 slots each), 5 units: 5×5 kEuro = 25 kEuro;
- Compute nodes, 6 ATCA modules (5 ATCA carriers, 20 AMC cards, spares): 240 kEuro;
- ATCA crate for compute nodes and spares: 15 kEuro;
- HPC interface card and spares: 30 kEuro
- Interfacing PC hosting HPC interface and Infiniband cards, reserve and spare: 10 kEuro;
- DAQ server, reserve and spare: 10 kEuro;
- Optical-fibre infrastructure: 210 kEuro (infrastructure within the PANDA hall) + 60 kEuro (infrastructure between the PANDA hall and Green Cube).

In summary we estimate the total cost of the PANDA DAQ to be 920 kEuro.

We acknowledge financial support from the Science and Technology Facilities Council (STFC), British funding agency, Great Britain; the Bhabha Atomic Research Centre (BARC) and the Indian Institute of Technology Bombay, India; the Bundesministerium für Bildung und Forschung (BMBF), Germany; the Carl-Zeiss-Stiftung 21-0563-2.8/122/1 and 21-0563-2.8/131/1, Mainz, Germany; the Center for Advanced Radiation Technology (KVI-CART), Groningen, Netherlands; the CNRS/IN2P3 and the Université Paris-Sud, France; the Czech Ministry (MEYS) grants LM2015049, CZ.02.1.01/0.0/0.0/16 and 013/0001677, the Deutsche Forschungsgemeinschaft (DFG), Germany; the Deutscher Akademischer Austauschdienst (DAAD), Germany; the Forschungszentrum Jülich, Germany; the European Union's Horizon 2020 research and innovation programme under grant agreement No 824093, the Gesellschaft für Schwerionenforschung GmbH (GSI), Darmstadt, Germany; the Helmholtz-Gemeinschaft Deutscher Forschungszentren (HGF), Germany; the INTAS, European Commission funding; the Institute of High Energy Physics (IHEP) and the Chinese Academy of Sciences, Beijing, China; the Istituto Nazionale di Fisica Nucleare (INFN), Italy; the Ministerio de Educacion y Ciencia (MEC) under grant FPA2006-12120-C03-02; the Polish Ministry of Science and Higher Education (MNiSW) grant No. 2593/7, PR UE/2012/2, and the National Science Centre (NCN) DEC-2013/09/N/ST2/02180, Poland; the State Atomic Energy Corporation Rosatom, National Research Center Kurchatov Institute, Russia; the Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung (SNF), Swiss; the Stefan Meyer Institut für Subatomare Physik and the Österreichische Akademie der Wissenschaften, Wien, Austria; the Swedish Research Council and the Knut and Alice Wallenberg Foundation, Sweden.

List of Figures

5	2.1 Schematic of the future FAIR layout incorporating the current GSI facilities on the left; on the right the future installations, the SIS 100 synchrotron the storage and cooler ring complex including CR and HESR and the Super FRS experiment being some of the new parts. $\bar{\text{PANDA}}$ is positioned right in the center of the image inside the HESR.	6	3.4 Time difference between the last hit of an event and the first hit of the next event in the MVD	14
10	2.2 [2] Time dependent macroscopic luminosity profile $L(t)$ in one operation cycle for constant (solid red) and increasing (green dotted) target density ρ_{target} . Different measures for beam preparation are indicated. Pre-cooling is performed at 3.8 GeV/c. A maximum ramp of 25 mT/s is specified for acceleration and deceleration of the beam.	6	3.5 Integral of the time difference between the last hit of one event and the first hit of the next event as percentage value for all events.	14
15	2.3 Aerial view of $\bar{\text{PANDA}}$ with the Target Spectrometer (TS) on the left side, and the Forward Spectrometer (FS) starting with the dipole magnet on the right. The antiproton beam enters from the left.	9	3.6 In the top of the figure, particle distributions are reported as a function of particle momentum and polar angle. In the bottom part of the figure the 1D profile is projected onto the polar angle. The red line results from the elastic scattering process that is implemented in the DPM generator (MVD TDR).	16
20	2.4 Side view of $\bar{\text{PANDA}}$ with the Target Spectrometer (TS). The antiproton beam enters from the left.	10	3.7 In the top of the figure, particle distributions, obtained with 4×10^4 UrQMD events, are reported as a function of particle momentum and polar angle for antiproton-N and antiproton-Au collisions. In the bottom part of the figure, the 1D profile is projected onto the polar angle, obtained with a statistical sample of 2×10^6 (MVD TDR).	16
25	2.5 Side view of $\bar{\text{PANDA}}$ forward Spectrometer (FS). The antiproton beam enters from the left.	11	3.8 Extracted average count rates for antiproton-proton reactions obtained with the DPM generator. The distribution over all individual channels of the micro-strips and the pixel front-end with the highest occupancy are shown in the frames on the left at top and bottom respectively. Arrows indicate the position inside the MVD. Results for all readout chips (FE) in the different detector layers are shown on the right (MVD TDR).	17
30	3.1 Measured antiproton-proton total (black) and elastic (red) cross section and calculated as the difference between the two the inelastic (blue) cross section in the momentum range of the PANDA experiment	13	3.9 Simulation of P-bar p reactions giving the number of hits per event and per cm along the straw tubes in the innermost layer of the PANDA STT. The target position is at z=0cm. (Figure from PANDA STT TDR).	18
35	3.2 Simulated time between two consecutive anti-proton target interactions with a mean time between events of 500 ns, a continuous beam with 1600 ns and a gap of 400 ns	13	3.10 The readout structure of the barrel DIRC subsystem using TRB3 boards.	18
40	3.3 Time difference between two consecutive hits in the MVD for 1000 simulated background events	14		

5	3.11 Hit rate in the target EMC from the DPM background generator at $p\bar{p} = 14$ GeV/c. Top panel – barrel EMC, bottom panel – forward end-cap EMC. [EMC TDR, p 36.]	19	4.12 The photograph (upper panel) and a sketch (lower panel) of the set-up used to test the synchronous serial optical-link connection.	32
	3.12 Readout architecture for the forward tracker in the High Resolution mode.	21	4.13 Stability of the recovered phase. During the measurement the front-end module was periodically reset. The insert shows a zoom of the region around 30 ns.	32
10	3.13 Readout architecture for the forward tracker in the High Luminosity mode.	22	4.14 Difference of arrival times of synchronisation commands at different front-ends.	33
	3.14 Event rates at sensors of the first detector plane at 1.5 GeV/c for a average luminosity of 1×10^{32} cm ⁻¹ s ⁻²	23	4.15 Measurement of the jitter induced by the measurement set-up.	33
15	4.1 The overview of the PANDA readout system.	26	4.16 Scheme of the setup used for the long-term stability test of the SODANET.	33
	4.2 Block diagram of the PANDA Data Concentrator.	26	4.17 Time-difference between the same signal measured by two TDC channels (T_{TDC}).	34
20	4.3 A preliminary layout of the PANDA Data Concentrator.	26	4.18 SODANET jitter (ΔT).	34
	4.4 ATCA compliant FPGA based compute node. (top): carrier board; (bottom) xFP daughter board. The architecture is based on XILINX [®] Virtex FPGAs. A single CN can support up to 16 optical links and a total of 18 GB DDR2-RAM. The FPGAs are equipped with embedded PowerPC - CPUs, running Linux operating systems for slow control functions.	27	4.19 Stability of SODANET synchronization. The lower panel shows stability of the system once the hardware warmed up to the working temperature.	34
25	4.5 PCIe High-Flex readout card [1].	28	4.20 Effect of the warming-up process of one FEE on the synchronization stability (upper panel) and the temperature of the FEE FPGA (lower panel). Curves with different colours represent two independent measurements.	34
30	4.6 The basic concept of direct FPGA-GPU data-link tunnelled through InfiniBand network [1].	28	4.21 Hardware set-up used for long-term stability-tests of the SODANET synchronization protocol.	35
	4.7 Illustration of the data-flow of the traditional approach (dashed line) and direct FPGA-GPU communication (top), and round-trip latency measurements. (Figure taken from [1].)	28	4.22 Long-term stability of SODANET synchronization at constant room temperature. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.	35
35	4.8 The topology of the SODANET network for the PANDA experiment.	29		
40	4.9 Structure of the SODANET packet as it is sent through a serial link. Each block corresponds to 1 byte	29		
	4.10 Measured propagation time (T_d) of a synchronization command from the source to each FEE is used to correct local timestamps.	31		
45	4.11 Header structure of a SODANET data-packet.	31		

5	4.23 Long-term stability of SODANET synchronization at constant room temperature measured with additional SODANET hub in the network. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.	36	4.30 Example of a single track displayed in X-Y view (left) and ?-Z view (right). In the ?-Z view, the blue solid line represents the fit result from the first iteration using central position of each hit. The red dashed line shows the fit result after the second iteration.	45	5
10	4.24 Long-term stability of SODANET synchronization at forced temperature change of the SODANET source. Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.	36	4.31 The Compute Node used in this study, equipped with 5 Virtex 4 FX60 FPGA chips.	45	10
15	4.25 Long-term stability of SODANET synchronization at forced temperature change of one of the SODANET end-points (TDC). Blue, green, yellow, red and black curves, right axis, correspond to temperatures measurements of different FPGAs. Pink curve, left axis corresponds to the measured time difference between TDC events measured by two different TRB modules.	36	4.32 Simplified block-diagram of the algorithm in VHDL.	45	15
20	4.26 Schematic view of the data flow to the different processing stages within the PANDA experiment	38	4.33 Device utilization summary for the design.	46	15
25	4.27 The overview of the PANDA readout system starting at the Data Concentrator (DC) level.	41	4.34 Example of a single event . The red circles represent the tracks found by the algorithm.	46	15
30	4.28 The overview of the PANDA readout system employing commercial hardware at the burst-building level.	42	4.35 The tracking efficiency and the resolution as a function of the transverse momentum.	47	20
35	4.29 Illustration of the track finding algorithm. The black circles represent the individual straw tubes. The red solid curve indicates the trajectory of a charged particle. The red dashed circles inside tubes represent the isochrones.	44	4.36 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.	47	25
40					
45					
50					

- 4.37 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. 48
- 4.38 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. 49
- 4.39 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. The right figure shows only the results of the time-based simulation. 50
- 4.40 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. 50
- 4.41 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. 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. 52
- 4.42 Same as Figure 4.41, 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 the text. 53
- 4.43 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. 54
- 4.44 Schematic overview of the setup that was used to test the implementation of the double-pass cluster-finding algorithm on a prototype of the read-out 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. 54

5	4.45	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.	55	5.3	Invariant mass spectra of the relevant signal resonances in the regions of interest after combinatorics and additional filtering for the five benchmark channels: (a) $\Lambda^0 \rightarrow p\pi^-$ (and charge conjugates), (b) $J/\psi \rightarrow e^+e^-$, (c) $J/\psi \rightarrow \mu^+\mu^-$, (d) $\bar{p}p \rightarrow e^+e^-$, (e) and $\bar{p}p \rightarrow e^+e^-\pi^0$. The black hollow histograms show the distributions for signal events, the red hatched histograms the corresponding distributions for background (DPM) events. Clearly the background levels are significantly reduced in all cases compared to Fig. 5.1.	5	10	10
15	4.46	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 <i>same eight inputs</i> , based on their superburst number, but has its <i>own output</i> . Hence, the FPGAs use 32 inputs and 16 outputs. More details can be found in the text.	55	5.4	Final detector's setup during the testbeam time at COSY.	15	15	64
20	5.1	Invariant mass spectra of the relevant signal resonances in the regions of interest after combinatorics for the five benchmark channels: (a) $\Lambda^0 \rightarrow p\pi^-$ (and charge conjugates), (b) $J/\psi \rightarrow e^+e^-$, (c) $J/\psi \rightarrow \mu^+\mu^-$, (d) $\bar{p}p \rightarrow e^+e^-$, (e) and $\bar{p}p \rightarrow e^+e^-\pi^0$. The black hollow histograms show the distributions for signal events, the red hatched histograms the corresponding distributions for background (DPM) events.	61	5.5	Photograph of the SADC prototypes connected to the Proto9.	20	20	66
25	5.2	Selection criteria for $\Lambda \rightarrow p\pi^-$ (a_1) - (a_4), $J/\psi \rightarrow e^+e^-$ (b), and $J/\psi \rightarrow \mu^+\mu^-$ (c), $\bar{p}p \rightarrow e^+e^-$ (d), $\bar{p}p \rightarrow e^+e^-\pi^0$ (e_1), (e_2). The black histograms show the distributions for Monte Carlo truth matched signal candidates (the truth match is applied to emphasize the difference between signal and background shapes), the red hatched area the one for background combinations. The blue dashed vertical lines indicate the filtering criteria. Plot (e_1) shows the invariant $\gamma\gamma$ mass to form π^0 candidates, where the shown window size already represents a selection of about ± 3 standard deviations around the peak position.	63	5.6	Photograph of the TRBv3 board. Two FPGAs on the right side perform functions of the DC. Another two on the left side: the top one is the UDP server, and the bottom one is the SODANET source.	25	25	66
30	5.7	Diagrams of readout schemes which were tested during the beamtime at COSY. The top two, (a) and (b), were used for tests only with Proto9. The bottom diagrams represent the cases when other subdetectors were connected in addition to the EMC prototype.	67	5.8	Distribution of detected hits between ADC channels.	30	30	68
35	5.8	Distribution of detected hits between ADC channels.	68	5.9	Energy distribution of hits coming from the ADC channel. Light pulser peak is located on the right side. Proton peak can be seen on the left side.	35	35	68
40	5.9	Energy distribution of hits coming from the ADC channel. Light pulser peak is located on the right side. Proton peak can be seen on the left side.	68	5.10	Discrepancy between ADC channels in case of proton peak position before alignment.	40	40	68
45	5.10	Discrepancy between ADC channels in case of proton peak position before alignment.	68	5.11	Divergence of time difference between EMC ADC channels and TOF system before alignment.	45	45	68
50	5.11	Divergence of time difference between EMC ADC channels and TOF system before alignment.	68					

5	5.12 Deposited energy distribution for a selected crystal. This distribution was obtained by combining the data from two SADC channels which correspond to the LAAPDs connected to this crystal.	69
	5.13 Time difference between two LAAPDs of the same crystal.	69
10	5.14 Time evolution of the time difference.	69
	5.15 Time difference-deposited energy correlation for the single crystal.	69
	5.16 Time difference after application of the energy cuts.	70
15	5.17 Profile representation of the time difference evolution.	70
	5.18 Time difference in the case of two DCs. The time shift caused by wire length difference is observed.	70
20	5.19 Time difference profile in the case of two DCs.	70
	5.20 Time difference between TOF and ADC channels in logarithmic scale.	70
25	5.21 Time difference between TOF and ADC channels.	71
	5.22 Cluster energy distribution in the case of time cuts.	71
	5.23 Time difference between TOF and ADC channels if energy cut was applied.	71
30	5.24 Time difference evolution (EMC+TOF).	71
	5.25 Time difference between TOF and FT sub-systems.	71
35	5.26 Time difference between TOF and FT in linear scale.	72
	5.27 Time difference between FT and all EMC channels.	72
40	5.28 Time difference between FT and all EMC channels in linear scale.	72
	5.29 Cluster energy distribution in the case of FT-EMC event coincidence when time cut was applied.	72
45	5.30 EMC-FT Time difference in the case of cluster and energy cuts	72
	5.31 EMC-FT Time difference evolution.	73

List of Tables

	3.1	Number of crystals with expected hit-rate for the barrel EMC.	20
	3.2	Summary of the EMC readout.	21
5	3.3	Estimated data rates per tracking station in High Resolution mode.	21
	3.4	Estimated data rates per tracking station in High Luminosity mode.	22
	3.5	Total hit rate for 320 HV-MAP sensors at $L = 2 \times 10^{32} \text{ cm}^{-1}\text{s}^{-2}$	23
10	3.6	Data rate at $L = 2 \times 10^{32} \text{ cm}^{-1}\text{s}^{-2}$ and 95% hit detection efficiency. The numbers do not include SODANET headers or other overhead.	23
15	5.1	Efficiencies for signal and background events after combinatorics.	60
20	5.2	Efficiencies for signal and background events after combinatorics and additional filtering, and the background suppression factors. The numbers marked with * represent the lower border of the one standard deviation confidence interval.	62
25	5.3	Efficiencies for signal and background events after combinatorics and additional filtering, and the background suppression factors. The numbers marked with * represent the lower border of the one standard deviation confidence interval. (DAY1 scenario with 75% EMC).	63
30			