

Development of the Readout Controller
for the
CBM Micro-Vertex Detector

Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich Physik
der Johann Wolfgang Goethe-Universität
in Frankfurt am Main

by **Borislav Milanovic**
from Indjija, Ex-Yugoslavia (Serbia)

Frankfurt, 2015
(D30)

vom Fachbereich Physik der
Johann Wolfgang Goethe-Universität als Dissertation angenommen.

Dekan: Prof. Dr. René Reifarth

Gutachter: Prof. Dr. Joachim Stroth

Prof. Dr. Peter Senger

Datum der Disputation: 25.6.2015

*I would like to dedicate this thesis to my family
for their infinite patience and support*

Abstract

The upcoming CBM Experiment at FAIR aims at exploring the region of highest net baryonic densities reproducible in energetic heavy ion collisions. Due to the very high beam intensities expected at FAIR, unprecedented data regarding rare observables such as charm quarks and hyperons will be accessible. Open charm mesons are particularly interesting, since they support the reconstruction of the total charm cross-section in order to search for exotic phenomena, e.g. a phase transition towards the quark-gluon plasma which is predicted by several theoretical models. Open charm studies will be performed via secondary vertex reconstruction with a suitable Micro-Vertex Detector (MVD). The CBM-MVD is currently in the development and prototyping phase with primary design goals concentrating on spatial resolution, radiation hardness, material budget, and readout performance.

CMOS Monolithic Active Pixel Sensors (MAPS) provide an excellent spatial resolution for the MVD in the order of few μm in combination with a low material budget (50 μm thickness) and high radiation hardness. The active volume of the devices is formed from the epitaxial layer of standard CMOS wafers. This allows for integration of pixels together with analogue and digital data processing circuits on one single chip. This option was explored with the MIMOSA-26¹ prototype, which integrates functionalities like pedestal correction, correlated double sampling, discrimination and data sparsification based on zero suppression combined with a pixel matrix of $\sim 2 \text{ cm}^2$. The pixel array composed of 576 lines of 1152 pixels is read out in a column-parallel rolling-shutter mode. One discriminator per column and the digital data processing circuits are located on the same chip in a 3 mm wide area beneath the pixel matrix allowing for binary hit encoding. This area also contains the circuits for pedestal correction and the configuration memory, which is programmed via JTAG². The preprocessed digital data is read out via two 80 Mbit/s LVDS links per sensor, which stream their data continuously based on a low-level protocol.

Within the scope of this thesis, a readout concept of the CBM-MVD is proposed and studied based on the current MIMOSA sensor generation. The backbone of the system is formed by the Readout Controller boards (ROCs) featuring FPGA³ microchips and optical links. Several ROC prototypes are considered using the synergy with the HADES⁴ Experiment. Finally, the TRB3⁵ board is selected as a possible candidate for the initial FAIR experiments. Furthermore, a highly scalable, hardware independent FPGA firmware is implemented in order to steer and read out multiple MIMOSA-26 sensors. The reconfigurable firmware is also designed with the support for future MIMOSA sensor generations. The free-streaming sensor data is deserialized and error-checked, prior to its transmission over a suitable network interface. In order to demonstrate the validity of the concept, a readout network similar to the HADES Data Acquisition (DAQ) system is developed. The ROC is tested on the HADES TRB2 boards and data is acquired using suitable MAPS add-on boards and the TrbNet protocol.

In the context of the CBM-MVD prototype project, a readout network with 12 MIMOSA-26 sensors has been prepared for an in-beam test at the CERN SPS facility. A comprehensive control system is designed comprising customized software tools. The subsequent in-beam test is used to

¹MIMOSA - Minimum Ionizing MOS Active sensor

²JTAG - Joint Test Action Group

³FPGA - Field Programmable Gate Array

⁴HADES - High Acceptance Di-Electron Spectrometer

⁵TRB3 - Trigger and Readout Board version 3

validate the design choices. As a result, the system could be operated synchronously and dead-time free for several days. The readout network behavior in a realistic operating environment has been carefully studied with the outcome the the TrbNet based approach handles the MVD prototype setup without any difficulties. A procedure to keep the sensors synchronous even in case of a data overflow has been pioneered as well. After the beam test, improvements and conceptual changes to the readout systems are being addressed which allow an integration into the global CBM DAQ system.

Kurzfassung

Das künftige CBM Experiment wird an der gerade im Bau befindlichen FAIR Beschleunigeranlage Messungen in Schwerionenstössen durchführen, um so Ereignisse mit höchsten Baryonendichten zu vermessen. Aufgrund äusserst hoher Strahlintensitäten sind so noch nie dagewesene Forschungsergebnisse im Bereich der Charm-Quarks- und Hyperonen-Studien zu erwarten. Insbesondere die Vermessung von Open-Charm-Verteilungen wird eine wichtige Rolle spielen um den totalen Charm-Wirkungsquerschnitt zu errechnen und so diese dichte Phase der Materie auf exotische Phänomene zu untersuchen, wie Quark-Gluon-Plasma, Phasenübergänge und die Zustandsgleichung des frühen Universums kurz nach dem Urknall. Hierbei können die Open-Charm-Teilchen durch die Rekonstruktion ihrer Zerfallsprodukte und sekundärer Vertizes durch einen Mikro-Vertex Detektor (MVD) identifiziert werden. Daher wird für den MVD neben einer exzellenten Ortsauflösung auch eine schnelle Auslesezeit, hohe Strahlenhärte und ein kleines Materialbudget benötigt. Zurzeit befindet sich der CBM-MVD in einer Prototypphase.

Die in einem CMOS-Prozess hergestellten Monolithischen Aktiven Pixel-Sensoren (MAPS) weisen eine exzellente Ortsauflösung von wenigen μm auf und dies in Kombination mit einem sehr dünnen Formfaktor von $50 \mu\text{m}$ Dicke. Ihr aktives Volumen wird durch eine Epitaxialschicht der gewöhnlichen CMOS-Wafer geformt. Dies ermöglicht die Integration der Pixel zusammen mit ihren analogen und digitalen Ausleseschaltkreisen direkt auf demselben Chip. Von dieser Möglichkeit wurde durch das Design des MIMOSA-26¹ Sensors ausgiebig Gebrauch gemacht, welcher, unter anderem, einen integrierten Schwellenabzug, korreliertes Double-Sampling, die Diskriminierung der analogen Signale und eine finale Datenkompression verbindet. Die Pixelfläche des Mimososa-26 umfasst 576 Zeilen von jeweils 1152 Pixeln mit einer Grösse von insgesamt $\sim 2 \text{ cm}^2$, wobei die Pixel in einer Zeile parallel ausgelesen werden (Rolling-Shutter Modus). Pro Spalte befindet sich auf demselben Chip jeweils ein Diskriminator in dem 3 mm breiten, nichtsensitiven Bereich unterhalb der Pixelmatrix. Diese Region enthält auch die notwendigen Spannungsumsetzer, z.B. zur Schwellengenerierung, welche von einer JTAG² Schnittstelle angesteuert werden. Die verarbeiteten, digitalisierten Daten werden über zwei LVDS Datenleitungen mit jeweils 80 Mbit/s ausgelesen. Dieser Datenstrom muss kontinuierlich ausgelesen werden und liegt in einem 16-Bit Datenformat vor, der im folgenden weiter verarbeitet, und mit den Datenströmen der anderen Sensoren des MVD zeitlich korreliert werden muss.

Im Rahmen dieser Arbeit wurde ein Auslesekonzept des CBM-MVD Detektors erforscht basierend auf gegenwärtigen Anforderungen der MIMOSA Sensorfamilie. Den Kern des Systems stellen Readout Controller (ROC) Ausleseboards dar, bestückt idealerweise mit FPGA³ Mikrochips und optischen Konnektoren. Unterschiedliche Prototypen sind betrachtet worden mit Hilfe der Synergie aus dem HADES⁴ Experiment, bevor schliesslich ein passender Kandidat für die FAIR Experimente mit dem TRB3⁵ Ausleseboard gefunden werden konnte. Des Weiteren wurde eine hochskalierbare, hardwareunabhängige FPGA Firmware entwickelt um mehrere MIMOSA-26 Sensoren zu kontrollieren und auszulesen. Die rekonfigurierbare Firmware unterstützt ebenfalls künftige MIMOSA-Generationen. Dabei wird der kontinuierliche Datenstrom jedes Sensors zunächst deserialisiert und nach Fehlern überprüft, um die Daten sodann an die

¹MIMOSA - Minimum Ionizing MOS Active sensor

²JTAG - Joint Test Action Group

³FPGA - Field Programmable Gate Array

⁴HADES - High Acceptance Di-Electron Spectrometer

⁵TRB3 - Trigger and Readout Board version 3

passende Netzwerkschnittstelle weiterzugeben. Hierbei wurde ein Auslesenetzwerk Prototyp geschaffen um die vorliegende Implementierung zu testen, wobei existierende Hardware (TRB2-Ausleseboards) und ein Netzwerkprotokoll (TrbNet) von HADES zum Einsatz kamen. Die Firmware wurde in den rekonfigurierbaren FPGA Mikrochips der HADES TRB2 Platine instanziiert und die Daten mittels spezifischer Erweiterungsplatinen aufgenommen und über das optische TrbNet Netzwerk ausgelesen worden. Die zur Verfügung stehende Auslesebandbreite betrug 100 MB/s, welche allerdings mit zusätzlicher Hardware erhöht werden kann.

Dieses System wurde im Rahmen eines CBM-MVD Prototypen eingesetzt. Insgesamt wurden 12 MIMOSA-26 Sensoren, angeordnet in fünf Ebenen, betrieben und ausgelesen. Neben dem notwendigem Auslesenetzwerk wurde ein vollwertiges Kontrollsystem entworfen bestehend aus individuellen Software Tools. Das Auslesekonzept konnte in einem Experiment am CERN SPS erfolgreich validiert werden. In diesem Experiment konnte das System synchron und ohne Unterbrechung für mehrere Tage angesteuert und ausgelesen werden. Das Auslesenetzwerk wurde so unter realistischen Arbeitsbedingungen genauestens überprüft mit dem Resultat dass der gewählte Ansatz mit TrbNet keinerlei Schwierigkeiten darstellt. Eine Methode, die Sensoren im Falle einer Netzwerküberlastung ohne Verlust der Synchronität auszulesen wurde ebenfalls getestet. Nach der Strahlzeit wurden Verbesserungen und konzeptuelle Änderungen ermittelt, welche den MVD für eine Integration in das finale CBM Auslesesystem vorbereiten sollen.

Contents

1	Physics Motivation and the CBM Experiment	1
1.1	New Phases of Matter	1
1.2	Heavy Ion Experiments at FAIR	4
1.2.1	Physical Background	5
1.3	The CBM Experiment	7
1.3.1	Detector Setup	8
1.3.2	Data Acquisition and Event Selection	10
1.3.3	Detection Capabilities	11
1.3.4	Micro-Vertex Detector	12
1.4	Thesis Overview	13
2	CMOS Technology and Applications	15
2.1	Introduction to CMOS	15
2.1.1	MOSFET Transistor and CMOS Circuits	16
2.1.2	Basics of the CMOS Manufacture Process	18
2.1.3	Latch-Up Effect in CMOS-Based Circuits	19
2.2	CMOS Pixel Sensors	21
2.2.1	Principle of Particle Detection in Silicon Sensors	21
2.2.2	Monolithic Active Pixel Sensors	23
2.2.3	Novel Technologies for MAPS Sensors	25
2.3	Field Programmable Gate Array	27
2.3.1	Theoretical Description	27
2.3.2	FPGA Introduction	28
2.3.3	FPGA Architecture	28
2.3.4	Radiation Effects	30
3	MIMOSA-26 Sensor Specifications	33
3.1	General Description	33
3.2	Digital Sensor Architecture and Operating Principles	34
3.2.1	JTAG Bus Interface	36
3.2.2	Sequencer	36
3.2.3	Comparators	38
3.2.4	Sparse Data Scan	41
3.2.5	Memory and Output	42
3.3	Sensor Operation and Characteristics	42
3.4	Next-Generation MIMOSA Sensors	44

4	Conceptual Design of the MVD Readout	47
4.1	CBM Micro-Vertex Detector at FAIR	47
4.2	Performance Requirements and Limitations	49
4.2.1	Simulation Method	49
4.2.2	Applied Sensor Model	51
4.2.3	Expected Data Rates at SIS-100	52
4.2.4	Sensor Limitations Study	54
4.2.5	Expected Hit Distributions	57
4.3	Interfaces to CBM Modules	58
4.3.1	CBM Network Protocol	60
4.3.2	FLES Interface	62
4.4	Proposed Readout Architecture	64
4.4.1	Choice of Technology	64
4.4.2	MVD Readout Topology	65
5	Readout Controller Development	69
5.1	Design Specifications	69
5.1.1	General Requirements	69
5.1.2	TrbNet Network Protocol	71
5.2	FPGA Firmware Design and Implementation	72
5.2.1	General Overview	72
5.2.2	Basic Modules	74
5.2.3	Sensor Specific Modules	75
5.2.4	Network Specific Modules	76
5.2.5	Proposed Extensions for CBMnet	79
5.3	Readout Controller Prototypes	80
5.3.1	TRB2-Based ROC Prototypes	80
5.3.2	TRB3-Based ROC Prototype	82
5.4	Development of a Readout Network Prototype	84
5.4.1	General Network Operation	85
5.4.2	MIMOSA-26 Front-End Electronics	86
5.4.3	JTAG Controller Implementation	87
5.4.4	Central Control Unit Implementation	88
5.4.5	TrbNet Hubs	89
5.4.6	Laboratory Tests and Results	90
6	In-Beam Tests and Results	93
6.1	Test Setup at the CERN SPS	93
6.1.1	MVD Prototype Core Module	93
6.1.2	Beam Telescope Setup	94
6.1.3	Power Distribution	95
6.1.4	Computer Network	95
6.1.5	Summary	96
6.2	Applied Software Toolkit	96
6.2.1	Event Builder	96
6.2.2	Sensor Threshold Settings	97
6.2.3	Slow Control and Monitoring Utilities	97

6.2.4	Safety Measures	100
6.3	Measurements	100
6.3.1	Preparations	101
6.3.2	Performed Tests	101
6.4	Results	101
6.4.1	Initial Results	101
6.4.2	Stability Analysis	102
6.4.3	Sensor Synchronization	104
6.4.4	Overload Study	106
6.4.5	Summary and Evaluation	106
6.5	Conclusions After the Test	108
7	Summary	111
	Zusammenfassung	113
A	Correlated Double Sampling with Clamping	121
B	Reduced MIMOSA-26 Bonding Scheme	123
C	Proposal of the Microslice-Container Encoding for the MVD	125
D	Readout Controller Firmware Parameters	127
E	Readout Controller Error- and Status-Bits	129
E.1	Data Checker	129
E.2	Chain Controller	130
F	CBMnet User Interface	131
G	MIMOSA-26 Front-End Electronics Connectors	133
H	Grounding Scheme of the Readout Network Prototype	135
I	Software Tools for Sensor Threshold Settings	137
I.1	Threshold GUI	137
I.2	Threshold Finder	137
J	HADES Event File Structure	139

List of Figures

1.1	Evolution of an ultra-relativistic heavy-ion collision	2
1.2	Phase diagram of strongly interacting matter	3
1.3	FAIR facility expansion and the planned CBM building	5
1.4	Kinetic freeze-out points of the phase diagram	6
1.5	Predicted particle cross-sections at FAIR using HSD	6
1.6	Expected particle yields per Au-Au collision at 25 AGeV	6
1.7	Horn structure of the K^+/π^+ ratio	7
1.8	CBM tracking stations inside the magnet	8
1.9	CBM detector setup	10
1.10	Basic structure of the CBM DAQ system	11
1.11	Different stages of a particle collision and the CBM tracking capabilities	12
1.12	MVD setup for the secondary vertex reconstruction	14
2.1	PMOS transistor	17
2.2	The NOR gate realized in CMOS	18
2.3	Lithographic process	20
2.4	Latch-ups in CMOS electronics	20
2.5	Structure of the MAPS pixel region	24
2.6	MIMOSA-26 in-pixel electronics	25
2.7	Depleted regions of normal- and high resistivity sensors	26
2.8	Radiation hardness of MIMOSA-26AHR	26
2.9	Triple-well MAPS	27
2.10	SOI technology	27
2.11	Example of a full adder implementation on the FPGA	29
2.12	Typical FPGA architecture	29
2.13	FPGA-slice architecture	30
3.1	MIMOSA-26 geometry and readout characteristics	34
3.2	MIMOSA-26 architecture with the data flow	35
3.3	JTAG chain example	36
3.4	MIMOSA-26 sequencer connectivity to other components	37
3.5	Rolling shutter timing diagram with the clamping operation	37
3.6	MIMOSA-26 comparator architecture	38
3.7	Signal voltage generation in the MIMOSA-26 pixel	39
3.8	MIMOSA-26 zero suppression (SUZE-1)	41
3.9	MIMOSA-26 data format	43
3.10	MIMOSA-26 sensor characteristics	44

3.11	Geometries of MIMOSA-26, MIMOSA-28 and MISTRAL sensors	45
4.1	MVD layout inside the magnet	48
4.2	Example of the implemented SUZE-2 encoding algorithm	51
4.3	Expected data rates at SIS-100, average case	53
4.4	Expected data rates per frame at SIS-100, worst case	53
4.5	Distributions of the sensor limitations	55
4.6	Number of data words generated by the most exposed FSBBs	56
4.7	Lost data resulting from the SUZE-2 logic	57
4.8	Average hit density of the most exposed sensor layer	59
4.9	Data distribution per sensor within the MVD	59
4.10	Efficiency of individual sensor classes	59
4.11	CBMnet jitter cleaner	60
4.12	CBMnet message structure	61
4.13	FLES structure	62
4.14	FLIB architecture and the first prototype	63
4.15	FLES interval building	64
4.16	Proposed layout of the MVD readout system	67
4.17	Global CBM readout scheme	67
5.1	TrbNet-based network architecture	72
5.2	General ROC FPGA layout	73
5.3	ROC readout chain architecture	74
5.4	Readout chain – input stage	75
5.5	Proposed data format	77
5.6	Readout chain – formatter	77
5.7	Readout chain – chain controller	78
5.8	Extensions for CBMnet	79
5.9	TRB2 board	83
5.10	MAPS add-on for TRB2	83
5.11	General Purpose add-on for TRB2	83
5.12	TRB3 board	84
5.13	Structure of the readout network	86
5.14	FEE prototypes	87
5.15	CCU layout	89
5.16	TrbNet Hubs	90
5.17	MIMOSA-26 output observed with a logic analyzer	90
5.18	Data from a MIMOSA-26 sensor irradiated with an X-ray source	92
5.19	JTAG chain with 8 sensors and 2 m length	92
5.20	Results from the test with a β -source	92
6.1	MVD prototype module	94
6.2	MIMOSA-26 beam telescope and layout of a reference plane	94
6.3	Architecture of the test setup	97
6.4	Sensor matrix with fake-hits at three different threshold settings	99
6.5	Applied monitoring tools	99
6.6	First in-beam test results	103

6.7	Spill structure after increasing the beam intensity	103
6.8	Missing frames of the high data rate study	103
6.9	Synchronization of the network components	105
6.10	Timestamp deviations of the three ROCs in the setup	105
6.11	Missing frames during the network overload phase	107
A.1	Correlated Double Sampling with clamping	121
C.1	Proposed MC encoding for the MVD	126
F.1	CBMnet user interface	131
G.1	RJ45 Ethernet cable modification	133
G.2	Pin assignment of the Converter Board	134
G.3	Pin assignment of the JTAG Queue Board	134
H.1	Global grounding scheme	136
I.1	Threshold GUI	138
J.1	HADES event file structure	139

List of Tables

1.1	Primary MVD design goals	13
3.1	Characteristics of different MAPS generations in comparison	45
4.1	List of applied simulations	50
4.2	Summary of studied limitations	54
5.1	FPGA processing time per package for the upcoming MIMOSA generations . . .	70
5.2	ROC resource consumption on the FPGA	84
6.1	Analysis of the timestamp deviation	105
6.2	Summary of performed in-beam DAQ tests	107
B.1	Reduced MIMOSA-26 bonding scheme - part 1	123
B.2	Reduced MIMOSA-26 bonding scheme - part 2	124
D.1	Readout chain settings for MIMOSA-26	128
E.1	Error bits - data checker	129
E.2	Error bits - chain controller (status bits)	130

Chapter 1

Physics Motivation and the CBM Experiment

1.1 New Phases of Matter

Quantum Chromo-Dynamics (QCD) is the theory of the strong nuclear force. It has the immense challenge to determine the laws of physics that regulate the behavior of subatomic particles which can neither be isolated nor observed (directly) in the real world. The particles are named quarks and form together with leptons (e.g. electrons) the fundamental building blocks of the visible universe. The force mediators of QCD are termed gluons due to their ability to confine quarks and 'glue' them together into compound states called hadrons. No quark can be observed freely in the laboratory. They are subject to an SU(3) gauge invariant theory [1] where the quarks and their corresponding antiparticles, the antiquarks, carry one of the three color charges or anticharges, respectively. All hadrons are 'colorless' objects meaning that the charge fields are confined within.

Being fundamental and all-encompassing, the QCD faces a multitude of questions to be addressed. The theory itself, however, is presently not fully solvable [2]. But systematic approaches do exist for specific problems or certain kinematic regions. For example, the perturbative methods (pQCD) provide plausible approximations in the high energy regime. On the other hand, the chiral perturbation theory is applicable only in the low mass region. Lastly, the Lattice QCD (LQCD) allows the study of QCD unrestricted by the energetic region, but it does not support dynamical quantities and is restricted by its own measures of simplification and mathematics. Thus, many questions are still open and the QCD not entirely explored.

Recent research activities led to the strong assumption that a new form of matter has been discovered, the Quark-Gluon Plasma (QGP) [3, 4, 5, 6, 7, 8, 9, 10, 11, 12], where quarks and gluons are not restricted to one hadron anymore but interact freely in a deconfined state over a large, thermally equilibrated region. If heavy ions get exposed to extreme temperatures above ~ 160 MeV, they presumably start melting and their quark and gluon contents mix facing a phase transition. Indications exist that a strongly interacting, perfect fluid is created.

The RHIC¹ collider in Brookhaven, USA, has thoroughly studied the new found state of matter, which was also present within the first microseconds after the Big Bang. The 'little' Bangs are now studied in particle colliders with sufficiently large beam energy. But the effects indicat-

¹RHIC – Relativistic Heavy Ion Collider




Figure 1.1: Different stages of an ultra-relativistic heavy-ion collision. Two ions (green and orange) collide producing an almond-shaped overlap region (red) (1). The large impact energy is sufficient to create new particles which rescatter and reach thermal equilibrium (2). Then, if the energy density and the volume are sufficiently large, the existence of a QGP phase is predicted (3). Due to the large pressure, the QGP is expanding and cooling down, similar to the early universe according to the Big Bang theory. Hadrons start emerging (4) and at some point stop interacting (5).

ing the QGP phase have not been found in proton-proton (p-p) or proton-nucleus (p-A) collisions. Heavy ions, preferably gold (Au) or lead (Pb), are needed instead. In these energetic nucleus-nucleus (Au-Au or Pb-Pb) collisions, a large volume is covered with high energy density which allows creation of new quark-antiquark pairs. The newly founded particles rescatter with each other and reach very quickly nearly perfect thermal equilibrium in this QGP dominated phase. The system temperature can easily exceed 300 MeV ($3.5 \cdot 10^{12}$ K). After the system expands and cools down, the hadrons begin to emerge which is known as the chemical freeze-out. Hadrons are present in a highly excited form of hadron gas [13]. The partonic¹ degrees of freedom are vanishing out of the system and only hadronic interactions remain. After a short while, even these cease to exist and the hadrons obtain their final state which can be measured with a suitable particle detector. This step is called the kinetic freeze-out. All five stages of the collision are depicted in Fig. 1.1.

RHIC experiments have discovered that no rapid phase transition exists at large collision temperatures, “but rather a gradual evolution from dominance of hadronic towards dominance of partonic degrees of freedom” [7]. This supports many QGP models, however a clear proof is still missing. Prominent QGP-phase indicators are observed anomalies in the sense of heavy quarkonia melting and jet quenching. Quarkonia are hadrons comprised of quark-antiquark pairs, e.g. the J/ψ , χ_c and ψ' mesons which all contain charm-anticharm ($c\bar{c}$) valence quarks². At present beam energies, they can be created exclusively during the initial stage of the collision. Due to the strong matter effects in the QGP, the $c\bar{c}$ bond is broken and the charm quarks rebind with lighter quarks during the chemical freeze-out [14]. Therefore, the initial charmonium is destroyed producing lighter hadrons with charm content. Such anomalous quarkonium suppression is considered as the main indicator of the QGP formation at RHIC energies [14, 15], but it can not be excluded yet that some other effect is responsible for the observed suppression. Furthermore, particle jets could get fully absorbed by the QGP [15, 16]. A jet (or di-jet, tri-jet) is a cascade of particles produced in the same direction originating from hard scattering processes on the partonic level. For example, two jets of opposing directions (‘back to back’) due to momentum conservation are often found in energetic heavy-ion collisions. But one of the two jets is found to lose its energy or gets entirely absorbed in the dense, hot medium which can be treated as an indicator of the QGP formation. Both effects are not present in p-p and p-A collisions.

¹The term ‘partons’ is used to denote both, quarks and gluons.

²Valence quarks are the effective constituents of hadrons which determine their quantum numbers. Sea quarks, on the contrary, remain purely virtual.

After the recent studies in combination with many theoretical predictions, the current picture of the QCD phase diagram related to the QGP formation is shown in Fig. 1.2. The diagram is related to the thermodynamic aspects of matter, since thermodynamic observables can be easily extracted from the final particle-state kinematics.

The temperature T dependence of the phase transition has already been demonstrated by RHIC and recent LHC¹ experiments, which is in good agreement with LQCD calculations and advanced hydrodynamic models [2, 18, 19]. The other chosen thermodynamical quantity is the baryochemical potential μ_B , which allows together with temperature the parametrization of pressure, and with it the equation of state (EOS) of nuclear matter. Both, T and μ_B are intensive parameters and therefore phase independent. They are specified simultaneously with the beam energy [20]. Following trends from the last century, particle accelerators have been designed to support increasingly higher beam energies, but with this strategy only the high-temperature region of the phase diagram can be explored. Therefore, several new facilities are being manufactured and some running experiments upgraded to reach into the less explored, low energetic, high μ_B region. Atomic nuclei are situated at low T , moderate μ_B in Fig. 1.2 (purple line). By increasing the temperature above approx. 160 MeV the matter enters a QGP phase following a smooth cross-over phase transition. Recent theoretical models additionally suggest that at zero T and high μ_B , the matter could enter the QGP stage as well, but this time following a rapid first-order phase transition [2, 21]. Thus, it appears possible that by compressing nuclear matter above some critical μ_B value, the QGP gets formed again. Such matter is found, e.g. in the core of neutron stars which experience densities above six times higher than that of atomic nuclei. If density is increased even further, the more exotic color superconductor phases are predicted.

Though, even more appealing is the existence of a critical point. Further LQCD calculations

¹LHC – Large Hadron Collider




Figure 1.2: The phase diagram of strongly interacting matter. The first order phase transition (red line) as well as the critical point have not been discovered yet. Therefore they divide the diagram into one part (to the left) which has been studied and well understood, and another (to the right, including the separation line and the critical point) which is open for speculations. Source: [17].

suggest its existence where the crossover phase transition changes into the first order transition. Some models even suggest that this might be a tri-critical point where three phases meet together, under the assumption that quarkyonic matter related to chiral symmetry truly exists [2, 21]. The search for this critical point is therefore highly motivated by the possibility to increase the knowledge about the QCD and improve its leading tools and models.

Chiral symmetry [2, 13], which has gained much attention lately, successfully predicts the low mass hadron spectrum with only few input parameters. Due to the breaking of chiral symmetry, the quarks acquire their mass by interactions with the chiral condensate, i.e. the QCD vacuum. Now, if the QGP phase is so dense that no chiral condensate is present inside, the chiral symmetry would get partially restored and the innermost particles could become massless. The chiral condensate related effects could therefore serve as order parameters to determine the phase transition and validate this significant theory.

The phase diagram will be investigated by many heavy-ion experiments around the world. The LHC collider at CERN, Geneva, currently concentrates on extreme temperatures. The RHIC collider is adapting its technological equipment to utilize lower beam energies and scan through the phase diagram (beam energy scan). Besides RHIC, the SPS¹ accelerator at CERN, will investigate the large μ_B region with the NA61 experiment. Also NICA², a new facility in Dubna will focus on the same, unexplored region. Furthermore, in the scope of the FAIR³ research program a new specialized detector is being developed with the purpose to explore the high μ_B region and focus more on rare probes which can not be measured with high accuracy by other facilities. The Compressed Baryonic Matter (CBM) experiment, as the name suggests, will create peak baryonic densities to explore the phase diagram and search for deconfinement, chiral symmetry restoration and the critical point. Being the subject of this thesis, the FAIR complex together with the CBM detector will be described in following.

1.2 Heavy Ion Experiments at FAIR

The present GSI⁴ Helmholtz Center, situated in Darmstadt, Germany, will be substantially upgraded during the following years. The SIS-18 synchrotron will be accompanied by the new SIS-100 and SIS-300 accelerators⁵. The new facility will run under the name FAIR - Facility for Antiproton and Ion Research [22, 23, 24]. The FAIR complex, as shown in Fig. 1.3, follows an interesting idea to locate both of the new accelerators in the same underground tunnel, one above the other. In such constellation, several experiments can be performed in parallel. The circumference of the tunnel will be approx. 1100 m. The SIS-100 and SIS-300 accelerators can deliver beam energies of up to 11 AGeV and 35 AGeV for Au beams, respectively. Thus, SIS-300 is expected to reach into the region of highest baryonic densities. However, the great advantage of the new facility lies not in the energy, but in beam quality and intensity. Both accelerators are expected to provide 10^{11} ions per second and 10^{13} protons per second.

¹SPS – Super Proton Synchrotron

²NICA – Nuclotron based Ion Collider fAcility

³FAIR – Facility for Antiproton and Ion Research

⁴GSI – Gesellschaft für SchwerIonenforschung (heavy-ion research association)

⁵The numbers in the name of SIS accelerators denote the magnetic rigidity in [Tm].




Figure 1.3: The planned FAIR facility. UNILAC and SIS-18 will be adopted from the present GSI facility to serve as pre-boosters. All the other components (red) have to be built from scratch. The CBM building responsible for heavy-ion experiments will comprise two detectors, HADES and CBM, operating simultaneously. Source: [24].

Under such favorable conditions, two fixed target detectors will study the low-temperature region of the QCD phase diagram - HADES¹ and CBM (see Fig. 1.3). The HADES detector [25, 26] is foreseen to continue its research on low-mass hadron spectrum with beam energies up to 8 AGeV [27]. The CBM detector is presented in section 1.3.

Besides heavy-ion studies, the FAIR facility will also conduct other related experiments. The APPA² collaboration concentrates on atomic and plasma physics, as well as their applications in bio-, medical- and material science. NuSTAR³ comprises a multitude of experiments regarding the structure of nuclei, nuclear astrophysics and radioactive ion beams. Lastly, the PANDA⁴ experiment will focus on hadron spectroscopy, strange and charm physics, as well as hypernuclear physics with antiproton beams. To support these experiments, the FAIR complex will feature additional storage rings and a proton linear accelerator.

1.2.1 Physical Background

According to [20], the region of highest net baryon densities for Au-Au collisions is estimated around 30 AGeV incident beam energy in the fixed target frame, which is accessible by FAIR. The calculations are based on a model derived from the results of recent heavy-ion experiments. The model is based on the fact that kinetic freeze-out properties need to obey some fundamental thermodynamical principles. The result is shown in Fig. 1.4.

In addition, a well established model, the Hadron String Dynamics (HSD), can be used to estimate the particle yields expected at FAIR energies. The results are shown in Fig. 1.5 and 1.6. As can be seen, probes containing charm quarks are produced very rarely in Au-Au collisions at 25 AGeV ($\sqrt{s_{NN}} \approx 7$ AGeV). But due to the high beam intensity, the effect of the low statistics

¹HADES – High Acceptance Di-Electron Spectrometer

²APPA – Atomic, Plasma Physics and Applications

³NuSTAR – Nuclear STructure, Astrophysics and Reactions

⁴PANDA – antiProton ANnihilation at DArmstadt




Figure 1.4: The kinetic freeze-out points in the QCD phase diagram for Au-Au collisions obtained with parametrized thermodynamic calculations. The region of maximum achievable net baryon density (and highest baryochemical potential) is found around 30 AGeV ($\sqrt{s_{NN}} \approx 8$ AGeV). Source: [20].

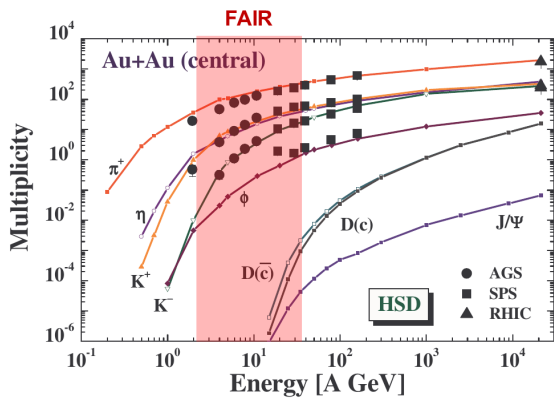


Figure 1.5: The predicted particle cross-sections using the HSD model. The AGS, SPS and RHIC data for lighter mesons are well described. Source: [28].




Figure 1.6: The expected particle yields per collision for Au-Au collisions at 25 AGeV calculated with the HSD model. The particles below the horizontal line could not be measured with high statistics up to now. Source: [29].

can be alleviated. This allows unprecedented studies of rare charm probes at maximum baryonic density. The huge advantage of produced charm particles is that they are created only during the initial stage of the particle collision at FAIR, prior to possible QGP formation. Being relatively heavy, their path is also less disturbed by lighter hadrons. However, during the QGP phase they are influenced by collective flow. Additionally, the charmonia may dissolve. Hence, charm quarks serve as ideal probes of the collision fireball and the QGP stage, as already indicated in section 1.1.

One of the first experimental approaches to the onset of deconfinement has been performed by the NA49 experiment at SPS. During the years 1999 - 2002, NA49 performed Pb-Pb collisions at 20, 30, 40 and 80 AGeV [30, 31], motivated by the statistical model which had predicted the onset of deconfinement in the region of $\sqrt{s_{NN}} = 7.5$ AGeV. Due to an increased number of degrees of freedom, the model estimated an enhanced strangeness production with respect to lighter mesons at the phase transition boundary. The point of 7.5 GeV is in good agreement with today's 8 GeV center of mass energy required to reach highest net baryon densities. One




Figure 1.7: The horn structure related to the onset of deconfinement shows a peak at ≈ 8 GeV center of mass energy. The strangeness production with respect to pions shows an anomaly related to the phase transition. The slope has been predicted by a statistical model based on deconfinement. Source: [30].

of the key results, augmented with recent measurements from RHIC and LHC, is given in Fig. 1.7. The strangeness production indeed shows an anomaly in the predicted region which can, up to date, only be explained by a statistical model with deconfinement [31]. Further anomalies related to the same region measured by NA49 and recently by RHIC suggest a phase transition of matter towards partonic degrees of freedom. FAIR will also access that region with the CBM Experiment. Charm particles are believed to exhibit a similar anomalous behavior at the phase boundary.

1.3 The CBM Experiment

As already indicated, the CBM experiment [2, 29, 32, 33, 34, 35] will explore the QCD phase diagram at highest net baryon densities to search for signs of the phase transition. Moreover, in-medium modifications of hadrons and the onset of chiral symmetry restoration can be addressed. Additionally, CBM will focus on new forms of strange matter, e.g. hypernuclei, double hypernuclei, and baryon-hyperon and hyperon-hyperon interactions. Also, open questions regarding di-baryons and metastable states can be addressed, e.g. $uuddss$, as predicted by the LQCD [35]. The studies are now possible due to the high beam intensity of the upcoming FAIR facility providing sufficiently high probability for these rare events to get detected. But in return, the CBM detector will have to support the extraordinary high collision rates.

The CBM detector [36, 37, 38, 39] will study heavy-ion collisions of 2-45 AGeV incident beam energy and is intended to support collision rates of up to 10 MHz. The expected data rates are predicted to be in the order of 1 TB/s, however they can be reduced down to ~ 1 GB/s with a sophisticated high-level trigger. Presently, the detector is still in the development phase. The physics program is expected to begin in 2019. The CBM detector will operate together with the HADES detector in the same experimental hall. The physics program of both detectors overlaps in the low energy region allowing for experimental cross-checks.




Figure 1.8: The CBM tracking stations inside the magnet. Source: [38].

1.3.1 Detector Setup

The CBM detector is designed with a fixed-target setup. A variety of subsystems is utilized to cover a wide range of possible heavy-ion studies at FAIR. The detector acceptance covers polar angles of 2.5° - 25° with full azimuthal coverage. The electronic components need to be radiation hard in order to sustain the expected radiation dose. Additionally, the sensors and support materials need to be designed as thin as possible to reduce multiple scattering. Moreover, fast operating sensors are required to support the highest beam intensities.

Tracking System

In order to detect individual particles and their decay topologies, particle tracks through various stations of silicon detectors need to be reconstructed. A superconducting dipole magnet with a magnetic field of up to 1 T is foreseen to bend the tracks of charged particles, enabling also the reconstruction of their momenta. Hence, the tracking system is placed inside the magnet. It contains two sub-detectors: the Micro-Vertex Detector (MVD) and the Silicon Tracking System (STS). They will apply silicon pixel sensors and silicon strips, respectively, for charged particle tracking. The number of planned stations is four for the MVD and eight for the STS.

An isolation envelope with dimensions of $1400 \times 2000 \times 1100 \text{ mm}^3$ is required to separate the tracking stations from the magnet. Both tracker will operate in vacuum, however the MVD will be placed in its own vacuum vessel. The entire tracking system is shown in Fig. 1.8.

Particle Identification

Different particle species require specialized detection systems, hence the CBM detector will feature following subsystems:

- **RICH:** Electrons usually move faster than other particles, due to their extremely low mass. In the Ring Imaging CHerenkov (RICH) detector, a radiation gas is placed with certain properties. The speed of light inside the gas is smaller than in vacuum. When particles, mostly electrons and positrons, are moving faster than the speed of light in that medium, they emit Cherenkov light in the shape of a cone along their path. The Cherenkov light wavelength is usually situated around the UV spectrum and can be reflected with a curved mirror onto an array of photo electrodes. The typical ring shapes are therefore found

whenever electrons (or in some cases pions) pass through, allowing their identification and coordinate determination. Recent studies pointed out that also pions which trigger the Cherenkov light can be sufficiently well separated from leptons with kinematic cuts, thus the RICH could also serve as a charged pion identifier.

- **MUCH:** Muons are directly detectable particles which lose the least energy when traversing through matter. They can be detected following a very simple principle. By placing a very thick lead plate to absorb all the hadrons, electrons, positrons and photons, the muons will be the only particle species that remains after the absorber. The MUon CHamber (MUCH) detector uses lead absorbers followed by a three-layer tracker, e.g. based on Gas Electron Multipliers (GEM), to detect and track muons. There are six such blocks planned for CBM. The MUCH tracker is needed to identify tracks found in the silicon tracker as muons and reconstruct their invariant mass spectra.
- **TRD:** Particles traversing an inhomogeneous medium (a radiator) are producing transition radiation which can be detected, e.g. via Multi-Wire Proportional Chambers (MWPC). The amount of radiation depends strongly on the particle species, as it is higher for relativistic leptons than for hadrons. This effect can be utilized to distinguish fast electrons from pions in the Transition Radiation Detector (TRD). The detector can additionally contribute to charged particle tracking. Currently, ten layers organized in three stations are planned with a total coverage of 585 m^2 .
- **TOF:** Pions, kaons and protons can be distinguished with a Time Of Flight (TOF) detector. Due to their relatively large mass differences, the time required for detector traversal differs as well. The momentum is obtained from the bending radius of the track inside the magnet. Together with the time of flight information, the mass of the particle can be determined distinguishing the denoted particles. Time information is acquired by first measuring the start of the collision with a suitable start detector. Afterwards, Multi-gap Resistive Plate Chambers (MRPC) situated further away provide sufficiently high timing resolution to obtain the final time of flight measurement. MRPC will cover an area of $15 \times 10 \text{ m}^2$ placed 10 m downstream from target.
- **ECAL:** The Electromagnetic CALorimeter (ECAL) is used to detect photons from light vector meson decays. It will be located 12 m from the target.
- **PSD:** The centrality and the reaction plane of the collision can be reconstructed by measuring the projectile spectators¹. The measurement is important for collective flow studies and will be performed by a hadronic calorimeter called the Projectile Spectator Detector (PSD).

In addition, several beam detectors will also be applied to detect a collision and to determine the beam position. More details on the current status of the subsystems can be acquired from [37, 38, 39]. All hadrons will be reconstructed from their decay topology and momenta. Because of the thick absorbers from the MUCH detector which do not allow full event reconstruction, the CBM program needs to be divided into two mutually exclusive parts: the muon setup and the electron setup. MUCH and RICH detectors will be mounted on a movable support structure and exchanged for each setup, respectively. The planned electron- and muon-setups are depicted

¹Spectators, in contrast to the participants, are nucleons of the ion that did not collide with the target.




Figure 1.9: The full CBM detector setup. Upper version will support full event reconstruction of all observables based on hadrons and electrons. Lower version is designed specifically for di-muon detection from vector-meson decays. Source: [29].

in Fig. 1.9. The start version of the CBM detector for SIS-100 will not contain all the detector stations, as planned for SIS-300. The TRD will contain only three to four out of its ten layers and the MUCH will start with only two absorbers and two tracking stations [35].

1.3.2 Data Acquisition and Event Selection

Eight sub-detectors have been mentioned in the previous section, each with its specific internal architecture. The data of interest are rare probes which occur with low probability. The use of a trigger is therefore advisable, however practically unfeasible. Many probes have to be reconstructed from their decay topology originating from the decay vertex which is displaced some mm (D-mesons) up to some cm (hyperons) away from the primary collision vertex. For this occasion, a full track reconstruction involving many subsystems (MVD, STS, RICH or MUCH, TRD) is required, which is a very challenging task [40, 41]. Up to 1000 particles need to be considered per collision with 85% fake combinatorial space points in the STS. A cellular automaton is developed for this task, followed by an application of the Kalman filter [42]. Thus, delimited by the computational complexity the CBM detector is designed as a free-running system. Autonomous, self-triggered Front-End Electronics (FEE) will stream the data with a rate of ~ 1 TB/s. The possibility of a physical trigger is currently ruled out, but the application of a high-level trigger is still possible. Therefore, the free-streaming data will be processed at runtime by a computer farm to reduce the data volume by several orders of magnitude leading to an effective storage rate of ~ 1 GB/s. The process is named First Level Event Selection (FLES) [43]. The FLES system is designed as a multicore-based computing cluster which utilizes graphics cards for track reconstruction.

The FEE will deliver asynchronous data messages on activation by a particle, which need to be organized into the corresponding events prior to track reconstruction. This is a highly challeng-




Figure 1.10: The generalized CBM DAQ scheme. Data is transmitted from the detectors to the computing cluster (FLIB/FLES). A designated timing network (TNet) provides synchronous control signals, which are forwarded to the ROCs, front-ends and sensors via the synchronous, deterministic CBMnet protocol. All front-ends receive a control message (e.g. a timestamp) at the same time. Source: [43].

ing task, since overlapping events need to be considered originating from different subsystems. Furthermore, a high precision time stamping mechanism is necessary for this occasion. More details can be found in [34].

Unifying the different sub-detectors with one readout network where synchronization of all components stands in the foreground is not a trivial task. Recent studies, however, suggest that bidirectional, high-speed optical links for clock distribution, time synchronization, control messages and data readout can be used for this task [44]. The applied network protocol is termed CBMnet [45]. Specific hardware components, based mostly on n-XYTER chips [46] are designed for various subsystems to prototype the readout. Data is forwarded to the Readout Controllers (ROCs) where time stamping on the synchronized global level can occur. Some systems might require a combination of certain data packages from different ROCs, in which case a Data Combiner Board (DCB) is planned for post-processing. Afterwards, the data can be transmitted for event reconstruction. A FLES Interface Board (FLIB) is used to buffer the data and prepare it for the FLES system. The generic full-scale data acquisition (DAQ) scheme can be viewed in Fig. 1.10.

1.3.3 Detection Capabilities

The detection capabilities and the physics program are large. Hyperons and other, more exotic particles containing strange quarks can be determined via their decay vertices from reconstructed STS tracks. They usually decay after few cm and can be easily tracked, also without the MVD. Additionally, the decay products can be distinguished by the TOF detector. Furthermore, di-electrons can be used to analyze low-mass vector-mesons, which can serve as an ideal tool to probe the collision fireball. For that, the standard CBM setup with the RICH is used (see Fig. 1.9). Next, the muon setup is designed to detect di-muons from rare charmonium and lighter vector-meson decays. Lastly, the open charm studies can be performed by observing the hadronic decay channel of D-mesons with the MVD. A list of all indirectly measurable particles based on track reconstruction is given in Fig. 1.11. Additional studies on collective flow, correlations and fluctuations are supported as well. As this thesis focuses on the MVD development, the MVD detector is introduced in the following section.

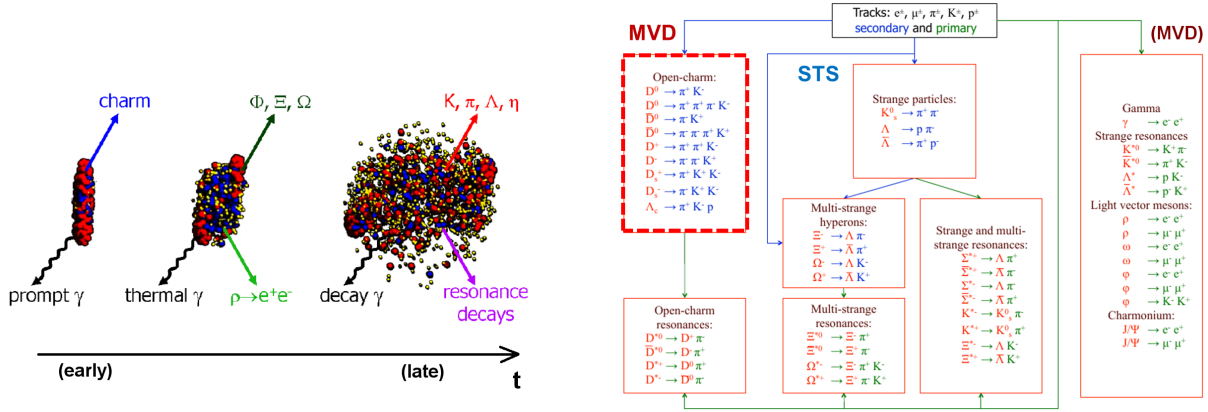


Figure 1.11: Different stages of the collision produce different particle spectra. Charm will be produced only in the early stage. The comprehensive list containing all observables reproduced by the tracking algorithm via their decay topology is shown to the right. The MVD will be mainly responsible for open charm reconstruction via their secondary decay vertices. In addition to that, tracking of di-leptons from the primary vertex and detection of photons via γ -conversion is planned as well. Source: [22, 38].

1.3.4 Micro-Vortex Detector

The CBM-MVD detector [47, 48, 49] relies on CMOS¹ pixel sensors for charged particle detection. Its main field of application is the secondary vertex reconstruction of rare probes. Apart from that, the MVD allows for tracking all charged particles, especially interesting of which are di-leptons emerging from vector-meson decays.

The secondary vertex reconstruction is necessary in order to obtain D-meson decays and contribute to the total charm cross-section. The mean decay path of D-mesons lies between 122.9 μm and 311.8 μm [50]. With a sufficiently large Lorentz boost, they might reach ~ 1 mm after the target. Nevertheless, they decay before they can be measured by any detector. But the coordinate at which they decay, the secondary decay vertex, is displaced from the primary collision vertex. Thus, with a sufficiently precise track reconstruction, the tracks of decay particles originating from the secondary vertex can be distinguished from the huge background produced by the primary vertex. The process is outlined in Fig. 1.12. For the secondary vertex resolution of $\lesssim 70$ μm , a spatial resolution of ≤ 5 μm is required [48]. Additional kinematic cuts are required to clarify the signal (see [51] for details).

The MVD needs to provide the necessary vertex resolution to reconstruct these secondary particle tracks (≤ 5 μm). Apart from that, the detector needs to fulfill many additional design goals. Being the first detector after the target, the material budget may not exceed $0.3 X_0$ for the first, and $0.5 X_0$ for the following stations in order to reduce multiple scattering which impairs the momentum resolution of all particles. The detection efficiency needs to reach 99.5% in order to detect rare probes. Since CBM relies on high collision rates, the readout time needs to be accordingly fast (≤ 30 μs). Lastly, the radiation hardness needs to be optimized in order to sustain forefront particle bombardment. All these primary criteria are summarized in table 1.1.

¹CMOS – Complementary Metal-Oxid Semiconductors

Requirement	Goal
Spatial resolution	5 μm
Material budget	0.3 % - 0.5 % X_0
Detection efficiency	99.5 %
Time resolution	30 μs
Ionizing radiation hardness	3 MRad
Non-ionizing radiation hardness	10^{14} n _{eq.} /cm ²

Table 1.1: The principal CBM-MVD design goals required for open charm measurement. Additional prerequisites involve: sufficient cooling, vacuum operation, mechanical stability, and power consumption.

1.4 Thesis Overview

This thesis focuses on the design of a highly scalable, free-streaming MVD readout system. The sensors, as well as the main readout components are based on CMOS microcircuits. Therefore, an introduction into the CMOS technology is given in chapter 2. Besides describing the basic CMOS properties, the chapter introduces CMOS pixel sensors and the reconfigurable FPGA¹ microchips. Furthermore, in chapter 3, the MIMOSA-26 sensor is used to introduce the principles and basic characteristics of the applied MVD sensor technology. The sensor serves as a basis for all further MVD studies, as it allows for developing the first MVD prototype and the corresponding readout electronics. Chapter 4 specifies the necessary design parameters for the MVD readout system. At first, the most constraining performance requirements are obtained from recent detector simulations, followed by the specification of crucial CBM interfaces. After discussing the integration of the MVD into the global CBM readout system, a generalized MVD readout concept is proposed.

Chapter 5 presents the implementation of a readout system fulfilling the basic prerequisites derived from the preceding studies. In order to reduce costs and time, hard- and software from the HADES detector has been used abundantly. An adaptable FPGA firmware module is developed to support the readout of MIMOSA-based sensors. The module is integrated into the first three readout network prototypes which are presented and analyzed within the same chapter. One of the network prototypes has found application at CERN, operating a system of 12 MIMOSA-26 sensors during an in-beam test which is discussed in chapter 6. After describing the CERN setup, the applied software toolkit with all the security features is presented. A comprehensive list of the performed studies is included, however only the data acquisition related part is evaluated. This thesis concludes with a summary in chapter 7.

¹FPGA – Field Programmable Gate Array

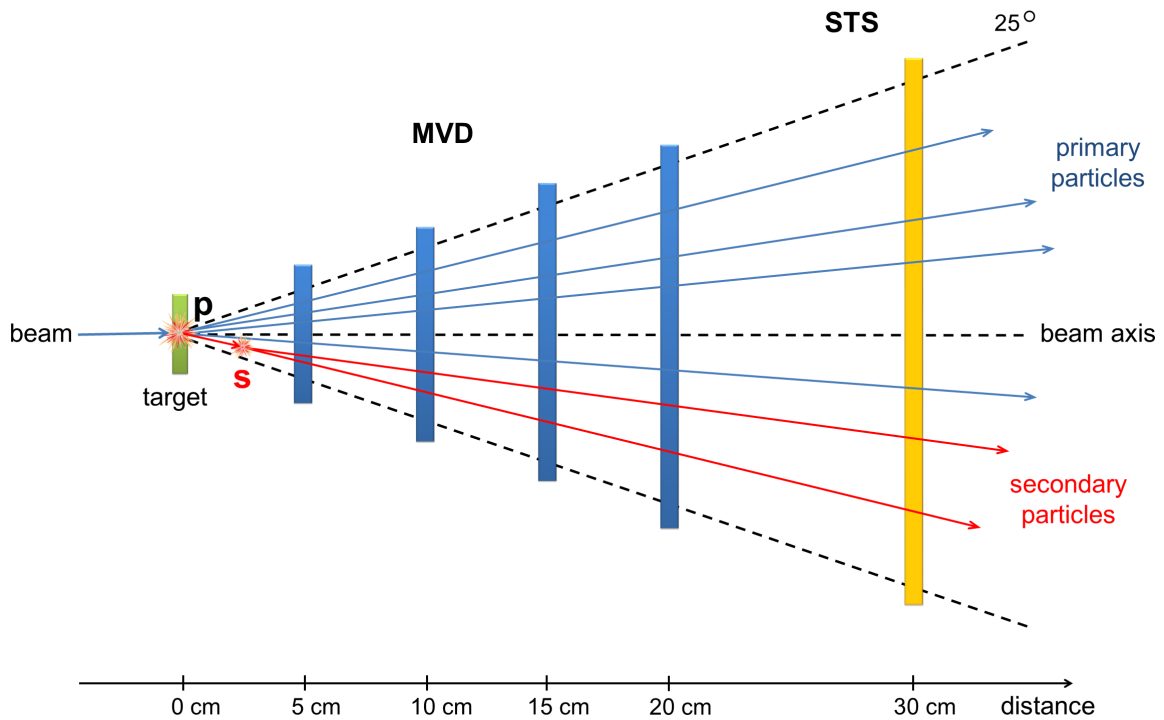


Figure 1.12: The principle of the MVD secondary vertex reconstruction. After the collision, many particles originate from the primary collision vertex **p**. However, D-mesons decay within few 100 μm , before reaching the first MVD station. Their hadronic secondary decay vertex position **s** can be reconstructed with the help of the MVD. Based on the MVD spatial resolution, primary and secondary vertices can be distinguished. The bending of charged particle tracks inside the magnetic field is not displayed for simplicity reasons.

Chapter 2

CMOS Technology and Applications

This thesis shows an implementation of a readout system of CMOS-based pixel sensors. The readout itself utilizes CMOS components as well. Therefore, an introduction to the CMOS production process, the motivation behind it and some applications for pixel sensors and reprogrammable microchips are presented in this chapter.

As first, an introduction to CMOS technology is presented in section 2.1 which explains the MOSFET transistors, the planar process and latch-up formation in CMOS devices. Afterwards, section 2.2 outlines the basic principle of particle detection in silicon and introduces the CMOS Monolithic Active Pixel Sensors (MAPS), together with some recent improvements in the field. Lastly, section 2.3 encompasses the Field-Programmable Gate Array (FPGA) microchips which are frequently used in experimental physics for data acquisition and control systems.

2.1 Introduction to CMOS

CMOS stands for Complementary Metal-Oxide Semiconductors [52] and denotes a production process based on the idea that emerged in the 1960-ies and nowadays dominates the electronics market. The basis are MOSFET¹ transistors, the essential elements in all modern electronic devices.

The advantage of the CMOS technology is that it applies two MOSFET types simultaneously (PMOS and NMOS, see the following section) to substantially reduce the power consumption and noise. The disadvantage of this approach is an increased sensitivity to hazards [53] and the doubled number of applied transistors². Also, wells have to be etched in the silicon structures due to the reasons which are explained in section 2.1.2. However, the unwanted effects can be corrected with modified domino logic [52]. Domino logic incorporates alternating blocks of PMOS- and NMOS-based logic circuits. The final device is in complementary MOS technology, though its individual blocks contain only one transistor type.

¹MOSFET – Metal-Oxide Semiconductor Field Effect Transistor

²However, the load resistors applied by the pure NMOS and PMOS technologies are not needed anymore. Thus some resources are recovered.

2.1.1 MOSFET Transistor and CMOS Circuits


The MOSFET transistor [52, 54] has a large area of applications. Mostly used as an electronic switch, it allows the design of complex digital circuits. But also applications as an amplifier are found frequently in analog devices. The MOSFET transistor has three basic connectors: the source, the drain and the gate¹. The gate is acting as a (digital or analog) switch. It is usually realized with polysilicon which is a moderate conductor. Once an appropriate voltage is applied, the gate creates a conductive channel between the source and the drain allowing charge carriers to pass through. Thus, the switch is closed and the electronic circuit is conductive. The amount of applied gate voltage over some threshold value controls the size of the channel². There are however saturation effects, as well as some unwanted parasitic effects which should be taken into account when developing microelectronic circuits. From a different perspective, the gate controls the resistance of the channel from infinity to some practical, low value. If the gate voltage is switched off, the conductivity is interrupted and resistance rises to infinity. This is the basic principle of MOSFETs in the enhancement mode (normally-off). Many other types exist, e.g. the depletion type MOSFETs (normally-on) where a gate voltage has the exactly opposite effect.

Certain material properties are required before the MOSFET can be put to use. Using a doping process, the silicon conductive capabilities can be altered. Silicon has four valence electrons in the atomic shell. By implanting elements with five valence electrons (donators, **n-doping**), e.g. phosphorus, arsenic or antimony, the silicon crystalline structure is modified. The four silicon valence electrons create a bond with the four donator valence electrons, but one electron in the phosphorus remains loosely bound, thus it can participate in charge transport. On the other hand, by implanting elements with three valence electrons (acceptors, **p-doping**), e.g. boron, aluminum or indium, only three valence electrons can create a bond. Thus, again, one valence electron is loose, only this time from the silicon atom. But it is not used for charge transport, instead, it is more energetically favorable for the acceptor atom to capture the loose electron and form four bonds with the neighboring silicon atoms, leaving the initial silicon atom with a defect electron, a so-called hole. In such p-doped materials, these holes are used for charge transport.

The MOSFET transistor requires differently doped junctions between the source, the substrate and the drain regions. For example, in a p-n-p doped MOSFET, the majority charge carriers in the n-doped substrate are electrons. The p-doped source and drain use holes for charge transport but due to the n-doped region in between, the holes from the source can not reach the drain since they immediately recombine in the substrate. Thus, the transistor is off (normally-off mode). The gate is placed above the region between the source and the drain. A thin silicon dioxide layer (oxide) galvanically isolates the gate from the p-n-p junction. However, by applying a negative gate voltage, an electric field is created that reaches into the substrate and counters the particle recombination. The loosely bound valence electrons get pushed away from the gate electrode leaving the region depleted of its majority charge carriers. If the gate voltage is high enough, i.e. exceeding some threshold value, the depletion region will be sufficiently large to allow holes to travel through the polarized n-doped but now positively charged material. This effect is called **channel inversion**. An inverted, positive channel in n-doped substrate is created by minority charge carriers due to depletion. For a fixed gate voltage the size of the channel obtains certain

¹The bulk connector is not mentioned here for simplicity reasons.

²This effect is used for the purpose of the amplifier.



(a) In the unbiased mode, the charge carriers can not reach the drain electrode. They recombine immediately in the substrate. Hence, the transistor is closed.

(b) Only after applying an appropriate gate voltage, the conductive channel gets formed depleting the substrate region of its majority charge carriers. The transistor is now open.

(c) A characteristic PMOS transfer function.

Figure 2.1: An illustration of the PMOS transistor.

dimensions allowing only a certain amount of charge carriers to pass. Thus, there are additional saturation- and parasitic effects (see [52] for details). Since positive charge carriers, the holes, are used for transport the p-n-p transistor is named **PMOS**. Equivalently, an n-p-n transistor uses electrons to traverse the depleted region, therefore it is named **NMOS**. One practical difference between the two types is that the electrons are more mobile than holes which allows the NMOS transistor to switch faster. The PMOS transistor is shown in Fig. 2.1.

Any conventional CMOS device uses alternating NMOS and PMOS transistors to perform its task. An example is given in Fig. 2.2 showing the NOR gate [52] realized in CMOS technology. The NOR gate implements the logic function:

$$\text{NOR} : \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$$

$$\text{NOR}(0, 0) = 1, \text{ else } 0$$

which is '1' only when both inputs are '0'. The logical '0' is some negative voltage sufficient to create the channel in a PMOS transistor. Conversely, the logical '1' is the positive voltage needed to activate the NMOS transistor. Note that a '0' does not have any effect on NMOS, and vice versa, if the voltages are in the conventional range of few volts. If either of the inputs provide a logical '1', the path from the supply voltage U to the output will be cut off by the PMOS transistor and the ground will be connected by the NMOS instead. However, if both inputs are '0', the PMOS transistors will supply the voltage to the output pulling it up to a logical '1'. Since a '0' does not affect NMOS, the ground will be disconnected. The NMOS transistors form a pull-down network, while the PMOS transistors form a pull-up network to realize the NOR gate in this design.

In principle, the component dissipates power only during switching, i.e. as long as the charge carriers need to pass from the source to the drain. Once they pass, the output potential is fixed and no current can flow anymore. This is a huge advantage with respect to other technologies, e.g. bipolar-transistor based logic families or pure NMOS logic¹. The length of the channel be-

¹The NMOS and PMOS technologies use pull-up and pull-down resistors to define their logic functions, which

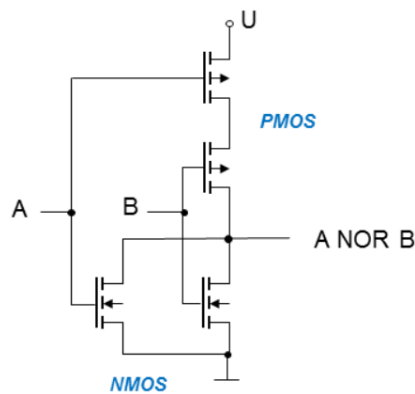


Figure 2.2: The CMOS implementation of a NOR gate. Source: [53].

tween the source and the drain determines its basic characteristics and serves as an international measure of the CMOS process precision. Currently, the channel length reaches down to a value of 22 nm allowing less current to flow hence minimizing the power consumption and supporting higher clock frequencies. Smaller transistor size also allows higher logic densities.

2.1.2 Basics of the CMOS Manufacture Process

Moore's law [55], which predicts a doubling in computing resources every 1.5 years, has been kept valid for decades mostly thanks to improvements in the so-called **planar technology**. Due to optimizations of the chip yield and the underlying feature size it is widely applied in the microelectronics sector, including CMOS-based pixel sensors for nuclear physics.

The planar process [52, 53, 54, 56] gradually incorporates miniature CMOS structures on a silicon wafer. In the first step, raw silicon needs to be obtained with an impurity concentration of around 10^{-9} . This is performed by melting conventional sand and processing it chemically. Hereby, the desired doping concentration can be achieved by admixing the donators or acceptors into the molten silicon. The pure silicon rod is then obtained either with the Czochralski process [52] or the float zone method [56]. Subsequently, the rod is sliced into wafers with a thickness of few hundreds micrometer. Afterwards, the wafer is stored at temperatures from 800 - 1200 °C in an oxygen rich environment. This allows the formation of a protective silicon dioxide layer (oxide) to suppress impurities on the surface. Now the wafer is ready to commence the CMOS integration process.

Nanometer- to micrometer-sized structures are selectively doped in order to create the source and the drain, as well as other CMOS relevant structures, e.g. diodes, vias and wells. This is done by performing a series of **photolithographic steps**:

1. The wafer oxide is covered fully with a photoresist, e.g. using sputtering.
2. A mask and optical lenses are used to irradiate microstructures on the photoresist with UV light. After exposure to UV light, the areas can be removed, freeing the layers below for further processing. The areas which were not exposed to UV light due to the mask geometry remain protected.

may dissipate additional power even when the transistors are inactive.

3. One of the following steps can now be performed:

- (a) Etching: the free areas can be opened even more to access deeper layers below the photoresist. For example, some steps require removing the protective oxide layer or a fraction of the silicon. This can now be performed on photoresist-free places with certain acids, e.g. hydrofluoric acid (HF).
 - (b) Doping via diffusion: the etched wafer is placed in a furnace and heated up to 1200 °C. The dopant is provided either as a gas or a liquid and enters the wafer via diffusion. It will enter all the areas that have been etched free during the previous iteration. For this step, the photoresist will melt, therefore a thick oxide has to be used instead to protect the areas which should not be doped.
 - (c) Doping via implantation: the dopant atoms are ionized, accelerated and shot into the wafer. The photoresist can act as a shield, protecting closed areas. The depth can be adjusted with the ion energy. The doping is thus more uniform than with diffusion. However, the ions produce large defects in the crystalline structure and are not electrically active. Therefore, the wafer is treated thermally after the implantation for dopant activation and annealing [56].
 - (d) Oxidation: the oxidation procedure has to be repeated in order to cover the processed areas with a protective layer.
 - (e) Metallization: the contacts are connected together using Aluminum. Often, several metallization layers are required. Copper, although being the better conductor is not used due to the fact that it produces impurities in silicon near the contact. With aluminum, the effect is reduced.
 - (f) Passivation: the areas are protected against mechanical damage and chemical contamination with special substances, e.g. silicon nitride, polyimide or silicon oxide [56].
4. In order to complete the entire microchip, the above procedure needs to be repeated with a different mask. Thus, a new cycle including all the previous steps is started. One iteration allows to process only one microstructure, e.g. the gate of some transistors. With the subsequent cycle, the drain could follow (after the appropriate etching). Microelectronic circuits require hundreds of these cycles until they are produced.

The photolithographic step is visualized in Fig. 2.3. CMOS processes are restricted either to n- or p-type wafers. Therefore, implementation of only one transistor type NMOS for p-type wafer or PMOS for n-type wafer is possible. To overcome this obstacle, so called **wells** are implanted to change the doping locally to p- or n-type. For example to implement the PMOS transistor in p-type base wafer, a well with n-type doping (n-well) is required.

2.1.3 Latch-Up Effect in CMOS-Based Circuits

The negative side-effect of well-implantation is the exposed risk to so-called latch-up effects [57]. In CMOS-based electronics there is a possibility of induced current flow into the substrate that can short-circuit the device. Therefore, latch-ups can be very dangerous.

An example is shown in Fig. 2.4. There, PMOS transistor requires an n-well due to the p-doped substrate. This is common for the CMOS technology. But it also gives rise to dangerous




Figure 2.3: The lithographic process uses UV or X-ray light originating from a point source. The light is optically directed through a mask and projected onto the wafer, creating miniature structures which can be doped or metallized during subsequent steps. Multiple identical microchips are produced on the same wafer within the process. Source: [53].




Figure 2.4: The latch-up can occur when the n-well creates dangerous bipolar transistor structures T1 and T2, as shown in the left image. If one of them gets activated by chance, an infinite feedback-loop could be initiated producing a short-circuit. The schematics of the produced latch-up circuit is shown to the right. The circuit is identical with a thyristor, only this time the thyristor gate gets activated randomly.

structures T1 and T2. They represent bipolar transistors and do not require a depletion zone to get activated. Some voltage drop in the basis¹ of T1 is sufficient to induce some holes into the substrate. The holes are there minority charge carriers and can reach long distances. Therefore, with a low probability, they can open the basis of T2 and induce a current from the n-well to the ground. In return, the induced current will open T1 even more which creates an infinite feedback-loop. This produces a short circuit between V_{DD} and the ground. The connection has low impedance and the device will dissipate more power which can lead to rapid overheating. The logic function of the device is changed as well.

Main cause of latch-ups are random voltage spikes, but also radiation can induce charges into the harmful structures. The risk can be minimized by modifying the layout of the circuit or by changing its material properties. For example, by decreasing the substrate resistivity with an increased doping concentration, the induced charge carriers could diffuse away before the critical voltage builds up. An epitaxial layer on top of the substrate can also improve latch-up immunity.

¹The bipolar transistor has three connectors – emitter, basis and collector in analogy to source, gate and drain.

2.2 CMOS Pixel Sensors

The technological benefits of the CMOS technology can be applied to develop CMOS-based sensors for imaging and particle detection. They have found applications in various fields of science, e.g. nuclear physics, solid-state physics, astrophysics, biology, and medicine.

2.2.1 Principle of Particle Detection in Silicon Sensors

When charged particles pass through matter, they lose energy by interacting with the electromagnetic fields inside atoms and also due to scattering with the electrons and nuclei. The energy loss is quantitatively described in [50]. Due to the fact that valence electrons are loosely bound in silicon atoms, the impinging particle creates a cloud of electron-hole pairs along its path. The charge carriers are diffusing around until they lose the surplus energy and recombine. If they are collected before the recombination, the impinging particle can be detected.

The silicon material is always doped in order to decrease its resistivity and allow the charge carriers to reach longer distances. There are high-resistivity CMOS processes with a low doping concentration and low-resistivity ones with high-doping concentration. Silicon is more suitable to detect particles than gas detectors for several reasons [54]. Charge carriers require little energy to get activated (~ 3.6 eV), therefore more of them can be used for transport and consequently detectors are thinner. Electrons have high mobility allowing fast detection within some nanoseconds. Silicon is very dense, so particles will lose more energy per unit length. Also the material rigidity is excellent demanding less-complex support structures.

The energy loss of particles inside matter can be expressed in terms of radiation length X_0 . By definition, this is the mean distance over which a high energetic electron loses all but $1/e$ of its energy due to bremsstrahlung [50]. For silicon, the X_0 amounts to 9.370 cm. For a thin layer Δx the energy loss can be calculated:

$$\Delta E = E \cdot \frac{\Delta x}{X_0}$$

where E is the particle energy [54]. On average, 80 electrons are freed in thin silicon material per micrometer of traversing MIP¹.

Silicon pixel sensors need to fulfill several requirements in order to be applied for charged particle detection in physics experiments. In section 1.3.4, the principle of secondary vertex reconstruction has been outlined. This process is termed heavy flavor tagging and requires the detector to be placed in the vicinity of the target. Therefore, the sensors have to be very thin in order to reduce multiple scattering of traversing particles. Multiple scattering is an unwanted effect resulting from passage of particles through matter. The initial path is tilted by an angle

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta \cdot c \cdot p} \cdot Z \cdot \sqrt{x/X_0} \cdot [1 + 0.038 \ln(x/X_0)]$$

¹MIP – Minimum Ionizing Particle

with p and Z being the particle momentum and charge, respectively [50]. The thickness of the material per radiation length x/X_0 plays the major role in this equation. The thicker the material, the larger the scattering angle. With more detector stations the initial particle track gets more and more deteriorated which decreases track accuracy and the momentum resolution.

Another important issue emerges from the proximity to the collision vertex. The sensors are exposed to a substantial amount of radiation. There are two major types of radiation harmful to silicon sensors – ionizing and non-ionizing radiation [58]:

- **Ionizing Radiation:** The ionizing radiation is formed by ions and photons which create electron-hole pairs in the material. This effect, although contributing to particle detection is unwanted in regions outside the sensing area, especially in the oxide. The oxide is an insulator. The mobility of electrons, although being very low, is much larger than the mobility of holes in the oxide (up to a factor of twelve [59]). Therefore, the electrons can often leave, but the holes remain trapped inside. This can contribute to large electric fields generated by the space charges of the trapped holes which can change the oxide properties. For example, on the boundary between the oxide and the silicon some defects often emerge which induce electric fields into the MOSFET substrate so that a small conducting channel is opened, even when the gate is not activated. An increased leakage current is the consequence and the sensor noise profile enhanced.
- **Non-Ionizing Radiation:** The non-ionizing radiation is measured in neutron-equivalents per area (n_{eq}/cm^2). It has the negative effect that when fast neutrons (or protons) collide with a silicon nucleus, they push the entire atom out of the crystalline lattice. The atom diffuses through the material producing more defects along its path. When it nearly loses its entire kinetic energy, it creates a large defect area around the region where it stops. All the defects reduce the material's conducting properties.

Radiation hardness can not be achieved following one single design rule. Many factors must be considered instead, e.g. technology, design, fabrication procedure, radiation environment, and many more [60].

Some techniques have emerged to reduce the ionizing radiation effects caused mainly by trapped holes in the oxide [59, 60]. Guard rings are used to enclose the transistor, or specific areas and minimize the leakage current and some other potential parasitic effects caused by ionizing radiation [59]. Afterwards, some geometrical aspects can be considered during the transistor design, e.g. use of edgeless transistors, or the channel-stop technique as proposed by [60]. Also the method of Post-Oxidation Annealing (POA) under the right conditions can decrease the hole trapping rate to improve radiation hardness. After the oxidation step in the planar process (see section 2.1.2), the oxide layer can be thermally annealed at 1000 °C in an oxygen-rich environment with admixture of argon or nitrogen. This has proven to decrease the defects in the oxide which are responsible for hole trapping. The hole trapping can also be reduced by implanting electrons in the oxide. Deep electron traps can be used to counteract the holes after their creation, before they can reach the substrate boundary. An immediate benefit of the improvements in the general planar process can also contribute to ionizing radiation hardness. The transistor size is shrinking for decades leading to thinner gate-oxides, which in return must possess better quality to supply the high field intensity necessary for transistor operation. This substantially reduces

the probability of charge trapping. Recent studies show that thermal annealing can also be used to successfully recover the original, non-irradiated characteristics to some extent [58, 61].

However, none of these methods can be effectively applied to reduce the impact of non-ionizing radiation damage. But there are some possibilities to circumvent this problem. Pixel sensors, for example, can be operated at low temperatures to minimize the negative side-effects of both ionizing and non-ionizing radiation. In addition, the pixel pitch can be reduced to increase the radiation hardness [62, 63]. Apart from that, the sensors can profit from a high-resistivity sensing volume to focus the charge collection to the seed pixel. More details on the latter issue are given in the following section.

2.2.2 Monolithic Active Pixel Sensors

The necessity for thin, fast and radiation hard sensors with exceptional spatial resolution required for secondary vertex reconstruction has led to a new pixel technology based on the commercialized CMOS process. The Monolithic Active Pixel Sensors (MAPS) [64, 65, 66, 67, 68, 69, 70, 71] create a good compromise between many of the requirements imposed by nuclear physics experiments. In addition, the sophisticated CMOS processes allow fast deployment and substantially reduce the costs otherwise needed to research new, specialized technologies. The first CMOS wafer, the prototype, is costly but every subsequent wafer has costs reduced by one or two orders of magnitude. Monolithic means that the pixel microcircuits, as well as their readout logic are present on the same chip. Therefore, expenses invested into signal digitization and sensor readout are reduced, as well as the size and thickness of the foremost readout electronics.

Minimum Ionizing particle MOS Active sensors (MIMOSA) are a series of MAPS developed at the IPHC¹, Strassbourg. Being the state-of-the-art, they will be taken as an example to explain the progress in this field. During the years 1999 to 2012, over 30 MIMOSA versions have been produced. MIMOSA sensors detect particles via their energy loss in the sensing material, as outlined in section 2.2.1. In order to improve the detection efficiency, an epitaxial layer of 10 – 20 μm thickness [68] is grown on top of the highly p-doped substrate. With MIMOSA-26 from 2009, the series finally became suitable for large-scale experiments. The sensor is used to equip the MVD prototype and will be explained in detail in chapter 3. Apart from that, MIMOSA chips are applied in the Heavy Flavor Tracker of the STAR experiment. They have additionally been used to prototype the ILC² detector, and they are considered for the ALICE³ detector upgrade in future.

MIMOSA-26 is designed with an epitaxial layer thickness of 14 μm . The substrate has low resistivity below 1 Ωcm , whereas the epitaxial layer reaches around 10 Ωcm [70]. Due to differing electric potentials at the boundary between the substrate and the epitaxial layer, electrons generated by charged particles traversing through the epitaxial layer are trapped inside. Based on thermal diffusion, electrons move through the sensing material until they are collected over regularly implanted n-well diodes [67]. However, they might also recombine before they reach the diode. The lifetime of generated electrons is in the order of few μs in a common doping level of 10^{15} atoms/ cm^3 [66].

The pixels use active components, such as diodes and transistors for readout and amplification,

¹IPHC – Institut Pluridisciplinaire Hubert Curien

²ILC – International Linear Collider

³ALICE – A Large Ion Collider Experiment




Figure 2.5: The structure of the MAPS pixel region. The epitaxial layer is used to trap electrons inside. They can be collected with diodes. Hereby, a certain electric potential is modified within the in-pixel microcircuit allowing the detection of the particle. Due to charge dispersion, an impinging particle usually activates a cluster of several neighboring pixels. Source: [68].

respectively. In order to improve the readout speed and minimize the reset noise pedestals, correlated double sampling (CDS, [54]) needs to be applied together with digitization and zero suppression [67]. However, the restriction to one CMOS process and the realization of diodes in n-wells forbids the use of complementary technology. Additional n-well PMOS transistors would compete with the n-well diodes for charge collection, hence only NMOS can be used in the pixel region. Presently, MIMOSA-26 performs the digitization of pixel hits outside the pixel area. The pixels are read out row-wise (rolling shutter operation) by connecting them column-parallel to digitizers (1-bit comparators) [67, 69]. The structure of the MAPS pixel region is outlined in Fig. 2.5.

The CDS, however, can still be performed on the pixel level via **clamping** [54]. The clamping process is described in detail in Appendix A. MIMOSA-26 uses a specialized NMOS capacitor (MOSCAP) to capacitively couple the pixel with the readout path. The MOSCAP gate serves hereby as one capacitor plate, while the source and the drain are shorted together as the second. The clamping method is sufficient to perform CDS and obtain reasonable levels of signal-to-noise ratios for accurate particle detection. A source follower is used at the pixel output stage to forward the clamping result to further readout stages, e.g. digitization. In fact, MIMOSA-26 applies a second CDS during the digitization stage which is discussed in the following chapter. After the readout, the MOSCAP voltage on the diode side needs to be restored to its original value in order to detect another hit. MIMOSA sensors currently implement two different strategies for that. A constant, low voltage can gradually recharge the MOSCAP over time (SB, self-biasing mode), or an additional in-pixel transistor can be used to reset the potential immediately after the readout (3T, three transistor approach). MIMOSA-26 uses SB pixels. The simplified schematics of the pixel architecture is given in Fig. 2.6. The benefits of the entire CDS procedure are the reduction of low frequency noise, reset noise, thermal noise, channel charge injection, clock feed-through and column-wise fixed-pattern noise sources [67, 72].

Recent advancements have demonstrated that non-ionizing radiation hardness can be increased by an order of magnitude using a high-resistivity epitaxial layer [71]. First chip to confirm

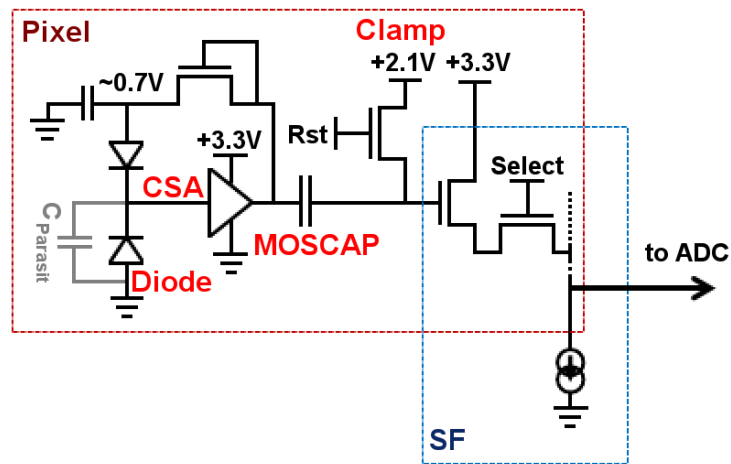


Figure 2.6: The simplified MIMOSA-26 in-pixel electronics. The charge collecting diode is located to the lower left (Diode). The charge is amplified with the Charge-Sensitive Amplifier (CSA) that loads an NMOS capacitance (MOSCAP). The MOSCAP is read out via clamping. At the pixel output, the Source Follower (SF) generates a low-impedance signal to the digitization stage (ADC). The second diode in the picture gradually restores the potential to its original value, before the hit (SB mode). Source: [71, 72].

this hypothesis is MIMOSA-25 [68, 70, 71] which features an epitaxial layer resistivity of $\sim 1000 \Omega\text{cm}$. After some additional studies (see [71] for details), a new version of MIMOSA-26 has been developed with a high-resistivity, $15 \mu\text{m}$ thick epitaxial layer. The chip is named MIMOSA-26AHR denoting the High Resistivity (HR) process. In conventional MIMOSA-26 sensors, the depleted area beneath the diode could not extend to more than a fraction of a micrometer [70]. With the new sensor, the depletion region depends on the applied diode voltage. But even for low values, the zone extends deep into the epitaxial region, according to simulations. Simulations of the depleted region in both sensor versions are presented in Fig. 2.7. Electrons which diffuse into the depleted area are exposed to an electric field and experience a force directed towards the diode. Charge collection is not a purely random process anymore and the charge collection efficiency is increased in the HR version. The HR sensors exhibit therefore improved signal-to-noise values [71]. This contributes greatly to non-ionizing radiation hardness, as shown in Fig. 2.8.

In conclusion, MIMOSA sensors feature excellent signal-to-noise ratios in the range from 15 up to 70. The detection efficiency for MIPs can reach over 99,5%. Spatial resolution down to $1 \mu\text{m}$ has been measured, depending on the pixel pitch [67]. They are a fully featured system-on-chip that can be operated at room temperature. Though, the radiation hardness improves when they are cooled down [58]. Furthermore, the sensors can be thinned down to reach the desired material budget goals. Currently, a thickness of $50 \mu\text{m}$ has been reached ($\approx 0.05\% X_0$). Therefore, the MAPS provide an attractive combination of pixel granularity, material budget, readout-speed, radiation tolerance, and power consumption [67].

2.2.3 Novel Technologies for MAPS Sensors

The MIMOSA-26AHR, as well as other MIMOSA versions have presently one disadvantage. Their in-pixel electronics are restricted to one MOSFET transistor type only (e.g. NMOS). Therefore, the level of logic complexity is delimited and the chips have not reached their full potential.

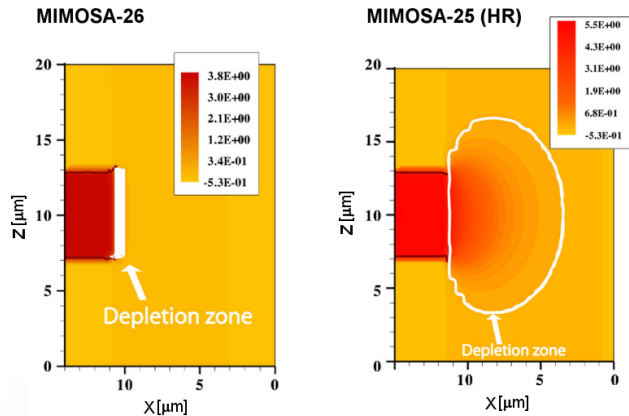


Figure 2.7: The simulation of the diode voltage distribution inside the MIMOSA-26 and MIMOSA-25 epitaxial layer. Due to the larger depletion zone, the HR version collects electrons much more efficiently. This effect counters the non-ionizing radiation damage in silicon. Source: [70].

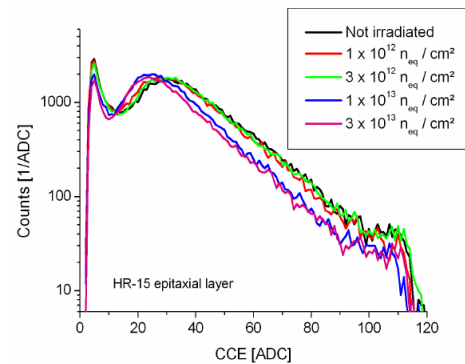


Figure 2.8: The MIMOSA-26AHR sensor are able to sustain large doses of $10^{13} \text{ n}_{\text{eq.}}/\text{cm}^2$. The charge collection efficiency (CCE) does not change significantly after irradiation. Source: [71].

This issue has been resolved in some alternative technologies, two of which will be presented in following.

One simple way to avoid this issue is found in the deep n-well MAPS technology [73, 74]. Here, a large n-well with the charge collecting diode is implanted deep into the silicon. Additional n-wells can be placed above. Due to sheer geometric aspects only a small, tolerable fraction of charge carriers reaches the upper n-wells.

In the so-called triple well technology the deep n-well contains another p-well for NMOS transistors. Fig. 2.9 shows the cross-section through the silicon wafer. All structures are embedded directly into the epitaxial layer. The deep n-well acts as a shield against substrate noise and allows for example full differential amplifiers together with the smaller n-wells in the vicinity. The limits are open to any CMOS microcircuit – ADC, shaper, latch, and many more. A large improvement of the signal-to-noise ratio, as well as the readout time is expected.

The second solution is more expensive and involves Silicon On Insulator (SOI) technology. SOI sensors, which are monolithic as well, circumvent the restriction to one transistor type by embedding a buried oxide layer (BOX) on top of a highly resistive substrate [54, 56]. In SOI, the entire substrate can be used as the active detection volume, the epitaxial layer is not needed anymore. Above the BOX, another wafer created in a low resistivity process is connected containing all the pixels and the readout electronics. There, PMOS and NMOS can be used without restriction and being built on an insulator, they can be spaced near each other without the negative parasitic effects. For that, every transistor is built on its own island of silicon eliminating latch-ups completely [54]. Contacts for diodes between the two silicon layers can be made with etching techniques and vias. While the SOI sensors certainly introduce many improvements they are more sensitive to ionizing radiation due to the buried oxide. They are also not part of a commercialized CMOS process and therefore expensive. Fig. 2.10 shows the cross-section of the SOI pixel.




Figure 2.9: The triple well MAPS uses a deep n-well to collect most of the generated charge carriers. Smaller n-wells near the surface are allowed for PMOS transistors. Source: [74].




Figure 2.10: The SOI technology uses a buried oxide layer (BOX) to insulate the CMOS part from the detection zone.

2.3 Field Programmable Gate Array

The last part of this chapter presents an interesting microelectronics device that forms the basis of nearly every detector readout system worldwide. In 1985, Xilinx company first introduced Field Programmable Gate Array (FPGA) microchips [75] comprised of logic cell arrays with customizable functionalities. Therefore, they can be reconfigured to fulfill various tasks. These multi-purpose microchips nowadays equip numerous embedded systems worldwide.

2.3.1 Theoretical Description

Digital electronic circuits can be well described using conventional logic [76]. First scientific approach to combine these two completely separate fields together was pioneered by Claude Shannon in his master thesis, submitted in 1937. The mathematical foundation of his theory originates from the work of George Boole in the 19th century, today known as **Boolean algebra**. Presently, Boolean algebra is widely used to describe digital electronics at the lowest levels.

Each Boolean variable can have only one of two values: true or false (1 or 0, respectively). Variables are called literals. They are combined together with a set of rules that is equivalent to a field of logic called Propositional Calculus. Operators combine two or more literals to produce an output, which is again true or false. For example, the AND operator combines two literals A and B corresponding to the logical thought: “If, and only if A and B are both true, the result is true.” An example of the NOR logic function and its realization in CMOS technology has been shown previously in section 2.1.1. By cascading several of these functions together, complicated logical networks can be created that represent some logical process. Such network is called a Boolean circuit and the operators are termed Boolean functions. A two-input Boolean function is defined as:

$$f_{Bool} : \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$$

There are 16 boolean functions (gates), in total, which operate on two literals. An inverter (NOT) is the only gate that acts on one literal, inverting its value. The Boolean functions are essential to determine the behavior of electronic devices. Every such function has an equivalent truth table.

From the computer-scientific perspective, it can be shown that the Boolean circuit is equivalent to a Turing machine [76]. A Turing machine is a theoretical automaton used in computer

science to describe algorithms. It is used only to determine the algorithm complexity class, not to calculate the actual result. The number of steps required to obtain the solution needs to be finite, hence the group of calculable algorithms is termed 'recursively enumerable'. But since Boolean networks resemble algorithms, and digital electronics devices are expressed through Boolean networks, it can be inferred that digital electronics devices can solve any deterministic and finite algorithm¹.

2.3.2 FPGA Introduction

The FPGA logic cells can implement Boolean functions. Their content, as well as the connections between them can be reprogrammed to fit the current task. This allows the implementation of a complex Boolean network in order to solve the given algorithm. In addition to that, the Input/Output cells (I/Os) can be adapted to support a large number of signal standards, e.g. LVDS and TTL. The FPGAs are thus generic, reconfigurable microchips suitable for a wide area of applications. They are a substantial part of every modern data acquisition and control system. Once soldered to a PCB² and connected to the peripheral devices, they can realize a variety of ASIC³ features without the need of a time consuming and expensive CMOS manufacture process. Merely by changing the logic content, a completely new design can be synthesized within minutes.

A simplified example⁴, showing the basic operation principle is presented in Fig. 2.11. The figure shows a full adder [52] implementing the bitwise addition. Two 1-bit inputs A and B are added with a carry signal C_{in} from the previous stage. The result is the sum S and the carry signal to the next stage. By cascading this basic component, N-bit adders can be easily created, as well as multipliers.

2.3.3 FPGA Architecture

A typical, modern layout, e.g. as found in the Lattice ECP3 FPGA [77], is outlined in Fig. 2.12. The logic cells, in terms of Lattice FPGAs called programmable function units, can fulfill many different functions. They can realize Boolean functions, distributed RAMs⁵ and shift registers, as explained in the following section. The I/Os (Lattice SysIO, Xilinx SelectIO) are also universal in the sense that they support a large variety of voltage standards (TTL, LVTTL, LVCMOS, LVDS, etc.). Every input cell can be reprogrammed into an output cell and vice versa. They can optionally incorporate input and output buffers. Routing inside the chip is programmable as well. For example, Xilinx FPGAs use General Routing Matrices (GRMs). They contain an array of switches that interconnects all the neighboring logic cells. It is possible to use long and short lines to interconnect the individual logic cells with each other. This allows for creating a bus throughout the entire chip by using long lines between cells which are far apart.

¹Moreover, a well-known hypothesis in computer science, the Church-Turing thesis, postulates that every calculable algorithm can be solved by a Turing machine and thus by a microelectronics device. However in some cases the algorithmic complexity would require too many resources or too much computation time for the device to be of practical use.

²PCB – Printed Circuit Board

³ASIC – Application Specific Integrated Circuit

⁴This example is designed for educational purpose only. The real FPGA would simplify the logic and pack it in few logic cells.

⁵RAM – Random Access Memory




Figure 2.11: The demonstration of the principles of an FPGA on the example of a full adder. The full adder truth table and the corresponding Boolean function is shown on the right-hand side. All visible FPGA components are programmable: the I/Os, the logic cells and the routing lines. By reprogramming the cells, a completely different component can be created. Source: [75].




Figure 2.12: A typical FPGA architecture.

The clock distribution inside the FPGA is also performed over several different clock drivers. The global clock is usually transmitted over long, low-impedance lines at the outer regions of the chip from where it can reach the inner cells over local (regional) clock drivers. Modern FPGAs use PLLs¹ and clock managers to skew, deskew, and modify the frequency and the phase of the input clock.

In order to provide sufficient on-chip memory, the FPGA contains multi-purpose block-RAMs. The blocks of memory can be combined together to form larger blocks. They can be used as single- or dual-port FIFO² buffers, ROMs³ and RAMs. The block-RAM functionality can be programmed using IP⁴ cores. These are vendor-specific, logical units which enable specific functionalities inside the FPGA. There are two types of IP cores: soft- and hard-cores. The soft IP cores can be implemented by reprogramming the logic cells or block-RAMs, whereas the hard

¹PLL – Phase Locked Loop

²FIFO – First In First Out

³ROM – Read-Only Memory

⁴IP – Intellectual Property

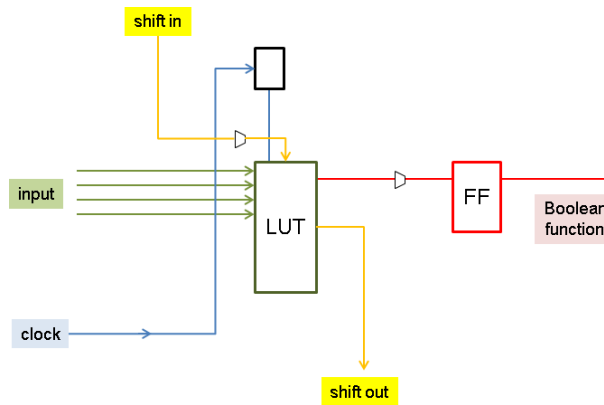


Figure 2.13: A modern FPGA slice can serve as a shift register, a RAM or a Boolean function. The central element, the look-up table (LUT), can be reprogrammed for different purposes.

IP cores require certain hardware components, e.g. an embedded microcontroller.

The FPGA cells are typically programmed using SRAMs. Therefore, the FPGA loses its functionality after shut-down and needs to be reprogrammed every time it is powered on. This can be done in master, slave or peripheral mode. The slave mode allows an FPGA daisy chain and the peripheral mode uses a microcontroller to load the bit-file containing the firmware. There are anti-fuse FPGAs as well, which do not lose their program. However, once programmed, they can not be reprogrammed anymore. The more expensive Flash FPGAs do not lose their code after power-off and can be reprogrammed a limited number of times.

FPGA Logic Cell

The logic cells are used primarily for the implementation of Boolean functions. Lattice ECP3 and Xilinx Virtex4 FPGAs both use 4:1 look-up tables (LUTs) for realization of their logic¹. The Boolean function is realized as a truth table. A LUT is in principle a 16 bit memory block where the four input signals are used as a decoder to obtain data from the memory. Each of the 16 possible '0' and '1' combinations accesses one bit of the memory. A flip-flop (FF) [52] is used in the output stage to more complex, synchronous logic networks. If not needed, the FF can be bypassed with a multiplexer.

Commonly, logic cells are divided into slices, the smallest programmable units in an FPGA. Each Virtex4 and ECP3 logic cell contains four slices, each with two LUTs. Besides the realization of Boolean functions, the slice can be used as a RAM or a shift register, based on the firmware. A typical structure of a slice is outlined in Fig. 2.13.

2.3.4 Radiation Effects

FPGAs are not as sensitive to radiation damage as MAPS, however they also experience a variety of different radiation-induced defects which can, in worst case, lead to the destruction of the microchip. Particularly, the Total Ionizing Dose (TID) which microchips are, in general, exposed to is responsible for steady degradation of voltage and switching characteristics of the applied transistors [78]. However, latch-up effects may also get triggered by radiation in FPGAs and

¹The 4:1 LUTs have four inputs and one output.

ASICs. The less dangerous, but by far more frequent Single-Event Upsets (SEU) and Single-Event Functional Interrupts (SEFI) can produce a temporary malfunction in the FPGA. Hereby, the impinging particle can cause a bit-flip in an SRAM cell of the FPGA fabric. Based on the cell location, the effect can be more or less serious. SEFI errors are related to the on-chip configuration and control registers. By changing a bit in these crucial registers, the entire FPGA can get modified, reset, or disabled. SEU, on the other hand, occur orders of magnitude more often [78] and corrupt the FPGA firmware by changing the configuration of routing matrices and multiplexers, or modifying LUT and block-RAM contents. More details on radiation effects can be found in [57, 78].

In order to overcome these risks, two methods are widely used that can eliminate SEU, increase fault-tolerance and thus contribute to a safer FPGA operation, as required in space-applications, nuclear facilities and heavy-ion experiments. Triple Modular Redundancy (TMR) creates a threefold copy of every logic element and signal line, including the clock. Therefore, if one of them is failing due to an SEU, the other two can proceed undisturbed. So-called voters are connected to the outputs of all three stages in order to determine which one is failing. They are realized in LUTs and need to be tripled as well. TMR requires a huge amount of additional resources and does not protect the device fully from all failures. Only SEU related to logic, routing and memory can be corrected [78]. Triplicating the entire FPGA chip is required to ensure SEFI tolerance. In addition, all FPGA cells can be repeatedly reinitialized during runtime (scrubbing) to guarantee that the correct firmware is present. However, this option needs to be supported by the FPGA model. More details can be found in [79].

Chapter 3

MIMOSA-26 Sensor Specifications

The MIMOSA-26 sensor unites a decade of MAPS research and development into one sophisticated microchip¹. It was considered the state-of-the-art for several years, fulfilling many design parameters required for the final CBM-MVD. Even though newer MIMOSA generations exist, their architecture is very similar to MIMOSA-26. Numerous MVD design parameters can be studied with this sensor, e.g. gluing, handling and bonding of MAPS sensors, thermal properties, cluster- and data- processing algorithms, and the MVD readout concept. These are only some of the reasons why the sensor is chosen to equip the first MVD prototype. In order to understand the state-of-the-art sensor and to provide an adequate insight into the present MAPS technology, this chapter will describe its digital architecture and discuss its many properties.


As first, section 3.1 presents a general description with a short summary of the main sensor properties. Section 3.2 presents afterwards the internal architecture. Hereby, the focus is set on the high-speed digital part. The data flow throughout the sensor, as well as digital components are discussed in detail. After that, section 3.3 describes the sensor operation, main characteristics and sensor limitations. Lastly, some recent improvements of the MIMOSA family are presented in section 3.4.

3.1 General Description

MIMOSA-26 is produced with the Austria Microsystems AMS 0.35 Opto (HiRes) process [80] and features four metallization and two polycrystalline silicon layers. The thickness of the epitaxial layer is $\sim 15 \mu\text{m}$ and the pixel pitch is $18.4 \mu\text{m}$. The sensor was produced in 2009 and its version based on the high-resistivity epitaxial layer in 2012. First application of MIMOSA-26 for tracking has been performed in the scope of the European Detector JRA1² beam telescope for the International Linear Collider (ILC) prototyping [69]. Afterwards, MIMOSA-28 was developed for the STAR detector and differs from MIMOSA-26 basically in geometry only. The internal architecture of both sensors is essentially the same.

¹The difference between MIMOSA-26 and MIMOSA-26AHR is merely in the high-resistivity (HR) epitaxial layer. The architecture of both chips is the same. Therefore, if not otherwise mentioned MIMOSA-26 will denote both chips in the following.

²JRA – Joint Research Activity



(a) Picture of not thinned MIMOSA-26 sensor glued and wire-bonded to a printed circuit board. Source: [82].

(b) The pixel matrix is read out in the rolling shutter mode where the readout time of each row is time-shifted against each other. The duration of each frame is equivalent to the integration time.

Figure 3.1: MIMOSA-26 geometry and readout characteristics.

MIMOSA-26 is a combination of two precursor test circuits: MIMOSA-22 and SUZE-1. The pixels are identical to one test-array of MIMOSA-22, whereas the intrinsic encoding of the particle hits is performed with SUZE-1 zero suppression circuitry. The total sensor size is approx. 2.95 cm^2 with an active area covering $\approx 2 \text{ cm}^2$, as shown in Fig. 3.1a. The sensor can be thinned down to $50 \text{ }\mu\text{m}$ without losing its detection capabilities. By design, it is divided into two main segments: the sensing area in form of a pixel matrix, and the readout circuitry. The total power consumption amounts to $\sim 730 \text{ mW}$, one third of which is distributed among the pixels, the rest within the readout logic block. The power density of $250 \text{ mW}/\text{cm}^2$ is therefore not equally distributed over the sensor. The sensor contains hot-spots especially in the lower part, below the matrix [81].

The pixel matrix contains $\sim 660 \text{ k}$ pixels distributed within 576 rows and 1152 columns. One matrix-row is processed within 200 ns in the so-called rolling shutter mode. All pixels in the row are processed in parallel (column-parallel readout). The integration time t_{int} is the time between two consecutive readout steps that the pixel has at its disposal in order to integrate the charge from the charge-collecting diode and detect a particle hit. Ideally, t_{int} is defined on the pixel-level, but since all pixels in one row are read out simultaneously it can be defined on the row-level as well. An example is shown in Fig. 3.1b. A full readout cycle is termed 'frame' in analogy to image processing. The integration time directly affects the frame rate ($1/t_{\text{int}}$) at which the sensor operates. The MIMOSA-26 chip operates at an internal frequency of 80 MHz and $115.2 \text{ }\mu\text{s}$ integration time. This leads to a frame rate of 8.7 kHz. The sensor data rate reaches up to 20 MB/s, however due to intrinsic limitations only a certain number of hits can be processed. The on-chip readout circuits are responsible for digitization, zero suppression, data formatting, and serial transmission, as explained in the following section.

3.2 Digital Sensor Architecture and Operating Principles

As shown in Fig. 3.2, the MIMOSA-26 chip can be divided into following functional units: JTAG¹ interface, sequencer, pixel matrix, comparators, Sparse Data Scan (SDS), multiplexer,

¹JTAG – Joint Test Action Group

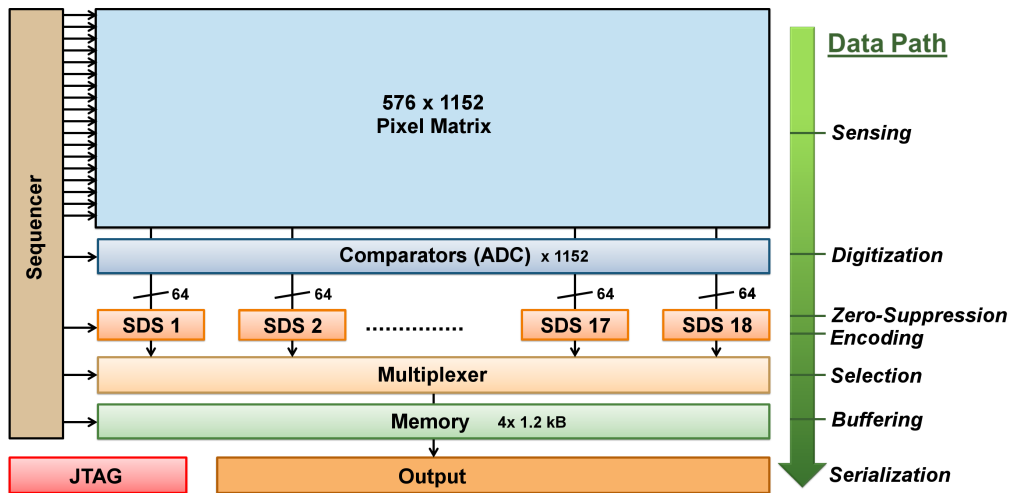


Figure 3.2: The internal MIMOSA-26 sensor architecture with the data flow. The size of the component blocks does not correspond to the actual size. All components except the pixel matrix are situated within the 1 cm^2 large readout area below the matrix.

memory, and serializer. Apart from that, the sensor also contains some analog readout circuits which are not discussed here. All the necessary sensor parameters are set via the JTAG interface. Then, the sequencer steers the readout steps required for the main sensor operation. The internal pixel architecture, as discussed in section 2.2.2, plays the main part with respect to particle detection.

In principle, each pixel provides an output signal in form of a voltage level. A particle traversing through the pixel epitaxial layer will modify the output voltage to some extent. The pixel voltages are then compared to some user-specified threshold value by the array of comparators. If the voltage level exceeds the threshold, a logical '1' is generated and a '0' otherwise. Thus, the comparators digitize pixel hits with respect to a predefined threshold setting. They operate independently, in parallel to form the first step of a three-stage pipeline to pre-process the data before transmission. The second pipeline stage uses the SDS combinatorial circuits to reduce the data load. Most of the pixels usually generate a '0' after the digitization and can be discarded from the data stream. This procedure is termed zero suppression. The SDS circuits encode the remaining digital '1'-s from the pixel matrix into 16-bit data packages. The last pipeline stage selects a certain number of packages using the multiplexer and stores them to the memory. Now the data is ready for transmission. Once the entire matrix is processed, the memory is read out and processed by serializers generating a bit-stream on two output channels.

The sensor provides its data continuously. Once started, it will operate in an infinite loop. The frame period is fixed to $115.2 \mu\text{s}$. However, the amount of useful data depends on the number of registered particle hits after the digitization process. Thus, the data volume varies with respect to the sensor occupancy and the applied threshold. The data is streamed with an internal 16-bit format. Additionally, the chip contains several test modes to analyze the pixel matrix, find defective comparators, check the SDS logic and test some voltage levels. In the following, the individual sensor components according to Fig. 3.2 are discussed.




Figure 3.3: The JTAG chain with N sensors.

3.2.1 JTAG Bus Interface

The MIMOSA-26 control interface implements the JTAG Boundary Scan standard, IEEE 1149.1 Rev1999. It supports 23 instructions which access 17 internal registers to configure the sensor [83]. Since JTAG is a serial bus, the sensors are daisy-chained and programmed together, as depicted in Fig. 3.3. Four signals are applied: Test Data Output (TDO), Test Clock (TCK), Test Mode Select (TMS) and Test Data Input (TDI). To form the chain, the output TDO of one sensor is connected to the input TDI of its neighbor. TCK transmits the JTAG clock which is usually in the kHz regime, but it can also operate at 10 MHz [84]. Another important JTAG signal is the TMS which sets the mode of the JTAG state machines. Detailed description of the JTAG procedure can be found in a related work [84]. All the registers are described in [83].

The applied JTAG interface has one major limitation: there is no direct way to access the registers of a particular sensor and verify whether the values written inside are correct. The register value can be read out only by writing a new value into the register. This might be taken into account regarding error-checking. There is no guarantee that the written value is correct. By checking the written value, another value must be written which is then, again, left unchecked.

3.2.2 Sequencer

The sequencer is the central control unit of the sensor. It synchronizes all readout steps from different internal components and provides the necessary clock and control signals, as shown in Fig. 3.4. The three stages of the pipeline (digitization, zero suppression and buffering) are all synchronized using the *cklatch* signal. The *cklatch* signal simply denotes when the data of one stage is ready to be processed by the next stage. It is provided to all three parts simultaneously. The pipeline can be understood as follows: while one matrix row is being digitized by the comparators, the SDS encodes the previous row at the same time, and the multiplexers store data from the row before that.

To explain all the individual parts of the sequencer in detail would go beyond the scope of this thesis. Some detailed information can be obtained from [83, 85, 86]. As an example, the rolling shutter operation will be explained in the following.

In order to reduce the power consumption, the pixels are deactivated by default. The *Power_On* signal which activates pixel-based amplifiers in a matrix row is deactivated. Merely two out of 576 rows are then activated at the same time. One row is selected for readout with the *Select* sig-

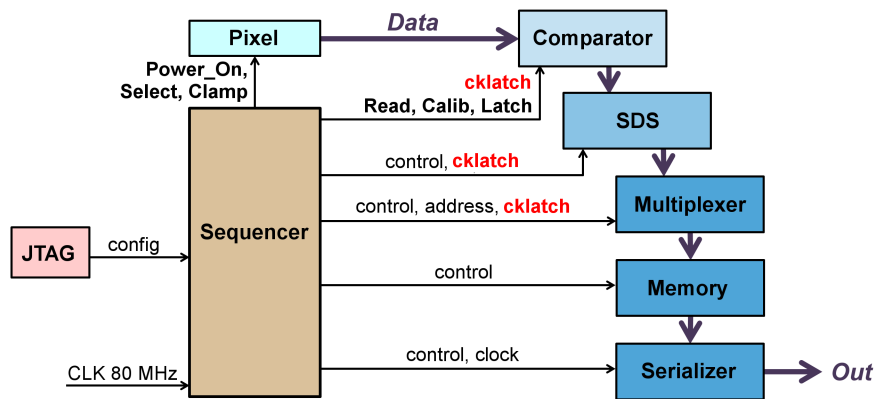


Figure 3.4: The sequencer is the core part of MIMOSA-26. It controls the rolling shutter operation, clamping, digitization and all the other readout stages. The component can be configured via several JTAG registers to optimize its performance.




Figure 3.5: The timing diagram of the rolling shutter and the digitization process. The row 575 is processed first, followed by 574. The pixels are read out in parallel by first applying the *Read* signal, then the *Clamp*, the *Calib* and the *Latch* signals in a sequence to all the in-pixel electronics within the same row. It can be seen that while the row 575 is read out, the row 574 is getting powered on to stabilize the voltages prior to the readout. Source: [83].

nal while the other is prepared to be read out. This is necessary because some preparation time is required to start up the amplifiers. Afterwards, the readout of the pixel voltage via clamping and the entire digitization process is performed within 200 ns. Immediately after that, the next row can be processed and the row after that is powered on to prepare for the readout. The cycle goes through all the rows and starts from the beginning in an infinite loop. The timing diagram of the procedure is given in Fig. 3.5. The *Read* signal initiates the readout of the pixel voltage. After that, the *Rst* signal resets the pixel source follower operating point [85] and *Clamp* performs the clamping step. As mentioned before, a particle hit will modify the readout voltage. The clamping voltage restores the readout potential to some practical value in order to detect another particle during the subsequent readout cycle. The signals *Latch* and *Calib* are used by the comparators only and will be explained in the following section.




Figure 3.6: The comparator structure is displayed together with the pixel. The variables i and j denote the row and the column, respectively. Each column of 576 pixels is multiplexed to one comparator. Therefore, the path between the pixel and comparator, as well as the corresponding fixed-pattern noise depend on the row i . The output is registered by a D-Flipflop. Source: [86].

3.2.3 Comparators

Located directly below the pixel matrix is an array of 1152 comparators [72], one for each column. They amplify and digitize the pixel output voltages with respect to an adjustable threshold value, which is set via JTAG. The width of each component is $18.4 \mu\text{m}$, as they have to match the pixel width. The power consumption amounts to $\approx 230 \mu\text{W}$ per comparator unit, or approx. 265 mW for the whole array¹. The comparators are coupled to each other via their bias voltage and some switches. In order to reduce the noise resulting from this coupling, they are divided into four segments [67]. Each segment contains 288 devices. The segments are referred to as blocks A, B, C and D. In literature, the term discriminator is often used to describe this component.

The schematics of the comparator together with the pixel circuitry is shown in Fig. 3.6. Every amplifier stage has a gain of 4. The pixel signal is amplified by a factor > 200 before it reaches the latch. The amplifiers provide a clearer signal and the latch performs the discrimination. More details on the individual components and the discrimination process can be found in [72].

In order to understand the component, signal timings from Fig. 3.5 need to be considered. Furthermore, Fig. 3.7 shows an example of signal progression at the MOSCAP capacitor, which couples the pixel charge with the readout path. V_{MOSCAP} and V_{Clamp} correspond to the voltages which are present at the readout side of the MOSCAP capacitor. In contrast, the V_{Hit} and $\Delta V_{\text{Recharge}}$ are voltages related to charges induced at the opposite side (diode side). They influence V_{MOSCAP} via capacitive coupling, which is why V_{Hit} has the opposite sign, by definition. The particle signal V_{Hit} is in the range between 0.1 mV and few mV after the pixelwise amplification. This demands very precise discrimination. Moreover, the discrimination time needs to be very short. It is in the order of the pixel readout time, which is 200 ns . The best possible solution has therefore been found in a fully-differential device with cascaded low-gain preamplifiers in deep inversion [72].

Fig. 3.7 additionally shows the in-pixel CDS results. The applied CDS procedure (see Appendix A) basically subtracts the MOSCAP voltages of two consecutive sampling points. The time between the sampling points is equal to the sensor integration time, i.e. $115.2 \mu\text{s}$. The

¹This is one third of the entire MIMOSA-26 power dissipation.

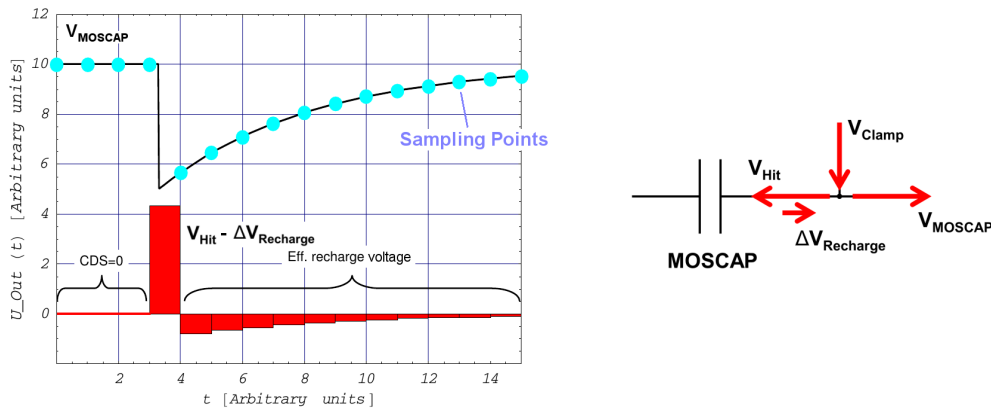


Figure 3.7: An example of the signal voltage progression at the readout side of the MOSCAP (V_{MOSCAP}). The voltage reaches its peak if no hit occurs and the readout potential is fully recharged by the clamping voltage. After the particle hit at $t \approx 3$, the voltage drops and the recharge current gradually pulls the voltage back to its nominal value. Below the voltage slope, the CDS result is displayed. All of the mentioned contributing voltages are summed up at the right-hand side. Source: [47].

recharge voltage and V_{Hit} have different polarity. Moreover, the recharge voltage gets gradually smaller with each cycle, if no particle hit occurs in between. Thus, the sign of the CDS result can clearly differentiate whether a hit occurs or not, in this idealized picture.

All 576 pixels in one column are multiplexed to one comparator. Therefore, there are Fixed Pattern Noise (FPN) effects which need to be considered. Pixels near the top of the matrix have a longer signal path than the pixels at the bottom. Consequently, the parasitic effects are higher and the voltage gets modified along the way disturbing the signal. The FPN alone can be in the order of the signal generated by the impinging particle and needs to be eliminated. This is accomplished by applying the CDS technique twice. First CDS attempt is performed within the pixel to obtain a readout signal and remove any pixel-based non-uniformities. The MOSCAP potential on the readout side will get repeatedly reset by the clamping voltage in order to obtain signals purely caused by impinging particles and without some pixel noise effects, or preceding readout cycles. Due to the SB pixel, the MOSCAP potential on the diode side will be gradually restored over time after each particle hit. Afterwards, the second CDS is used to eliminate the FPN. Following phases are required to perform the entire procedure:

- **Phase 1:** The pixel voltage level is read out by applying the *Read* signal. The result is applied to the capacitor pair $C1 - C1'$ within the comparator. The other side of the capacitor pair is kept at some fixed potential (V_{r2} from Fig. 3.6). V_{MOSCAP} can take a value from three different ranges. The base voltage is expressed in terms of V_{Clamp} , however, additionally this voltage gets modified with respect to a particle hit (V_{Hit}) and the subsequent, moderately slow recharge current ($\Delta V_{Recharge}$). Thus,

$$V_{MOSCAP} = V_{Clamp} - V_{Hit} + \Delta V_{Recharge}$$

as shown in Fig. 3.7. But the pixel voltage level is not stored alone. Also the FPN noise of the signal path V_{FPN} is memorized by the comparator. At the same time, the comparator uses the first differential amplifier to integrate the threshold value V_{Ref1} into the

equation. V_{Ref1} denotes the discrimination threshold voltage with some internal pedestal (V_{Ref2}) which can be set by the user via JTAG ($V_{\text{Ref1}} = V_{\text{Ref2}} + V_{\text{Th}}$). As a result:

$$V_{\text{phase1}} = V_{\text{Ref1}} - (V_{\text{Clamp}} - V_{\text{Hit}} + \Delta V_{\text{Recharge}} + V_{\text{FPN}})$$

After the readout, the switches S3 and S3' are opened in order to leave the potential at the right-hand side of the C1 - C1' pair floating, sensitive to any changes from the following readout phases.

- **Phase 2:** During this phase, the comparator is deactivated. Instead, the pixel restores the MOSCAP potential at the readout side by applying V_{Clamp} and resetting the in-pixel source-follower operation point. All information regarding the hit is removed from the MOSCAP readout side. Note that if the pixel is still recharging after a hit, this potential will not get fully restored, leading to:

$$V_{\text{phase2}} = V_{\text{Clamp}} - \Delta V_{\text{Recharge}}$$

Note that after this step, the *Clamp* switch will be opened again in order to leave the MOSCAP potential on the readout side floating and sensitive to any further hits which can then be registered during the next readout cycle in phase 1.

- **Phase 3:** Now, the second CDS can be performed without further ado. The comparator initiates the *Calib* signal in order to obtain the restored MOSCAP voltage, again, with the FPN noise source added to it. Eventually, some rest recharge voltage may be present due to a preceding particle hit. Additionally, the internal pedestal of the threshold voltage V_{Ref2} is provided to the differential amplifier in this stage:

$$V_{\text{Calib}} = V_{\text{Ref2}} - (V_{\text{Clamp}} - \Delta V_{\text{Recharge}} + V_{\text{FPN}})$$

But in this step, due to the opened switches S3 and S3' only the difference between the previously stored voltage and the currently read-out voltage is stored in the C1 - C1' pair. Therefore,

$$\begin{aligned} V_{\text{phase3}} &= V_{\text{phase1}} - V_{\text{Calib}} \\ V_{\text{phase3}} &= V_{\text{Ref1}} - V_{\text{Ref2}} - V_{\text{Hit}} + \Delta V_{\text{Recharge}} \end{aligned}$$

Due to the differential amplifier, the reference voltage V_{r2} is also eliminated. The original particle-induced voltage V_{Hit} in the range of few mV can be accessed without the FPN source and is compared directly with the pure threshold $V_{\text{Th}} = V_{\text{Ref1}} - V_{\text{Ref2}}$ without the pedestal V_{Ref2} . The signal is amplified and provided as the input to the latch. The recharge voltage $\Delta V_{\text{Recharge}}$ can never be mistakenly taken as the particle hit due to the opposite sign. The latch is simply blind against differences resulting solely from the recharge voltage due to the polarity of the CDS result (see Fig. 3.7 for details).

- **Phase 4:** Lastly, the latch performs the discrimination. If the input $V_{\text{Th}} - V_{\text{Hit}}$ is negative¹, then the latch will produce a logical '1' at the output denoting a particle hit, and '0' otherwise. Note that the voltage drop can also be caused by noise ('fake hit') and that the user-specified level of the threshold is vital for the detection efficiency.

¹The recharge voltage $\Delta V_{\text{Recharge}}$ is assumed to be negligibly small with respect to V_{Hit} in this step.




Figure 3.8: The zero suppression encoding algorithm (SUZE-1). A 'state' is a row-wise fraction of the cluster created by an impinging particle. Row i is selected for readout in this example. It contains two 'states' of four and three pixels, respectively. They are encoded in 16-bit by storing the address of the first pixel and denoting how many subsequent pixels are activated (binary code). Source: [87].

3.2.4 Sparse Data Scan

The previous step creates a digital representation of a particle hit and stores it in a D-type flip-flop. The task of the SDS circuitry is to encode these results by choosing a different representation of '0' and '1' in order to reduce the data load. The logical '0'-s which form the majority of comparator outputs will not be used any more (zero suppression). The logical '1'-s however, representing true particle hits, can later be dereferenced as (x,y) pairs where x and y denote the row and the column addresses in the ranges 0-575 and 0-1151, respectively. Due to a charge spread in the sensing volume, it is more likely that a particle activates a whole area of neighboring pixels which is called a cluster. Fig. 3.8 shows an example.

MIMOSA-26 encodes the data row-wise due to the rolling shutter operation. All logical '1'-s in one row resulting from the digitization stage are processed by a purely combinatorial circuit (see [85, 87] for details). A string of neighboring '1'-s is denoted as 'state'. In Fig. 3.8 for example, the row i is selected for zero suppression which contains two 'states'. Each 'state' is encoded as a 16-bit word containing the column address of the first pixel and the number of subsequent '1'-s. MIMOSA-26 only supports 'states' with up to 4 pixels.

The 1152 possible hits are divided into 18 banks which work in parallel. Each bank is responsible for 64 pixels and can only produce a maximum of six 'states'. The combinatorial circuit finds the proper hits and sets the encoding. First, the row address is encoded together with the number of found 'states' and an overflow bit¹ in the memory (also as a 16 bit word). Then, a multiplexer selects nine 'states' from all 18 banks and stores them in the memory (pipeline stage 3). Only a total of 1026 of such 1-4 pixel strings can be encoded during one frame.

¹The overflow bit marks if more than nine 'states' are present in the row, in which case they are discarded.

3.2.5 Memory and Output

MIMOSA-26 features four SRAM cells, each with 600×16 bit of storage space. Two of them are used for storing the data from the current frame. At the same time, the other two cells containing the data from the previous frame are read out. At the end of the readout cycle the sensor switches the pairs and proceeds with the next frame. The number of stored packages per frame changes dynamically with hit occupancy. However, it is memorized by an internal counter and provided to the subsequent readout stage.

The sensor uses two output ports. Each port is associated with one of the SRAM cells for readout. Two internal serializers obtain 16-bit packages from the memories and load them into their local 16-bit register, one at a time. The register content is then simply shifted out with the internal 80 MHz clock¹, contributing to a total bandwidth of 160 Mbit/s. Further details can be found in [83, 86, 87]. However, prior to reading out the RAM, some header information is provided, e.g. the frame length. The final data format is presented in the subsequent section.

3.3 Sensor Operation and Characteristics

The full bonding scheme can be obtained from [83], however, it is sufficient to use a reduced bonding scheme (using only the digital sensor interface) presented in Appendix B. The power dissipation is regulated internally after providing two 3.3 V supply voltage inputs for the digital and analog parts, separately. Additionally, the voltage of 2.1 V is required for clamping. This voltage directly affects the pixel and the comparator performance and needs to be stable during the entire operation.

A JTAG driver is necessary to program and operate the sensor, as explained in [84]. Besides the JTAG relevant signals, the sensor requires a differential clock input. The optimal performance is reached with a frequency of 80 MHz. Additional single ended lines can be used to start and reset the sensor, however it is also possible to start the sensor via JTAG. The sensor data is output serially via two LVDS channels D0 and D1. Additionally, the clock and a digital marker signal (MKD) are provided as well. The clock is used to deserialize the data and MKD marks the start of a new readout cycle.

Internally, all the adjustable voltage and current levels are driven from one single JTAG register (BIAS_DAC). The register also defines the individual comparator thresholds. However, they can be modified only for groups of 288 comparators, as a whole, and all of them apply the same V_{Ref2} setting (see section 3.2.3). Thus, VDISREF2 provides the internal pedestal and should be set to 1 V. After that, the individual thresholds for the four comparator blocks are set using VDISREF1A - VDISREF1D. The internal thresholds can be monitored by examining sensor output pads 3 - 17 (see table B.1). In addition to that, a voltage on pad 1 is related to an integrated temperature sensor.

Due to the constant integration time, the sensor provides frames at a rate of 8.681 kHz. If no hits were detected, only a small overhead of 128 bits per frame is given contributing to a data rate of 0.14 MB/s. However, under full load 19.93 MB/s can be reached. The intrinsic data format, as shown in Fig. 3.9, simplifies the data readout. The zero-suppressed data is packed into

¹For compatibility reasons the clock can be halved. Additionally, one output channel can be deactivated.




Figure 3.9: The intrinsic MIMOSA-26 data format. All packages are transmitted bitwise. In this example, 20 packages are provided, 8 of which form the frame overhead of 128 bits. The actual hits are encoded within the row- and column-packages. As can be seen, 9 'states' from 3 rows are given here. The data length amounts to 6 per channel.

16-bit words. But prior to the SRAM readout, the sensor first provides some general information initiated with a Header package. Then, the current frame number, as stored by an internal 32-bit counter, is provided on both output channels. The D0 channel contains the LSB¹, whereas the D1 contains the MSB² of the counter. Afterwards, the number of data packages for the given frame is provided on each output channel. The value is always the same on both channels. This is true, even if D1 does not contain a useful package at the end. In this case, the last package must be discarded. The dummy package can be identified due to a special encoding inside the data stream.

The SRAM data is formed by two distinctive types of packages. At first, the row address (0 - 575) is given with the number of useful 'states' (see section 3.2.4) and an overflow bit which marks whether the maximum number of possible 'states' per row has been exceeded. Afterwards, up to a maximum of nine 'states' with their column addresses are provided alternating on both output channels. The process is repeated for all following rows. Lastly, a Trailer package is transmitted. The Header and the Trailer packages are user specified unique bit patterns denoting the start and the end of a frame.

As a result of the encoding process and the output bandwidth, the sensor is subject to some **limitations**. The number of possible 'states' selected in one SDS bank is limited (see section 3.2.4). If there are more than 6 'states' in a bank, which is improbable yet not impossible, the data would be lost. Furthermore, only 9 'states' per row can be acquired during the SDS procedure, however the overflow bit can be used to detect the data loss in this case. Lastly, the stored data packages need to be transmitted during the period of one frame. Hence, only a maximum of 1140 data packages is encoded per frame, i.e. 570 words per SRAM cell.

The MIMOSA-26 sensor, as well as its high-resistivity version have been sufficiently studied [71, 88]. Their main characteristics were obtained regarding the detection efficiency, spatial resolution, noise levels and radiation hardness. Some of the results are presented in Fig. 3.10. The image shows the relationship between the noise and the detection efficiency. The Fake Hit Rate (FHR) is defined as the number of pixels which are activated by noise divided by the total number of pixels, obtained during one frame in the absence of any particle. According to the

¹LSB – Least Significant Bit(s)

²MSB – Most Significant Bit(s)




Figure 3.10: Preliminary MIMOSA-26 characteristics (left) compared to MIMOSA-26AHR test results (right). The HR sensor shows a clear improvement in all three measured categories. The detection efficiency reaches nearly 100 % at a FHR of 10^{-5} which was not possible with the standard sensor. Also the charge carriers diffuse to the neighboring pixels less often, which slightly improves the spatial resolution. Source: [71, 88].

studies, a reasonable detection efficiency can be achieved by tuning the sensor thresholds to produce a FHR of 10^{-5} . In that case, the spatial resolution also obtains an optimal value which is within the requirements for the CBM. Furthermore, the comparator transfer functions have been studied in a related work [89], which demonstrates their linear behavior. It can also be shown that the CDS procedure of the comparators successfully eliminates column-based FPN effects, but they are still present at the row-level. Moreover, radiation studies [71] have shown that the MIMOSA-26AHR is able to sustain radiation doses of several 10^{13} n_{eq.}/cm², which is very close to the final CBM goal.

3.4 Next-Generation MIMOSA Sensors

After demonstrating its suitability for particle detection, the MIMOSA-26 sensor has been re-designed to meet the STAR detector requirements driven by the expected hit multiplicity of $2.4 \cdot 10^5$ hits/cm²/s. The new sensor generation is termed MIMOSA-28 [86] and contains a larger pixel matrix with an increased pixel pitch. The internal clock frequency remains at 80 MHz, however the output bitrate is doubled to 160 Mbit/s per channel by using DDR¹ output registers. The RAM size has been increased in order to accommodate for the larger pixel matrix, leading to a total data bandwidth of 40 MB/s. The geometries of some recent MIMOSA generations are presented in Fig. 3.11, and their main characteristics are summarized in table 3.1. Both sensors, MIMOSA-26 and MIMOSA-28, apply the SUZE-1 zero-suppression circuitry to encode particle hits.

Recently, a new SUZE-2 circuitry became available which can encode an entire window of 4×5 pixels into a single 30-bit data word. The principle remains the same as for SUZE-1, i.e. the relevant row is encoded in one separate word and the windows are encoded afterwards in one or more following packages. The number of hits per SDS-bank and row is hereby limited

¹DDR – Double Data Rate




Figure 3.11: The geometrical parameters of current MIMOSA sensors with their corresponding active and passive areas, in comparison. In addition, there is a small dead area of up to 0.5 mm at the sides of each sensor. The MISTRAL sensor geometry is estimated from the relevant publications. Source: [83, 86, 90].

	MIMOSA-26	MIMOSA-28	MISTRAL	ASTRAL
integration time	115.2 μ s	185.6 μ s	30 μ s	20 μ s
frame rate	8.7 kHz	5.4 kHz	33.3 kHz	50 kHz
encoding	SUZE-1	SUZE-1	SUZE-2	SUZE-2
word size	16 bit	16 bit	30 bit	30 bit
output channels	1-2	1-2	3-6	3-6
bitrate per channel	40 - 80 Mbit/s	80 - 160 Mbit/s	320 Mbit/s	320 Mbit/s
output bandwidth	160 Mbit/s	320 Mbit/s	1.92 Gbit/s	1.92 Gbit/s
peak data rate	19.9 MB/s	40.0 MB/s	\approx 240 MB/s	\approx 240 MB/s
power consumption	250 mW/cm ²	150 mW/cm ²	200 mW/cm ²	85 mW/cm ²

Table 3.1: The characteristics of the most recent MIMOSA generations. The values for MISTRAL and ASTRAL are taken from related test chips MIMOSA-32, MIMOSA-34 and AROM0. Both are assumed to be composed of three independent FSBBs. The values are only approximate and can change in future. Source: [83, 86, 90].

as well. First SUZE-2 circuits will be integrated into MISTRAL¹ and ASTRAL² [90] sensors, which are being developed in the scope of the ALICE-ITS³ upgrade program. The access to a new CMOS technology (0.18 μ m TowerJazz [91]) has opened a plethora of possibilities for MISTRAL and ASTRAL. The smaller feature size allows higher logic density and better radiation tolerance. The new process additionally offers six metallization layers allowing the implementation of quadruple wells. This allows the integration of a deep p-well into the pixel which serves as a shield against charge loss from the epitaxial layer to the upper n-wells. Thus, the CMOS technology can be applied inside the pixel microcircuit enabling pixel-level comparators, which would reduce the overall power consumption and noise. This option is explored with the ASTRAL generation, whereas the MISTRAL generation follows the more conservative approach of a column-level comparator. Both of the new sensors support a rolling shutter readout mode scanning two rows simultaneously to double the readout speed. An output bandwidth of 320 Mbit/s per channel is anticipated. The sensors are being designed in a modular way. They will be composed of one or more Full-Size Building Blocks (FSBBs) which operate independently, yet in

¹MISTRAL – MImosa Sensor for the inner TRacker of ALice

²ASTRAL – AROM (Accelerated Read-Out MImosa) Sensor for the inner TRacker of ALice

³ITS – Inner Tracking System

parallel on the same chip. Current FSBBs are optimized for the ALICE requirements, however most of the design aspects coincide with CBM, e.g. the readout time, material budget, and radiation hardness. The main difference between the two experiments is the expected hit density. In contrast to ALICE, CBM is a fixed target experiment and will be therefore exposed to more impinging particles per area. The MVD hit density and its impact on the data rates are studied in the following chapter.

Chapter 4

Conceptual Design of the MVD Readout

The CBM Micro-Vertex Detector is presently in the prototyping stage. Its leading design principles have been extensively studied and are well known, up to now. This allows for developing a fundamental readout concept for the upcoming FAIR experiments.

Section 4.1 introduces the general MVD properties focusing on the detector layout and the planned applications at FAIR. Based on these principal considerations, key requirements regarding the detector performance can be elaborated. Thus, in section 4.2, a variety of simulations is performed using a virtual sensor based on the MIMOSA family. The studies allow an estimate of the data rates, sensor limitations, and particle fluxes of the most demanding SIS-100 experiments. Section 4.3 focuses on an appropriate MVD interface to the crucial CBM parts, i.e. the CBMnet and the FLES. Lastly, in section 4.4, a suitable readout architecture is proposed.

4.1 CBM Micro-Vertex Detector at FAIR

CBM is designed as a fixed target experiment with the MVD situated directly after the target. Impinging particles will first traverse the MVD, before they reach the remaining detector stations. The system exhibits therefore high demands with respect to radiation tolerance and material budget. The thickness of the stations needs to fulfill $0.3 - 0.5 \% X_0$ to reduce multiple scattering. Furthermore, the MVD will operate inside a dipole magnet of 1 T. This allows only a limited space for placement and influences subsequent maintenance as well. It is foreseen to place the MVD together with the STS sub-detector in a thermally isolated enclosure of dimensions $1400 \times 2000 \times 1100 \text{ mm}^3$ [92], as shown in Fig. 4.1.

The detectors additionally need to cover polar angles of $2.5^\circ - 25^\circ$ which form the CBM acceptance window. In a related work [93] it has been demonstrated that at least three planar MVD stations are necessary to reach the main physics goals. However, recent studies support also a fourth station in order to improve the overall tracking efficiency. Hence, the MVD model used in this thesis considers four stations placed at 5, 10, 15 and 20 cm away from the target employing next-generation MIMOSA sensors for particle detection. In total, 292 sensors with dimensions of $3 \times 1.3 \text{ cm}^2$ are considered [94], distributed in the order of 8, 40, 84, and 160 sensors per each station, respectively. The sensors are organized into ladders and operated independently from each other. However, they need to operate fully synchronously in order to facilitate the timestamp allocation and to support the subsequent track reconstruction. The MVD is being designed




Figure 4.1: The MVD is placed inside the CBM magnet. A vacuum vessel (orange) separates it from the STS. The image on the right-hand side illustrates the front-view of an MVD station. Source: [92].

in a modular way where each station can be further subdivided into quadrants. All quadrants are equal for a given station. They are rotated by 90° with respect to each other and placed around the beam pipe, as indicated in Fig. 4.1. Due to the relatively large, passive readout area, as known from MIMOSA-26, the sensors are arranged in two layers per station. The sensors in the second layer are shifted with respect to the first layer in order to cover the passive regions.

The tracking efficiency can be further increased by placing the MVD inside a vacuum chamber. The target is placed inside the vacuum as well, in order to avoid any additional separating panels. This reduces the secondary particle production and multiple scattering effects resulting from collisions of particles with air molecules and panels. On the other hand, it demands vacuum compatible detector components. The mechanical integration of the MVD is addressed in [94] which also proposes a vacuum vessel with dimensions of $650 \times 540 \times 285 \text{ mm}^3$. The vessel is designed to hold all the necessary support structures and to provide feed-throughs for cooling pipes and sensor cables. Cooling is performed with a liquid, e.g. liquid carbon-dioxide. All of the active electronic components will be placed outside of the vacuum vessel to minimize the radiation dose.

The CBM detector is designed as a multi-purpose device covering a wide range of physical observables. The MVD will adopt several different roles in the scope of the planned physics program. Moreover, it needs to support both FAIR accelerators.

At SIS-100, the MVD is mainly required for the reconstruction of secondary decay vertices of D-mesons emerging from p-A collisions [35] (see section 1.3.4). However, considering the heavy-ion experiments, the D-mesons are below their nucleon-nucleon production threshold. Nevertheless, the MVD could still be used to improve the overall tracking performance for low-mass vector meson and hyperon studies, in certain cases. Moreover, recent studies support a method to reduce the background of electron-positron pairs (di-electrons) [95, 96]. Di-electrons of interest emerge from vector-meson decays, however photons which produce di-electrons via γ -conversion when traversing detector material are known to increase the combinatorial background and complicate the subsequent analysis. The MVD could be used to suppress the γ -conversion background and therefore find application in all running scenarios, including heavy-ions. The heavy ions will involve some lighter systems (e.g. Ca-Ca), as well as Au-Au up to the maximum beam energy (14 AGeV and 11 AGeV, respectively). The maximum of the SIS-100 capabilities is represented by Au-Au collisions at 11 AGeV and p-A collisions at 30 GeV. The

desired collision rates are in the order of several 10 kHz and few MHz, respectively.

With the SIS-300 accelerator, which will be constructed few years after SIS-100, the CBM will finally reach its full potential. Higher beam energies and collision rates will be accessible allowing the study of principal CBM physics goals. A collision rate in the MHz range for Au-Au experiments (up to 35 AGeV) is desired according to current simulations [38, 39].

4.2 Performance Requirements and Limitations

The detection of rare observables in CBM requires high collision rates, which in return demand a fast sensor readout speed. This calls for a high-performance readout system with the ability to transport large amounts of data generated by particle hits in the detector stations. The points of interest are hereby defined as the expected data rates, as well as the limitations given by the current sensor technology standard.

The sensor and readout electronics performance is driven mainly by the progress in the commercialized CMOS technology, which follows a gradual improvement every few years according to Moore's law [55]. Hence, the technology available at SIS-300 is expected to outperform the technology applied at SIS-100. However, SIS-300 construction is scheduled several years after SIS-100 and lies far in the future. Therefore, this section will only cover requirements set by the p-Au and Au-Au collision systems at SIS-100 for the scheduled maximum beam energies and the proposed event rates.

4.2.1 Simulation Method

Heavy ion collisions at CBM energies can be sufficiently well simulated with the UrQMD¹ model [97, 98]. Version 3.3 of the model is applied within this thesis to obtain particle multiplicities and their phase space distributions directly after the collision. Two collision systems are particularly interesting at the SIS-100: the Au-Au collisions which will produce the highest hit occupancies, and p-Au collisions which will feature the highest collision rates. Both are studied extensively within this thesis. All events are prepared as minimum-bias collisions².

The initial particle tracks obtained from UrQMD simulations are propagated throughout the detector via the GEANT3³ software package. The CBMroot framework [51, 99] is utilized to obtain a realistic detector response model. The simulated CBM setup contains the magnet (version 12b), the beam pipe, the target, and the MVD. The MVD model features four stations with 292 sensors, partitioned into 876 independent FSBBs, as described in the previous section. The generated particles are influenced by the magnetic field and may create additional secondary particles in the detector materials. Upon traversing the active sensor volumes, particle hits can get detected by applying certain semi-empirical digitization methods. The MVD digitizer software [100] uses the charge deposited in the sensor epitaxial layer in order to determine the active pixels. Parameters from MIMOSA-26AHR studies are hereby included as a reference. A hit usually

¹UrQMD – Ultra-Relativistic Quantum Molecular Dynamics

²In terms of CBM, minimum bias collisions are collisions with random centrality, i.e. they include central collisions as well. However, the central collisions occur with low probability.

³GEANT – GEometry ANd Tracking

Collision System	Interaction Rate	N_{ions}	N_{coll}	Note
Au-Au, 10 AGeV	10 - 100 kHz	30 - 300	0.3 - 3	average case
Au-Au, 10 AGeV	10 - 100 kHz	90 - 900	0.9 - 9	worst case
p-Au, 30 GeV	1 - 10 MHz	3000 - 30000	30 - 300	average case
p-Au, 30 GeV	1 - 10 MHz	9000 - 90000	90 - 900	worst case

Table 4.1: Collision systems simulated. The interaction rate of the corresponding experiment does not change between the average and the worst cases, since the worst case is expected to occur on very small time scales and, on average, cancels out.

activates an area of several neighboring pixels, denoted as a cluster. The applied charge threshold is manually set to a moderately low value with the effect that the particles produce larger pixel clusters than the ones measured with MIMOSA-26AHR. The data is therefore slightly overestimated, in order to set a higher constraint on the studied observables.

The simulations are performed frame by frame. They are based on a sensor integration time of $30 \mu\text{s}$. In the planned experiments, the thickness of the CBM target will be fixed, hence the probability of beam particles to interact with the target will be constant as well. Therefore, the average beam intensity and the collision rate can be linked to each other by a constant factor. Both can be specified on the sensor frame level as well. The average number of beam particles traversing the target during one sensor frame is denoted by N_{ions} , whereas the average number of colliding ions per frame is denoted by N_{coll} . In the following, an interaction probability of 1% is assumed, hence only 1% of the simulated beam particles actually collide with the target. Therefore, N_{ions} and N_{coll} differ by exactly two orders of magnitude, i.e. a factor of 100, in the applied simulation model. Due to the fact that each collision occurs with a certain, fixed probability, statistical fluctuations have to be considered. In general, Au-Au collisions exhibit large particle multiplicities, therefore it is desired to keep N_{coll} low in order to distinguish the individual collisions and facilitate the subsequent track reconstruction. This allows the application of the Poisson distribution to obtain a realistic N_{coll} value for each of the simulated sensor frames. The standard deviation is therefore in the range of $\sqrt{N_{\text{coll}}}$. Moreover, the Poisson distribution is applied for N_{ions} as well, in order to add a certain level of randomness into the system. For larger beam intensities, or very high collision rates, the Poisson distribution obtains Gaussian characteristics.

The performed simulations are presented in table 4.1. As can be seen, each collision system is simulated twice, once without a security margin (average case) and once accounting for a security margin of three (worst case). The worst case is motivated by the measured beam intensity fluctuations at the SIS-18 synchrotron. The actual time structure of the beam depends on subtle details of the realization of the SIS-100 synchrotron, which are not yet fully known. However, a similar tendency as for the SIS-18 is predicted for the FAIR accelerators as well. On the time scale of 1 ms, SIS-18 measurements [101] indeed indicate a distribution which can be sufficiently well approximated with the Gaussian function. Moreover, the non-trivial maximum-to-average ratio related to the beam time structure can be obtained as well. It reaches a value of three at the $30 \mu\text{s}$ time scale. Thus, in the worst case, the beam intensity is expected to increase threefold with respect to the average, for a short duration in the order of one to several frames. It is therefore not sufficient to simulate the average case only. The security margin of three needs to be considered as well.

Each collision rate is simulated with 1000 frames. The number of ions passing through the target and the number of collisions for the i -th frame is obtained from the following equations:

$$N_{\text{ions}_i} = \text{Poisson}(N_{\text{ions}}), \quad N_{\text{coll}_i} = \text{Poisson}\left(\frac{N_{\text{ions}_i}}{100}\right)$$

In addition to the N_{coll_i} UrQMD events, N_{ions_i} beam particles passing through the target material in each frame are propagated with GEANT3 as well. These non-colliding particles are found to generate δ -electrons¹ in the target material, which are influenced by the magnetic dipole field and focused on certain detector regions. These MVD regions therefore exhibit higher hit occupancies. However, the effect is negligible for p-Au collisions, since the probability to generate a δ -electron depends on the charge (Z^2) of the impinging ion.

4.2.2 Applied Sensor Model

The simulations are performed with a virtual MIMOSA sensor named Mimosi1, which is based on the proposed MISTRAL sensor parameters and MIMOSA-26AHR characteristics. The active area covering $1 \times 3 \text{ cm}^2$ is divided into three FSBBs of 1 cm^2 size (see section 3.4). The FSBBs are treated as independent sensor modules with their own memory and output channels. The modeled FSBB pixel matrix contains 300×454 pixels of $22 \times 33 \mu\text{m}^2$ size. A high-resistivity epitaxial layer of $18 \mu\text{m}$ thickness is used. The hits are encoded with the SUZE-2 algorithm producing 30-bit words. The virtual sensor operates at an integration time of $30 \mu\text{s}$. Due to the expected bandwidth limitations, only 640 words per FSBB can be read out during one frame.

Data words are divided into **states** and **windows**. As shown in Fig. 4.2, the SUZE-2 algorithm is able to encode clusters up to the size of 4×5 pixels in one word named *window*. Windows always have the maximum size and may contain one or more particle hits. Larger clusters are partitioned into several, non-overlapping windows. This assures that each pixel is counted exactly once. Furthermore, a *state* encodes the address of the row containing one or more windows. Four neighboring rows are hereby grouped together and have the same address, forming 75 row

¹The δ -electrons are electrons which are knocked out from the atomic shell by the impinging beam ions.




Figure 4.2: An example of the implemented SUZE-2 encoding algorithm. Three clusters are shown, which are grouped into three different windows. The upper left corner (denoted by a red circle) determines to which row group the window is attributed. In this example, all three windows belong to the row group with base address 8, thus only one state needs to be created.

groups in total. Only one state per row group needs to be created, even if there are more than one windows found within.

Due to the data sparsification logic, the sensor is expected to have certain limitations, as indicated in [90]. Windows are encoded within groups of SDS-circuits, each of which is responsible for a 32-column wide segment of the pixel matrix. The Mimosi1 FSBB contains therefore 15 SDS-circuits in total¹ operating fully in parallel. Windows may expand over the boundaries of the individual SDS circuits. However, only 6 out of the seven possible windows per SDS-circuit can be encoded. Next, the SDS-circuits are grouped into two nearly symmetric banks, where only 9 windows per bank can be read out. Finally, the total number of windows per row group may not exceed 18. If that number is exceeded, at least one of the banks have an overflow as well. The bank overflow is denoted by a certain bit within the state word.

4.2.3 Expected Data Rates at SIS-100

The data rates were obtained from the encoded SUZE-2 data words, without the effect of sensor limitations. A fixed overhead of 120 bits per frame is added accounting for the Header, Trailer, and other format words. All simulations include a nominal fake hit rate of 10^{-5} . A fake hit usually produces 2×30 bits, one word for the SUZE-2 state and one for the window. However, more than one fake hit may occur within the same row group of the FSBB, in which case only one state word is generated. The same applies for real particle hits. This leads to the fact that the number of generated SUZE-2 windows scales linearly with the collision rate, however, the number of SUZE-2 states does not. Due to the increasing occupancy, more hits are recorded within the same row group. In the following, two cases are presented.

In the first study, only the average case from table 4.1 is considered. The data volume per frame generated by the individual FSBBs is summed up to obtain the total MVD data rate. For each collision rate, the results were averaged over all 1000 simulated frames and presented in Fig. 4.3. The plot gives an estimate on the data rate which will be provided by the MVD over time. The individual contributions from fake hits, the sensor format overhead, and SUZE-2 windows and states are shown separately. As can be seen, the data rate increases nearly perfectly linearly with the collision rate.

The second study concerns the worst case data rates which will be present at a time scale of $30 \mu\text{s}$, due to the maximum-to-average beam intensity fluctuations of three, as observed at the SIS-18. The data rates are obtained with the same method as in the first study. The resulting worst case MVD data rates are presented in Fig. 4.4. The plots include the average case (dashed line) for comparison. In addition to the average and worst cases, several upper limits on the data rate are presented as well, which are described in the following.

The simulated data rates are subject to statistical dispersions which can not be easily parametrized. Firstly, the centrality of the colliding ions is random and varies with each collision. In general, p-Au collisions generate less particles than Au-Au collisions. Thus, the centrality dispersion effects are intrinsically smaller for p-Au than for Au-Au simulations. Secondly, the applied Poisson distributions show a standard deviation σ of size $\sqrt{N_{\text{coll}}}$. Comparing the magnitude of σ to the expectation value, one can find that the ratio $\sqrt{N_{\text{coll}}}/N_{\text{coll}}$ is larger for lower values of N_{coll} . For example, for $N_{\text{coll}} = 1$ the standard deviation amounts to 100% of the expectation value, however for $N_{\text{coll}} = 10$ the magnitude is only 32%. Therefore, Au-Au simulations are

¹The last SDS-circuit is only 6 columns wide.




Figure 4.3: The average MVD data rates for the planned Au-Au and p-Au experiments at the SIS-100.

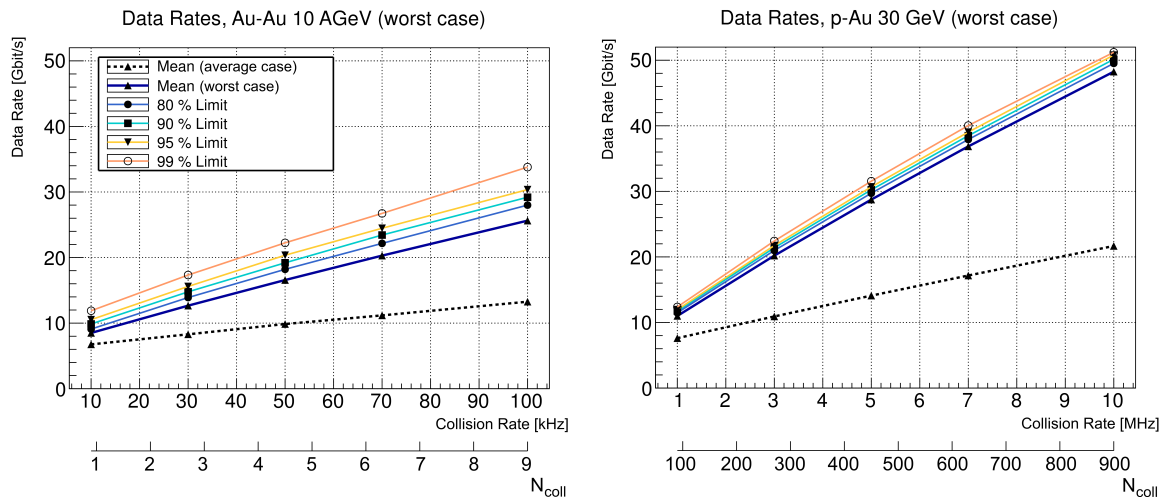


Figure 4.4: The worst case data rates per frame, assuming a threefold increased beam intensity. The presented N_{coll} values on the x-axis are denoting the increased intensity, however the collision rate refers to the original, average case.

featuring a higher statistical dispersion than p-Au. Hence, to provide a more constraining result, the data rates are additionally bounded by the upper limit derived from the relative statistical population of the 1000 simulated frames. Each limit in Fig. 4.4 represents a data rate which was not exceeded by the denoted percentage of the simulated frames¹.

The analysis has shown that the worst case increases the average data rate by a factor of 1.8-2.2, only. This appears to be in contradiction with the expected value of three which is used to scale the N_{coll} . In order to find the reason for this discrepancy, the FSBBs are analyzed separately with the result that the data load is not equally distributed among them. The particles are focused on certain detector regions, denoted as hot spots. Thus, only a fraction of the FSBBs carries the majority of the data. The most exposed FSBBs have shown an increase of ≈ 2.5 with respect to the average case. The idealized factor of three can never be reached in practice due

¹For example, the 99% limit was not exceeded by 990 frames per simulated collision rate.

Name	Maximum per FSBB	Limit
Data Overflow	1	640 data words
Row Overflow	75	18 windows per row group
Bank Overflow	2×75	9 windows per bank
SDS Overflow	15×75	6 windows per SDS module

Table 4.2: Four limitations are studied within this thesis. Maximum per FSBB refers to how often the limitation can be exceeded within the simulated frame, per each FSBB.

to the fact that the increasing number of particles also increases the probability to encode two or more clusters within the same row group. Therefore, the increased beam intensity does not affect all sensors, and those which are affected do not scale linearly with the intensity, hence the obtained worst case to average data rate ratio is not three, but rather two.

Without considering any sensor limitations, it can be concluded that the entire MVD will provide data rates in the range of 6-14 Gbit/s and 6-22 Gbit/s for the planned Au-Au and p-Au experiments, respectively. However, the beam fluctuations are expected to temporarily increase the average data rate by a factor of ~ 2 . In this case, the most exposed sensors will increase their output by a factor of 2.5. Therefore, in order to avoid data loss and bottlenecks, the MVD readout system needs to support a fluctuating data input of up to 50 Gbit/s and a security margin of three for the individual sensor output rates.

4.2.4 Sensor Limitations Study

The simulated FSBBs are mainly delimited by the properties of the SUZE-2 zero suppression circuitry. The internal memory of the applied sensor model is assumed to be sufficiently large. As observed in all MIMOSA generations with digital data output, the sensors are able to store more data than can be output during one frame. Therefore, memory limitations are replaced by the more constraining output bandwidth. In the current model, 640 Mbit/s are assumed per FSBB allowing a maximum of 640 data words per frame. Thus, a data overflow denotes whether this bandwidth limit has been exceeded. It can occur only once per FSBB for the simulated frame.

Moreover, three further limitations are studied which are related to the SUZE-2 logic. They are summarized in table 4.2. Since SUZE-2 reuses its logic components, they can occur more than once per frame.

The limitations can be studied for the individual FSBBs, sensors, stations, or they can be integrated over the entire MVD. Within this section, the individual overflows from all of the 876 FSBBs comprising the MVD are summed up and evaluated altogether. The preliminary results of the study suggest that mostly no overflows are to be expected as long as the detector is operated with up to 100 kHz interaction rate for Au-Au or 10 MHz for p-Au experiments. This picture, however, does change once the factor three security margin for beam fluctuations is considered. In such case, the limitations are exceeded more frequently. The total number of overflows registered in the MVD per simulated frame, considering only the worst case, is presented in Fig. 4.5. The distributions are particularly for lower collision rates centered around a value of 0, denoting that they did not occur during the given number of simulated frames. Thus,

4.2. PERFORMANCE REQUIREMENTS AND LIMITATIONS




Figure 4.5: The individual distributions of overflows within the MVD. Only the worst case, including the security margin of three, is considered.




Figure 4.6: The distribution of the generated data words among the most exposed FSBBs. Each point represents a value which was not exceeded by 99 % of the simulated frames. The first 24 FSBBs belong to the first station, the others to the second.

in most cases, the MVD would not lose any data and could be operated at the specified collision rates. However, an increase of the collision rate directly affects the probability of an overflow.


In general, the limitation effects are stronger in Au-Au simulations, than in p-Au. The data overflow saturates at a value of six for the highest chosen collision rate, showing a non-symmetric distribution. The value of seven is never reached. This signifies that there are non-uniformities in the system. The data load is not equally distributed among the FSBBs due to hot spots, as indicated in the previous section. This matter is studied in more detail in the following section. Furthermore, the bank and row overflows occur frequently as well. However, in most cases, the lost data can be backtraced to a certain pixel segment ($\leq 1 \text{ mm}^2$), thus minimizing the loss of information¹.

In the following, the limitations driven by the SUZE-2 logic and the effects of the output bandwidth are studied separately. This is motivated by the fact that a single data overflow can be responsible for the loss of the majority of data, far exceeding the SUZE-2 contributions. Therefore, the bandwidth limitation is studied first. The most occupied FSBBs are located within the 24 FSBBs of the first MVD station. The average number of data packets generated during one frame, without the effect of SUZE-2 limitations, is displayed in Fig. 4.6 for the studied worst case considering the tripled beam intensity. Only the first 30 FSBBs are shown. As indicated before, p-Au collision do not exceed the given output bandwidth. However, in Au-Au collisions, six FSBBs belonging to three different sensors are particularly exposed to more particles than others. These are FSBBs 0 and 1 of sensor 0, FSBBs 7 and 8 of sensor 2, and FSBBs 10 and 11 belonging to sensor 3². At higher collision rates, they are very likely to exceed the FSBB output bandwidth losing up to one half of the generated data, as is the case with the FSBB 8.

Regarding the SUZE-2 limitations, a comparatively smaller number of data is lost. All three SUZE-2 limitations are combined together to obtain the total number of overflows per frame.

¹If a bank or a row overflow occurs, a certain bit within the SUZE-2 state word can be set, denoting which bank is affected. Hence, the missing window can be backtraced to the corresponding bank. In worst case, the lost data can be localized to a segment of 7×185 pixels corresponding to an area of 1 mm^2 . However, this assumes that the bandwidth limitation is not exceeded.

²The systematic sensor numbering is obtained from the current CBMroot implementation.



(a) Average number of lost windows per frame. (b) The percentage of lost windows with respect to the total yield.

Figure 4.7: Lost windows caused by the SUZE-2 limitations. The worst case considers a tripled beam intensity with respect to the average case.

Afterwards, this number is averaged over all 1000 frames and displayed in Fig. 4.7. The result is shown twice. In Fig. 4.7a, the absolute value is shown denoting how many windows have been lost, on average, due to the combined SUZE-2 limitations. However, Fig. 4.7b shows the relative amount with respect to the total number of generated windows. Considering the low magnitudes in the plots, all data points below a certain limit¹ exhibit a large standard deviation. Since most of the simulated frames do not generate any overflows, nearly all data points below the indicated horizontal line are zero-compatible. Therefore, to improve the readability of the plots, the trivial deviation which extends down to zero is not displayed. Only three data points, which are not zero-compatible, are displayed with the standard deviation. The deviation represents the statistical uncertainty and can be reduced with additional simulations. Particularly in collisions with the average beam intensity (average case), the average number of overflows per frame is strongly fluctuating. In these cases, the overflow distribution is purely random and requires more statistics to define a more accurate average value.

However, in case of the tripled intensity (worst case), a certain trend can be observed. The number of lost windows increases exponentially with the collision rate. The difference between two neighboring data points reaches nearly one order of magnitude. Au-Au collisions are particularly affected. The ratio of lost windows compared to the total number of windows reaches 0.15 % for 70 kHz, and 1 % for 100 kHz, not including bandwidth limitations. These values show a very low uncertainty and can be used as a constraint for further MVD studies. Since the hit anisotropy tremendously affects the MVD, it is studied in the following section.

4.2.5 Expected Hit Distributions

The sensors positioned in the vicinity of the beam axis are expected to have higher hit occupancies due to the high forward-rapidity of particles emerging from the fixed-target setup. Moreover, δ -electrons, which are focused on certain detector regions by the magnetic dipole-field, increase the occupancy of particular sensors in Au-Au experiments [102, 103]. Both statements are confirmed within the studies presented in this section.

¹The limit is shown only for Au-Au collisions.

The average particle density in Au-Au and p-Au collisions for the most exposed sensor layer is shown in Fig. 4.8. The layer merely contains sensors 0, 2, 4 and 6¹. The hit distribution strongly depends on the collision system. Due to hot-spot regions in Au-Au experiments, some sensors will carry a dominant part of the load (sensor 2), whereas others will remain idle most of the time (sensor 6). A similar tendency is observed for p-Au collisions as well, however in this case the most occupied sensors are found in the vicinity of the beam axis.

The sensor occupancy is defined as the percentage of firing pixels per sensor. Due to the sensor limitations, only a maximum of 8.24 - 8.85 % can be reached, depending on the distribution of hits within the rows of the pixel matrix. To gain a better understanding of relative hit distributions between the individual sensors and to analyze their difference in performance unrestricted of any limitations, the overall data load without the limitations is examined next. All 292 sensors are analyzed separately. The data volume generated by the individual sensors is averaged over all collision rates and frames. The result is scaled to the percentage of the data volume generated by the entire MVD, and presented in Fig. 4.9. The vertical lines in the image separate different stations. As can be seen, the relative data load decreases with the distance from the target. The first MVD station clearly generates the highest data volume per sensor, regardless of the collision system or collision rate.

Based on the average data output from Fig. 4.9, the sensors are grouped into four classes. The class allocation is performed manually. The principal idea is that the class i contains the highest contributions from the station i . In particular, sensors 0, 2 and 3 for Au-Au collisions, and sensors 0, 2, 4 and 6 for p-Au collisions generate the highest data rates in the first MVD station. Thus, they are grouped into the occupancy class 1 of the corresponding experiment. Other sensors are grouped into different classes according to their average data output. Afterwards, the overall data collection efficiency of the individual sensor classes can be analyzed. It is defined as:

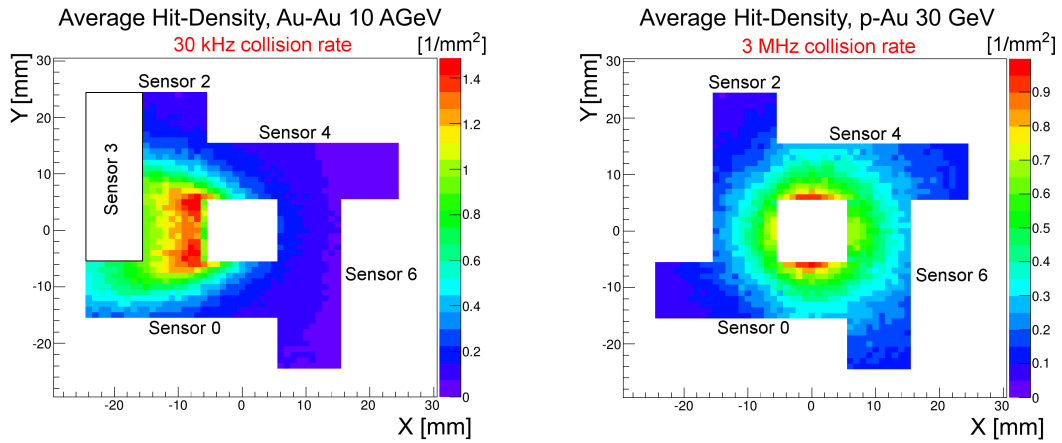
$$\text{Efficiency} = \frac{N_{\text{Total}} - N_{\text{Lost}}}{N_{\text{Total}}}$$

where N_{Lost} represents the number of data words which were lost due to any of the four studied limitations, and N_{Total} stands for the total number of data words that would be generated if no limitations were present. As can be seen in Fig. 4.10, class 1 is the primary source of data loss in the MVD. The performance of sensors in the δ -electron region is greatly reduced, hence Au-Au experiments will define the highest requirements on the final sensor model and the readout system. As shown in the previous section, the overflows mostly depend on the sensor output bandwidth, however an optimization of the SUZE-2 circuitry with respect to the bank and row overflow limit can improve the performance as well.

4.3 Interfaces to CBM Modules

The MVD will be integrated into the global CBM readout network. This calls for an implementation of the CBMnet network protocol and a suitable interface to the FLES computing farm.

¹Even sensor numbers belong to the first layer, whereas odd sensors belong to the second.



(a) The δ -electrons in Au-Au simulations are bent to the left by the magnetic field. (b) The hit distribution is radially symmetrical in p-Au collisions.

Figure 4.8: The average hit density distribution within the first layer of the first MVD station at 5 cm.




Figure 4.9: The contributions of the individual sensors to the total data volume generated by the MVD per frame. The result is averaged over all frames and collision rates. Each class contains sensors which have their data points in the given intervals, separated by the horizontal lines.




Figure 4.10: Performance of individual classes with respect to data loss.


4.3.1 CBM Network Protocol

The CBM detector is operated over a synchronous, high-performance, optical network named CBMnet [44, 45, 104]. The corresponding hardware for the CBM front-ends is currently under development. The principal design goals are led by the requirement for fast, compact, scalable and versatile detector network with a large focus on global synchronization and free-streaming data acquisition.

The network uses hierarchical point-to-point connections over a single optical transceiver to distribute the clock and collect the data. The clock needs to be appropriately recovered in order to supply the receiver and further nodes with the high synchronization level required by CBM. Therefore, the transceivers are started in an initialization phase, where they perform clock recovery and phase alignment. This requires low-level access to the transceiver hardware, e.g. their barrel shifters. During the initialization phase, the network is synchronized within one-bit accuracy of the optical transceivers and the clock is recovered with a peak-to-peak jitter of 80 - 100 ps, as tested with Xilinx FPGAs and Multi-Gigabit Transceivers (MGTs) [44, 45]. The clock frequency depends on the link speed. For 2.5 Gbit/s links, the recovered clock is 125 MHz [45]. Furthermore, an external jitter-cleaner device is used to reduce the clock jitter below 40 ps. The jitter cleaner hardware is shown in Fig. 4.11. The recovered, cleaned clock can then be used by the FPGA logic. This strategy allows the application of a common master-clock from a single source for the entire CBM readout network. User modules on the FPGA may be asynchronous, but the transmission via the CBMnet module should always take place with the recovered clock. By aligning the phases of all transceivers and by using one common clock, the network obtains deterministic characteristics.

CBMnet implements the first four layers of the OSI¹ layer model. There are three traffic classes, each assigned to a separate virtual communication channel. The channels share the optical transceiver with different priorities. All messages are therefore partitioned into detector data

¹OSI – Open System Interconnect






Figure 4.12: The structure of CBMnet messages. Start Of Packet (SOP), Start Of Slow Control (SOSC), as well as their corresponding terminators (EOP, EOP_C) are special comma characters, similar to DLM. The payload of DTM and DCM is always CRC stamped allowing the detection of transmission errors, in which case the message can be retransmitted. Source: [45].

(DTM¹), slow control (DCM²), and global network synchronization (DLM³). Some additional low-level control messages are automatically exchanged, e.g. acknowledge, idle, etc. They are encoded in 16 bit and form after 8b/10b encoding two special 10-bit characters. All network packets are transmitted in a 16 bit format and routing is performed within a 16 bit address space as well. The small packet size supports network load balancing and reduces the buffer sizes. DCM can have a sender and receiver address attached in front of the control message. Thus the network allows precise point-to-point routing of slow control and monitoring messages. Since all CBM front-ends operate in the data push mode, the DTM can have only the sender address placed in front of the data stream to identify the data source. Each DCM and DTM carries an 8-64 Byte large payload. The message structure is presented in Fig. 4.12.

The DLM, on the other hand, do not have any actual payload. They are transmitted as pure comma characters (2×10 bit, see [45]). As the name suggests, they have deterministic properties and can be used for signal latency measurement and subsequent synchronization. Once the optical link communication is established, the DLM can be transmitted with the priority request insertion. This means, that any active transmission on the link is immediately interrupted and the DLM comma characters sent instead. The interrupt takes place at the bit-level of the optical link. With this method, the latency of a DLM is always the same, within the accuracy of few hundred ps, based on the link speed. This fact can be used to synchronize all network nodes over the optical links. After the DLM, the interrupted transmission is continued.

There are 16 DLM types foreseen. For example, DLM0 is used for the initial latency measurements. It is transmitted to all nodes and reflected back to the sender. By counting how many FPGA clock cycles the message requires to return, the sender can create a latency table for all network nodes. According to this table, DLM1 and DLM2 can be sent to synchronize all timestamps and experiment-specific (epoch) counters. Afterwards, DLM1 is used to periodically check whether the synchronization remains during the detector operation. In order not to disturb the data flow too much, only one DLM is permitted each 35 clock cycles [45]. There are additionally six DLM types planned for user commands.

All DCM and DTM are CRC⁴ checked in order to avoid transmission errors. In case of

¹DTM – Data Transport Messages

²DCM – Detector Control Messages

³DLM – Deterministic Latency Messages

⁴CRC – Cyclic Redundancy Check

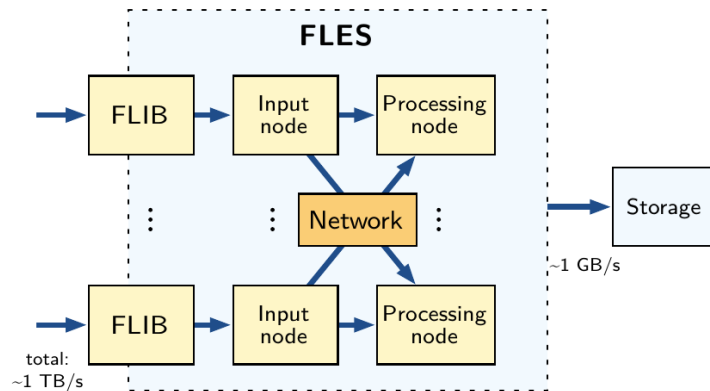


Figure 4.13: The FLES structure representing the final CBM readout stage. Source: [43].

an error, the broken message and all messages after the error are retransmitted. This ensures proper message ordering at the receiver. However, retransmission is only applicable under the assumption that transmission errors do not occur very frequently¹. Since DLM are vital for the detector operation, they are encoded, as all other control characters, using SEC-DEC codes². The effective bandwidth usage of CBMnet after 8b/10b encoding amounts to $\lesssim 73\%$. On a 3 Gbit/s optical link, this would allow up to ~ 2.2 Gbit/s of raw data, depending on data size.

4.3.2 FLES Interface


As already indicated in section 1.3.2, the free-streaming CBM front-ends will produce a data rate of ~ 1 TB/s. However, CBM is designed to measure rare events, thus a large portion of the data can be discarded. For this occasion, a high-level trigger is being developed to reduce the data storage rate down to ~ 1 GB/s. But in order to detect rare probes, they have to be fully reconstructed from their decay vertices. This calls for a sophisticated online event reconstruction algorithm. The First Level Event Selection (FLES) procedure applies a many-core computing cluster equipped with graphics cards for this challenging task. The computing cluster will use approx. 1000 input nodes to distribute the detector data to the computing nodes³. A high-throughput network structure is additionally required to handle the huge data load. For example, the InfiniBand network has been used for the initial test stage [38]. The FLES structure is outlined in Fig. 4.13.

Initially, the detector data needs to be collected first. For this task, a high-performance input and output PCB under the name FLES Interface Board (FLIB) [43] is being developed. The board applies Gigabit optical transceivers for data input from the detector stations. An on-board FPGA receives and pre-processes the detector data. Hereby, the input data links can be merged together, stored into memory and analyzed. A large memory of several GB is required to handle the input stream. An implementation of a Direct Memory Access (DMA) as a PC interface can be realized with the FPGA as well. The data can reach the PC memory via a PCI-express interface and a suitable PC driver. Thus, the FLIBs can be attached to conventional PCs receiving the


¹The primary source of errors are SEU.

²SEC-DEC – Single Error Correction - Double Error Detection. Hamming Codes with an optimal Hamming distance are used for this occasion.

³There were 60 000 processing cores estimated for the year 2010 [43].



(a) The FLIB is designed as a PC add-on card featuring several optical input links, an FPGA, large memory buffers and a PCI-express interface. Source: [43].



(b) The first FLIB prototype is based on the C-RORC board from ALICE featuring three high-density quadruple optical transceivers (12 links), a Virtex6 FPGA and a DDR3 memory. Source: [38].

Figure 4.14: The FLIB architecture and the first prototype.

TB/s large CBM data stream. The basic FLIB architecture is shown in Fig. 4.14. A prototype derived from the C-RORC¹ PCI-express add-on card from the ALICE experiment is currently under study [38].

Event Reconstruction

The entire FLES operation can be subdivided in two major parts. The first part, **interval building** [43], requires the entire CBM network for its operation. The self-triggered front-ends transmit their data without the ability to determine whether a collision in the target has occurred, or not. For example, some data packages could be attributed to a noisy detector channel, and some particles reach the detectors faster than others. Thus, the FLES does not receive a priori any information on when to start reconstructing the event. Instead, it needs to define certain time windows during which all data packets are grouped together and analyzed for possible event candidates. Since all subsystems run synchronously, a high-precision time-stamping mechanism is used to mark the corresponding data packets in time for subsequent interval building. The individual data containers are termed **Microslice Containers** (MCs) and consist mainly of a generic 32-byte header containing the timestamp, and of one or more detector-specific data segments [105]. All MCs are filled with data at fixed points in time defined by the control system, however their size varies with detector occupancy. They are combined together in the FLES into a larger data structure named **Timeslice Container** (TC). The TC contains a large fraction of MCs collected over a longer period of time, e.g. 100 μ s in order to account for different particle velocities and detector response times. The interval building procedure is outlined in Fig. 4.15. A proposal for the MC encoding of the MVD data is additionally given in Appendix C.

All MCs are pre-sorted and merged in the FLIBs. Afterwards, the sub-detector data which is relevant for the event reconstruction can be selected and read out via FLES input nodes. The data is forwarded through the network to a particular computing node responsible for a certain time window initiating a partial interval building of the corresponding TC. A particular TC contains many MCs holding the data from all sub-detector stations which is necessary to reconstruct the event at a given interval in time. TCs that overlap in time contain MC copies at the overlapping

¹C-RORC – ALICE Common ReadOut Receiver Card

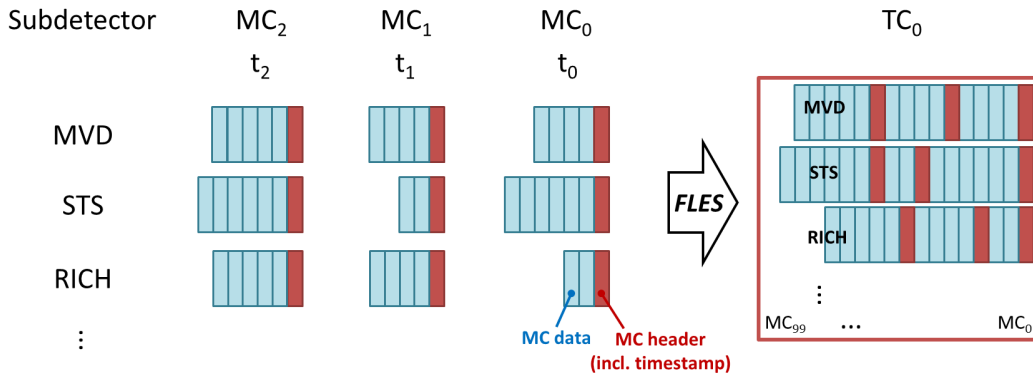


Figure 4.15: Different MCs are prepared at three given times t_0 , t_1 and t_2 (e.g. $1 \mu\text{s}$ apart) by the individual detector subsystems. Only one MC per subsystem is shown for simplicity. The data size is variable. Once an MC arrives in the FLIB, it is merged with others in order to create the larger TC data structure in the FLES, e.g. containing 100 MCs per subsystem. Subsequently, the TC can initiate the online event analysis at a designated computing node. Source: [106].

boundary and can be therefore analyzed completely independent from each other.

If a TC contains sufficient information, the computing node can initiate the second part of the FLES operation, the **online event analysis** [43] which runs locally on that node and involves the reconstruction of all particle tracks. The event reconstruction procedure is outlined in [40, 41]. With the presented strategy and a suitable scheduling routine, each computing node will be responsible for a different TC. However, each TC will usually hold more than one event due to a large time interval needed for a proper operation and due to occurring collision pile-ups at higher beam intensities. In order to improve the track reconstruction performance, a time-coordinate is added to the tracking algorithm. The procedure is termed 4D-tracking. If an interesting event is found, e.g. containing D-meson decays, the data from additional sub-detectors is requested and the full event is stored afterwards for offline analysis.

4.4 Proposed Readout Architecture

After defining the principal design parameters in the previous sections, an ideal readout system can be proposed.

4.4.1 Choice of Technology

The MVD readout system needs to support data concentration of ~ 300 sensors. A total input bandwidth of 50 Gbit/s should be available for the most demanding SIS-100 experiments. This number is derived from the largest simulated data rates, including a security margin of three to account for the worst-case beam fluctuations, as measured at the SIS-18 synchrotron. However, the data rates are not uniformly distributed among the sensors due to hot spots near the beam axis and due to δ -electrons. As the sensor occupancy classes depend on the experiment, the fundamental readout hardware needs to be easily modifiable to handle different sensor numbers, e.g. via add-on cards. The sensors in the outer MVD regions require fewer readout nodes with a high

fan-in, whereas the hot spot regions require more readout nodes due to the increased data load. This calls for a modularized and highly scalable readout system design. The high scalability supports early prototyping stages as well. Thus, customized front-end electronics can be used to provide basic readout capabilities in the proximity of the sensors, however more performant readout PCBs should be used for data processing afterwards. They are named Readout Controllers (ROCs), as defined by the standard CBM naming convention. The ROCs need to deserialize and decode sensor packages in order to recover the sensor data format. Moreover, FLES requires precisely time stamped data containers, hence the data needs to be pre-processed and re-encoded as well. In addition, the MVD readout network needs to operate fully synchronously. The CBM network protocol supports large-scale synchronization and provides the timestamps, as well as some additional control messages. Thus, the MVD readout system should support an appropriate interface to the rest of the CBM data acquisition system, ideally via a CBMnet implementation.

There are two possible technological solutions to design the readout hardware supporting all the aforementioned features, the ASIC- and the FPGA-based approach. An ASIC requires precise system specifications in order to reach its full potential. However, neither the sensor technology nor the FLES and CBMnet interfaces are fully specified up to now. Therefore, the use of such an expensive readout implementation is not advised. The FPGAs, on the contrary, support the early development stages and can also be prepared for the final experiment with a suitable firmware upgrade. This is the main reason for their application in numerous readout systems throughout the world, and the MVD should follow such a cost-effective approach.

In order to reduce the noise of the MVD readout modules, an application of optical transceivers for data transmission and communication is desirable. They considerably improve the data quality and should be used as early in the readout chain, as possible. Longer distances can be reached and the space consumption is reduced as well, with respect to copper cables. A data rate of several Gbit/s per fiber is presently found in most applications. Therefore, a moderate number of optical links would already suffice for the initial version of the MVD, at the SIS-100. The optical transceivers should directly interface the FPGAs for increased readout performance. This leads to a fundamental layout of the MVD readout system which is outlined in Fig. 4.16.

4.4.2 MVD Readout Topology

The generic CBM readout path which is applicable to several sub-detector is shown in Fig. 4.17, whereby the MVD is treated as a black box. The contents of the black box can be taken from Fig. 4.16. The entire data acquisition chain is subdivided into three parts. The first part is situated in the vicinity of the detector. Most systems, e.g. STS, TRD and MUCH, apply ASICs to amplify, digitize and read out the data. These detectors have a large number of channels, therefore the data is concentrated via Hubs and transmitted over optical links ('opto') to the second part. The second readout part is situated ~ 30 away in the CBM service building. It holds Data Processing Boards (DPBs) and the Experiment Control System (ECS). The DPBs are necessary to pre-process the large data stream involving several sub-detector systems and prepare the MCs. They will additionally implement the Detector Control System (DCS) receiving the master clock signal, as well as synchronization and control messages. The ECS is intended to steer the entire experiment. Control messages will reach the DCS via Ethernet. However, the common clock and the synchronization messages require a corresponding fully synchronized network architecture, e.g. the White Rabbit [107, 108]. Lastly, the pre-processed data reaches the ~ 700 m distant

'Green Cube' computing cluster in the third part, where the FLIBs and the FLES implement the online event analysis and store the corresponding event files.

The MVD differs from other sub-detectors in various aspects. Most importantly, the readout ASIC is integrated into the sensor pixel and the passive readout area, due to the monolithic technology. Additionally, the number of channels is very low. Thus, the Hubs and a specific DPB are not required. Instead, the sensors will be connected to several ROCs, as introduced in the previous section.

The optical Ethernet lines from the ECS can reach the ROCs to communicate all the control signals. One possible placement of the ROCs and the power supplies can be found near the HADES detector, approx. 10 m away from the MVD. However, it is advisable to place the ROCs in the CBM service building where they are protected from radiation. In that case, the distance between the FEE and the ROCs would require optical lines for data transmission, as well as radiation hard, optical front-end ASICs or an FPGA implementation of their transmission protocol. One possible solution, which is currently prototyped, can be found with the CERN GigaBit Transceiver (GBT) [110], however further studies are required.

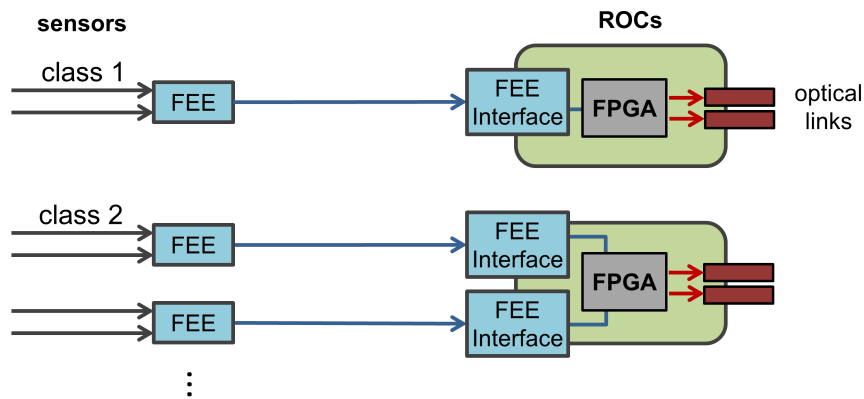


Figure 4.16: The MVD readout system is divided into modules. The radiation tolerant FEE provide some basic readout functionality, whereas Readout Controllers feature a powerful FPGA for pre-processing, one or more optical transceivers and suitable FEE interfaces. Higher occupancy classes, which generate less data and require less resources, can be supported via more FEE interfaces, e.g. designed as add-on cards. Therefore, the data output rate of each ROC can be kept at the same order of magnitude.

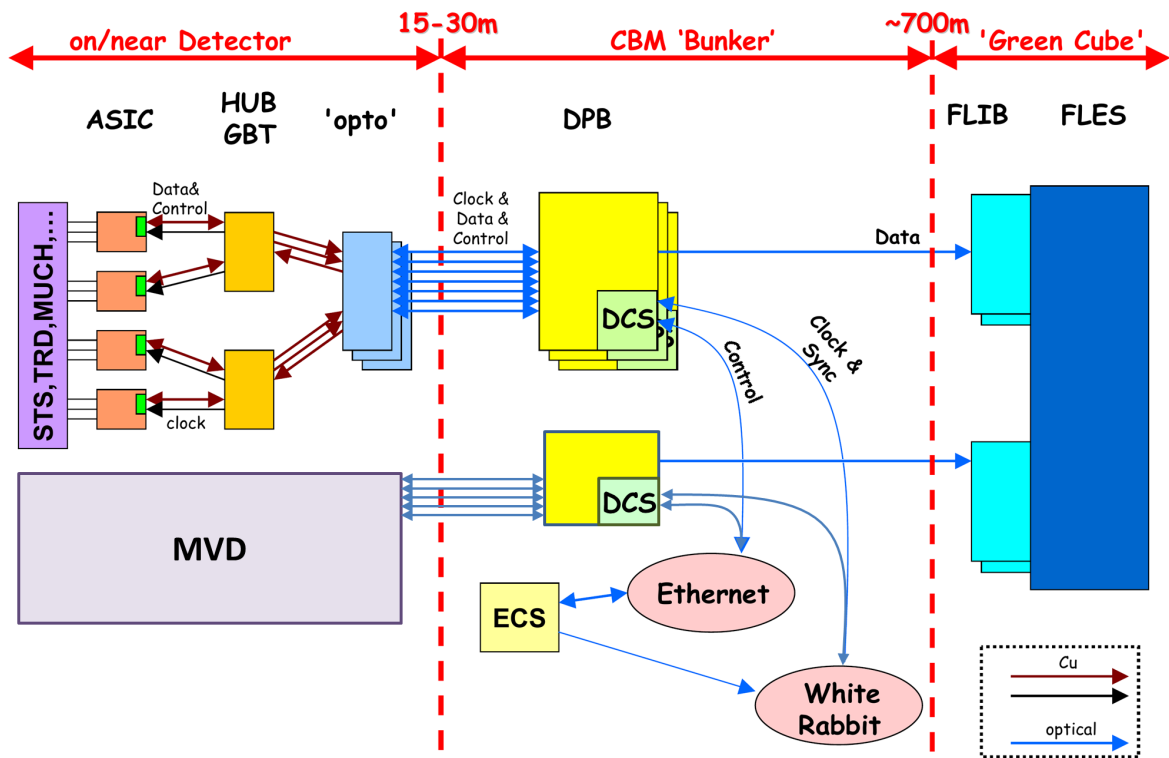


Figure 4.17: The CBM readout network distributes synchronization messages and the global clock via White Rabbit nodes to the subsystem DCS in the CBM service building ('Bunker'), in this picture. Most CBM sub-detectors use high-performance ASICs for data acquisition and control. The MVD, here depicted as a black box, can apply FPGAs for the readout due to the integrated readout functionalities within the sensor. The use of a DPB for the MVD is not planned up to now, as the data can be transmitted directly to the FLES. Source: [109].

Chapter 5

Readout Controller Development

Based on the studies from the previous chapter, a customized readout controller featuring an FPGA microchip, optical links, and add-on connectors would fulfill the basic requirements given by the CBM and the MVD electronics. In order to reduce hardware costs and invested development time, an already approved readout procedure has been considered and proven compatible. The HADES detector applies specialized TRB¹ boards and a corresponding network protocol, the TrbNet, for data acquisition. Since the hardware is already developed, debugged and tested in a large scale experiment, it is applied for the MVD development as well. In addition, a large variety of tools for monitoring, slow control and debugging became available.

As first, detailed ROC specifications and principal design requirements are presented in section 5.1, supplemented by an introduction into TrbNet. Subsequently, an FPGA firmware implementation for the MIMOSA sensor family is described in section 5.2. Afterwards, ROC prototypes based on the HADES TRB boards are introduced in section 5.3. Lastly, an implementation of a HADES-like readout network is used to validate the principal design choices, which is presented in section 5.4.

5.1 Design Specifications

The readout of MIMOSA sensors, despite the on-chip readout circuits, is not a trivial task. Detailed system parameters need to be specified in order to provide guidelines for the prototyping and testing stages.

5.1.1 General Requirements

As first, the MVD sensors can be operated independently from each other and will be therefore organized into ladders. However, the number of sensors in a ladder may vary, and the anisotropic hit distribution will occupy some sensors more than others. In order to distribute the data load equally among the readout nodes and reduce the number of applied ROCs, some FPGAs will be handling more sensors than others. Thus, the FPGA firmware, as well as the ROC hardware, need to support this level of **scalability**. In order to support future MIMOSA generations

¹TRB – Trigger and Readout Board

as well, a **reconfigurable** FPGA firmware would facilitate research and development. Thus, given some basic input parameters, the firmware needs to be able to automatically reorganize its internal architecture upon synthesis. Following design parameters are derived from recent sensor specifications for that occasion: integration time t_{int} , number of output channels N_{Channels} , word size M and readout frequency f_{Readout} . Additionally, some sensor limitations should be provided to support the analysis, e.g. the maximum number of data packages per frame. Lastly, the data format specifiers need to be included, particularly the Header and the Trailer bit patterns.

The sensors operate in a constant data push mode. The number of useful packages per frame depends on the hit occupancy. In worst-case, some sensors will reach their full bandwidth. Thus, in order to avoid bottlenecks, the crucial point for the readout system is to provide enough resources to process the data faster than the sensors output. Although the sensors provide the packages serially, once they are reconstructed within the FPGA they can be processed package-wise. Therefore, there is a small time margin available for the FPGAs to pre-process each package, until the following package arrives. With MIMOSA-26 and MIMOSA-28, exactly 16 sensor clock cycles are available per package at full sensor occupancy. With MISTRAL and ASTRAL, the word size will change to 30 bit, thus allowing a small time margin of 30 sensor clock cycles between the packets. The **package frequency** f_{Packet} of a MIMOSA sensor is defined as:

$$f_{\text{Packet}} = \frac{f_{\text{Readout}}}{M}, \quad t_{\text{Packet}} = \frac{1}{f_{\text{Packet}}}$$

While f_{Readout} denotes the serial data rate, f_{Packet} defines the package rate, i.e. it is specified on the MIMOSA word-level. The inverse of f_{Packet} defines the **average processing time** t_{Packet} which can be invested by the FPGA to pre-process the data per received sensor package. The expected t_{Packet} values of the upcoming MIMOSA generations are given in table 5.1.

In practice, more than one package will be provided to the FPGA per unit of t_{Packet} , depending on the number of sensor output channels. It is feasible to develop an FPGA firmware which operates on all packets in parallel. In this case, t_{Packet} can be used as a design constraint to specify the readout performance. However, in order to determine pixel-clusters, the received packets need to be compared to each other, hence perfect parallelization is lost. In such case, a time of $t_{\text{Packet}}/N_{\text{Channels}}$ can serve as a reasonable firmware design parameter.

Due to beam fluctuations, the sensor output rate could temporarily rise by a factor of 2-3. Large **input buffers** are therefore necessary to store the packets until the hit occupancy decreases. Furthermore, the data pre-processing needs to be divided into several steps. As first, the data needs to be deserialized and the packages recovered within the FPGA. Secondly, the pixel hits need to be decoded. Afterwards, the packages should be analyzed for errors and the hits encoded into a more-applicable data structure, e.g. pixel clusters. The precise timing information needs

	MIMOSA-26	MIMOSA-28	MISTRAL/ASTRAL
t_{Packet}	200 ns	100 ns	100 ns
$t_{\text{Packet}}/N_{\text{Channels}}$	100 ns	50 ns	16.7 - 50 ns

Table 5.1: The average processing time which needs to be supported by the ROC FPGA per sensor package in order to avoid bottlenecks. MISTRAL and ASTRAL values are approximated. The first entry assumes parallelizable operation over sensor output channels.

to be provided at the cluster-level. Lastly, the data needs to be re-encoded into a suitable FLES format.

The global CBM clock needs to be recovered from optical links and the ROC needs to communicate with the global CBM data acquisition system. Therefore, the implementation of a suitable **network interface** is necessary. However, at the initial stages of this thesis, the CBMnet interface was under development and hence not fully specified. Thus, an alternative network interface has been applied based on TrbNet. The TrbNet network protocol is designed for large-scale detector readout systems, as well as prototyping stages. It is present in a fully tested and debugged manner supported by a large set of control and monitoring tools which were successfully applied within the HADES Experiment. More details can be found within the following section.

Additionally, it is desirable to operate the sensors **synchronously**, in order to facilitate the time stamping procedure. This can be realized by driving the sensors with a common master clock and by providing a common start signal to all sensor ladders simultaneously. The start signal delay of few ps within a ladder can be neglected. Since all sensors have the same integration time, their rolling shutter operation and frame readout will be perfectly synchronous. Lastly, the ROCs need to implement a JTAG interface for **sensor programming**, as well as sufficient monitoring and slow control procedures to detect errors, e.g. loss of synchronization.

5.1.2 TrbNet Network Protocol

The TrbNet [111, 112, 113] is a generic network protocol which was initially designed during the HADES detector data acquisition upgrade and now widely used in many other experiments. Its purpose is to create a scalable, fast and reliable readout network for particle detectors.

The data within TrbNet is encoded through six different layers [111]. The top layer, called *the application*, provides the raw data that needs to be transmitted. The data is then segmented in smaller and smaller parts as it passes through lower layers. At the end, the transmission is carried out using 16-bit packets via the bottom layer termed *the media interface*. One TrbNet message contains 80 bits, therefore usually five packets are required for the transmission. The media interface is hardware dependent. An optical link is often used as the medium, but in some cases LVDS buses are applied as well. The entire network implements a tree-like structure, where the leafs (TrbNet Endpoints) encode the data for transmission and the inner nodes (TrbNet Hubs) connect all the parts together, distribute the messages and collect the data.

All devices in the network have a unique ID. Due to the DHCP¹-like nature of the address allocation, they can be accessed independently, group-wise or all together. The TrbNet communication is performed over three virtual channels with different priorities. They all share one common media interface. Therefore, only one channel is active at a time. The channel with the lowest priority is the slow control and monitoring channel. Messages sent from here will be delayed whenever more important packet types need to be transmitted. Second priority is reserved for data messages. The extracted detector data is transmitted here. This channel has the highest payload in TrbNet. Originally, TrbNet was developed for triggered detector systems, thus the channel with the top priority is the trigger channel. Triggers in TrbNet do not claim much bandwidth, so they can be sent fast and the data taking and monitoring can proceed nearly undisturbed. There is only one trigger sent at a time in TrbNet and every trigger needs to be answered by all TrbNet Endpoints by sending a *busy release* message. The error and status bits of all Endpoints are hereby merged together. In a system that does not require triggers, e.g. the

¹DHCP – Dynamic Host Configuration Protocol

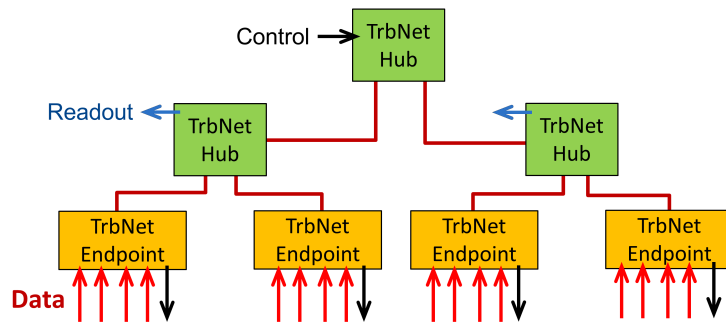


Figure 5.1: A generalized TrbNet-based network architecture. The central Hub is used to distribute control messages from one source and obtain the full network status in one place. Each Hub can merge and read out the data via Ethernet. The Endpoints are used to obtain the detector data and control the FEE. The communication within TrbNet takes place over bidirectional connections, e.g. realized via optical fibers, implementing three prioritized communication channels.

MVD, the trigger channel can be used to transmit other control messages of different types. A generalized TrbNet-based network is shown in Fig. 5.1.

TrbNet implements a diversity of security features. Handshakes are used wherever applicable. The data is CRC stamped and checked after each transmission allowing the detection of transmission errors. Furthermore, each Endpoint needs to answer every trigger message and every data-readout request. Thus failing components can be easily recognized via their unique IDs. The network is designed as deadlock free as possible and keeps running even when some components start failing. Further network errors can be excluded due to random-pattern matching. A random pattern is generated for each trigger and a corresponding readout request uses the same pattern to request the data for readout. Only if both patterns match at the Endpoint, the data is forwarded to the Hubs eliminating any possible event mix-ups.

5.2 FPGA Firmware Design and Implementation

After defining the basic requirements, an implementation of a generic FPGA module in VHDL¹ language is presented in this section.

5.2.1 General Overview

The current implementation gathers the sensor data, performs an error analysis, reformats the packages and transmits them via a TrbNet interface. In future, a cluster finding algorithm will reduce the data load and prepare the data for the FLES. Since all modules are implemented in the VHDL language, they are hardware independent and can be synthesized for any contemporary FPGA².

As first, the MIMOSA data is deserialized in order to reconstruct the sensor format. The data word size (e.g. 16- or 30-bit) is given as an input parameter. Frames are extracted by scanning

¹VHDL – Very high speed integrated circuits Hardware Description Language

²The code contains some FPGA dependent components, e.g. FIFO buffers and PLLs which need to be recreated when migrating to a different FPGA model. However, Virtex4- and ECP3-FPGA families are fully included.




Figure 5.2: The ROC is comprised of fully parallel readout chains. The network interface module needs to contain a full network protocol implementation (i.e. TrbNet or CBMnet).

the input stream for the Header package first, after which the frame information are gathered and several Finite State Machines (FSMs, see [52]) prepared for readout. The Trailer package marks the end of a frame, and with it the end of data processing. After a Trailer, all modules return to their IDLE state awaiting the next frame. The implementation is thus independent on the number of useful data packages and handles defective frames, as well¹. The entire procedure is encoded into a VHDL top-module named *the readout chain*. The FPGA can handle an arbitrary number of sensors by implementing more readout chains in parallel. Each readout chain is equipped with an independent *chain controller*. A simplified picture of the FPGA architecture is shown in Fig. 5.2.

The ROC firmware can be configured to support different sensor capabilities, including some newer MIMOSA generations. All the implemented design parameters are given in Appendix D. However, only the basic modules can be reconfigured with the given parameters. In Fig. 5.3, which presents a detailed readout chain layout, they are annotated with a (1). There are few other components which require a re-implementation for every MIMOSA generation (2), e.g. the data checker. Sensor limitations are varying between different MIMOSA types, hence the data checker can not be designed in a generic manner supporting all of the upcoming sensor options. Additionally, the applied network interface requires specific components (3) as well. Presently, only the TrbNet interface is fully implemented. The necessary CBMnet extensions are already planned and discussed in section 5.2.5.

Fig. 5.3 additionally depicts the data path through the readout chain. In following, all components are shortly introduced prior to their implementations in subsequent sections. Data is first transferred into the FPGA clock domain by the *input FIFO*. An implementation similar to an elastic buffer is realized to synchronize the data stream to the FPGA-internal clock. Then, the deserialization occurs within the *package generator*, which also detects incoming Header and Trailer packages. After that, data is written via the *data handler* to the *frame buffer* component, which is able to store several sensor frames, if necessary. Upon receiving a start signal from the chain controller, the *formatter* can pre-process and forward the data frame-wise to the network interface. Parallel to this basic readout procedure, the data is analyzed for errors and inconsistencies by the *data checker*. This provides sensor status information and allows the chain controller to initiate further readout steps.

¹There is only a finite number of checks possible in order to determine whether a frame is defective, or not. There are some errors as well, which can never be discovered. However, they usually occur with extremely low probability.




Figure 5.3: The ROC readout chain components can be partitioned into basic modules (1), sensor-specific modules (2), and network-specific modules (3). The free-streaming sensor data is stored in a buffer, which is read out frame-wise.

In case of high beam intensity fluctuations, a large data buffer can avoid data loss, storing several frames until the sensor occupancy decreases. There are two locations possible to place the frame buffer within the readout chain. As first, the raw data stream can be directly buffered at the input ports. However, this method does not allow any pre-processing operations on the data stream since the packages have not been reconstructed yet. Therefore, the frame buffer is placed after the package generator allowing early error-analysis and granting the ability to handle the data frame-wise, as explained in the following sections. Currently, the readout chain stores at least one full frame in the frame buffer, before it is read out. With this strategy, the entire frame can be analyzed for errors prior to starting the encoding process and further readout stages. If the frame contains too many errors, it could prove more efficient to discard it and free some resources, particularly during the high occupancy phase. Also, it allows placing the sensor status in front of the sensor data which is useful for performing various tests.

5.2.2 Basic Modules

The basic modules can be easily reconfigured to support different sensor parameters, as specified in Appendix D. The input stage of the readout chain consists of the input FIFO, the package generator and the data handler. All three stages are pipelined together according to Fig. 5.4.

The input FIFO is used to transfer the data bits from the sensor clock domain to the FPGA-internal clock domain. This procedure is referred to as clock domain crossing. In order to adapt the data to the new clock, an asynchronous dual-port FIFO buffer is implemented as a free IP core provided by the FPGA vendors Xilinx and Lattice. The implementation is similar to an elastic buffer. The data is buffered bitwise at the falling edge of the sensor clock. If the buffer is half full, it is read out. With this implementation, the FPGA clock needs to be faster than the sensor clock. However, for future, faster sensor versions DDR input flip-flops and shift registers are foreseen. Another important feature of the input FIFO is that it detects whether the sensor clock is still operational, via the *activity checker* sub-module. In case that the sensor clock stops ticking, an error is signaled directly to the chain controller.

The package generator deserializes the data stream recovering the sensor package structure. The sensor data bits are shifted into a register. There is one such register per sensor channel. The shift-register is checked and read out as one word, i.e. all bits are handled in parallel. Thus, it is

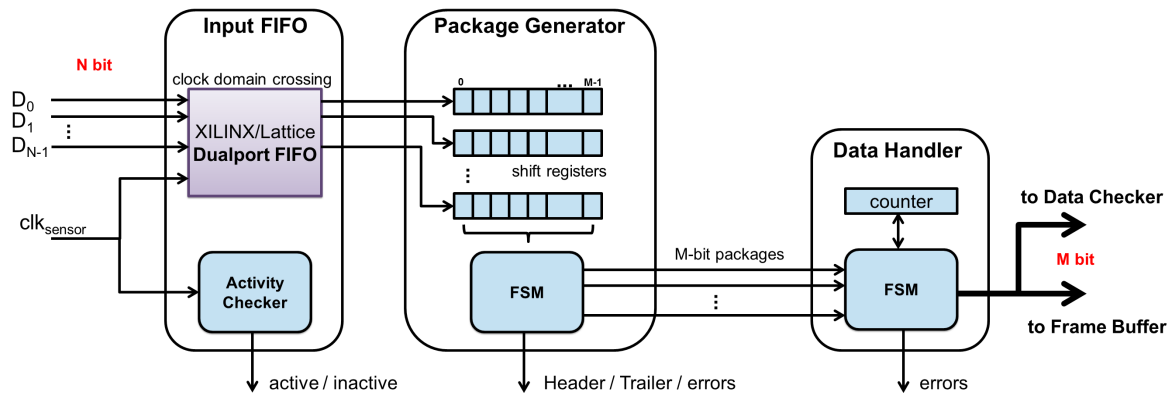


Figure 5.4: The input stage of the readout chain. Input data stream from N sensor channels is transferred into the FPGA clock domain by the input FIFO, deserialized based on the Header pattern by the package generator and then stored by the data handler. The final output contains M -bit words which are also forwarded to the data checker for a detailed analysis.

possible to wait until the Header package arrives and start the packaging process from then on. The internal sensor word-size M , e.g. 16-bit for MIMOSA-26, is well known and can be set as a design parameter. Every subsequent M shifts, another package is generated until the Trailer package arrives. Resets and errors are accounted for. All the decision logic is placed into an FSM.

The data handler is responsible for the appropriate frame buffering inside the ROC. The readout chain is designed to process the data frame-wise. If due to an error or a reset one frame is incomplete, the data handler ensures that even such Trailer-less frames are buffered correctly. This module has been simplified in order to reduce the system complexity and its sensitivity to errors. All data packages are counted and the total frame length is written at the end of the gathering process. Therefore, the number of packages for every frame is well known and the Header-Trailer determination can be entirely skipped, from here on. Again, the main control parts are placed into an FSM. Errors that signalize inappropriate sensor behavior are intercepted.

The frame buffer is used to store the data frame-wise. The large FIFO buffer provides additional storage space granting the FPGA sufficient time to pre-process the data and prepare for further readout steps. In order to keep track of where the frame starts and where it ends, an additional FIFO buffer is used for the frame lengths. Each frame length is 16 bits large. Thus the module uses two FIFOs, in total. The size of the data buffer can be selected with a design parameter. Up to 65 kB are currently supported, allowing the storage of nearly 29 full MIMOSA-26 frames corresponding to a time of 3.3 ms. All internal errors are accounted for. If the buffers are running full, the readout chain will automatically stop taking further data until the problem is resolved. The module is compatible to Xilinx Virtex4 and Lattice ECP3 FPGAs. The corresponding architecture can be specified with a design parameter.

5.2.3 Sensor Specific Modules

Currently, only **the data checker** requires a sensor specific implementation. Its purpose is to review every generated sensor package and check its integrity. The raw data packages are examined whether they contain a valid range, e.g. for the row and the column address of the pixel

hit, and if all bits are in place. Hence, most sensor and transmission errors can be detected early, by the FPGA. As a result, one 32-bit status word is created per frame containing all observable error and status messages. Current module version is optimized for MIMOSA-26, which could prove compatible to MIMOSA-28 after some minor modifications. Appendix E.1 summarizes all analyzed errors. Note that the absence of a status bit also implies an error. Therefore, a valid sensor frame always produces a static bit pattern which can be easily recognized by the user and subsequent analysis tools.

5.2.4 Network Specific Modules

Two particular modules need to be designed compatible to a given network interface. Only TrbNet is supported at the moment, however CBMnet will be included in near future as well. At the beginning of this thesis, TrbNet did not support free-streaming systems, hence a periodic *frame request* via the TrbNet trigger channel is required to read out the data properly. The frame request is a specialized message of a custom trigger type which is sent over the network to all ROC FPGAs from a central TrbNet controller. More details and the implementation of the controller are given in section 5.4. However, recently, an improved, synchronous TrbNet protocol has been developed as well. This thesis only implements the original version of TrbNet, thus the chain controller operates with frame requests.

The formatter is the last readout chain module before the data reaches the network. It is responsible for data formatting and subsequent transmission. The module is controlled entirely by the chain controller. The chain controller (see further in text) receives network messages and checks the sensor status and the number of stored frames in the buffer. Since a TrbNet message (e.g. the frame request) may block the entire readout network, it is always answered, regardless whether there is a full frame present in the buffer, or not. Hence, the formatter can be started in three modes: Normal, Empty and Null. The data transport to the network can have up to three stages, based on the mode:

- Mode **Null** does not provide any actual data to the network interface. It simply handles the corresponding TrbNet handshakes, if requested by the chain controller. This mode is used to answer the frame request, without sending any data. Furthermore, it can be used to delete a frame from the buffer and free some space, as explained further in text (see chain controller).
- Mode **Empty** provides only some basic information to the network constituting the ROC format header which is described in Fig. 5.5, variant A. However, only the sensor status, some debug information and the timestamp are sent, but no sensor data. The necessary TrbNet handshakes are exchanged as well.
- Mode **Normal** is the standard mode of operation. The formatter first transmits the format header, as in the 'empty' mode. After that, sensor data corresponding to the requested frame is sent from the frame buffer to the network interface. Some internal encoding algorithm or a cluster finder can additionally be integrated at this point to process the data.

Before sending the sensor data, the formatter may pre-process and re-encode the data stream. For testing purposes and laboratory set-ups, some meta information regarding the current status is placed in front of the sensor data as a header (variant A from Fig. 5.5). The start of the

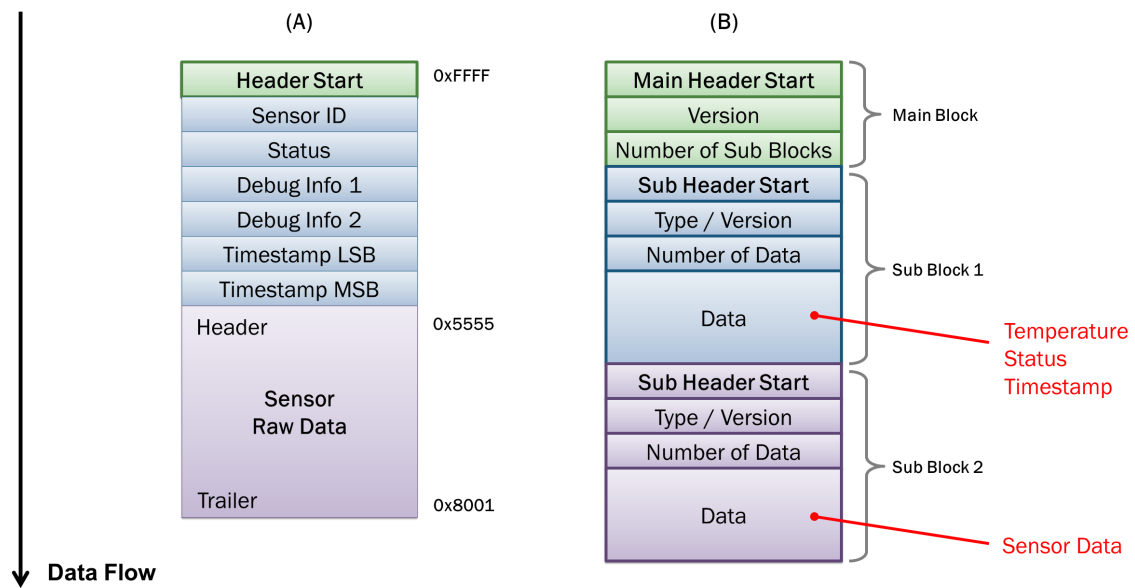


Figure 5.5: The static data format (A) and the planned dynamic format (B). Variant A simply places some meta information in front of the sensor data as a format header. The word 0xFFFF in hex marks the begin of the custom header. The actual sensor data, marked with 0x5555 and 0x8001 for Header and Trailer, respectively, is attached afterwards. Variant A does not contain any pre-processing algorithms at the moment. The more dynamic version B is planned as an extension supporting various sub blocks, versioning and reordering.

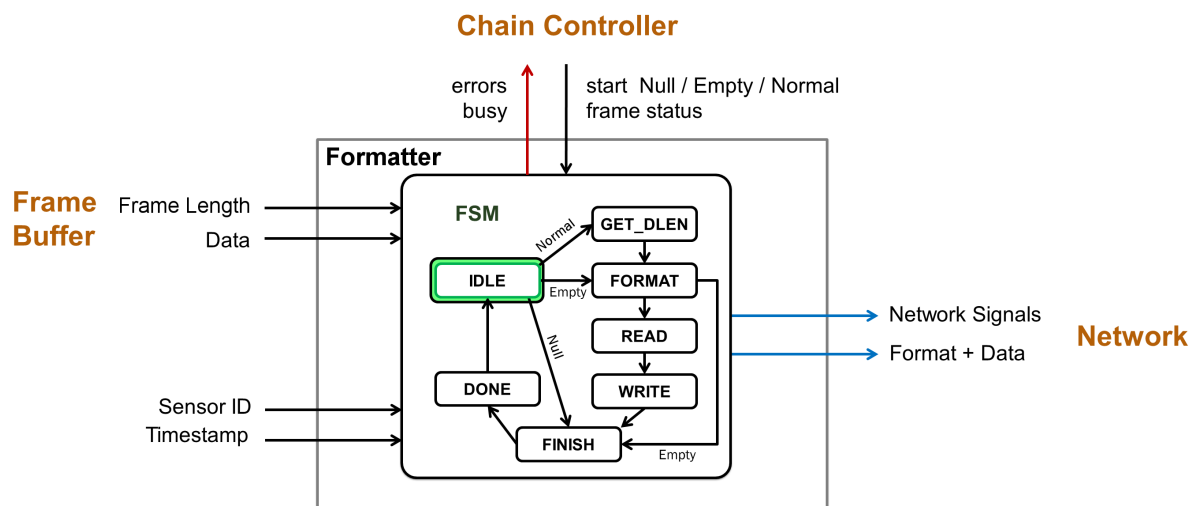


Figure 5.6: The formatter module provides some meta information (FORMAT), acquires frame data from the frame buffer (READ) and provides it to the readout network (WRITE). The chain controller sets the mode of operation (Normal, Empty, Null). Mode 'Empty' writes only the format header on demand, while mode 'Null' does not write any data to the network interface. A data pre-processing algorithm can be integrated into the FSM, e.g. between the READ and the WRITE state. The module supports only TrbNet at the moment.

customized header is denoted by a specific word, e.g. 0xFFFF in hex. Alternatively, a more dynamical reformatting of the data can be applied with the variant B from Fig. 5.5, where the data is organized within sub blocks. Each block has its own type and a different format following some predefined guidelines. The format is flexible, allowing for implementing new blocks and different block versions. However, the actual, time consuming pre-processing algorithms are not fully implemented yet. A related thesis currently focuses on the implementation of a cluster finder [114, 115] for the MVD. The module is compatible to the readout chain and can be easily integrated into the formatter. The simplified schematics of the formatter and its FSM are presented in Fig. 5.6.

The chain controller is the main component responsible for the dynamic data flow from the sensors to the DAQ system. It interfaces the network and nearly all of the readout chain modules. In order to support scalability, this component is integrated into every readout chain separately. Thus each readout chain can operate independently from the others, but requires a distinctive network address or ID. The component takes control of the data acquisition based on the network request and sensor status.

This module, as shown in Fig. 5.7, comprises several different parts. The frame counter keeps track of how many frames are stored in the frame buffer. Additional logic checks whether there are two consecutive Headers without the Trailer in between, in which case the sensor has been reset. An important feature is additionally provided by the timestamp generator which creates a timestamp whenever the Header arrives or the sensor stops operating. Therefore, each frame contains the exact arrival time of its Header or the time where the sensor got deactivated. The status generator gathers the relevant status and error messages from all components and integrates them into a 32-bit large FRAME_STATUS word, as presented in Appendix E.2. The

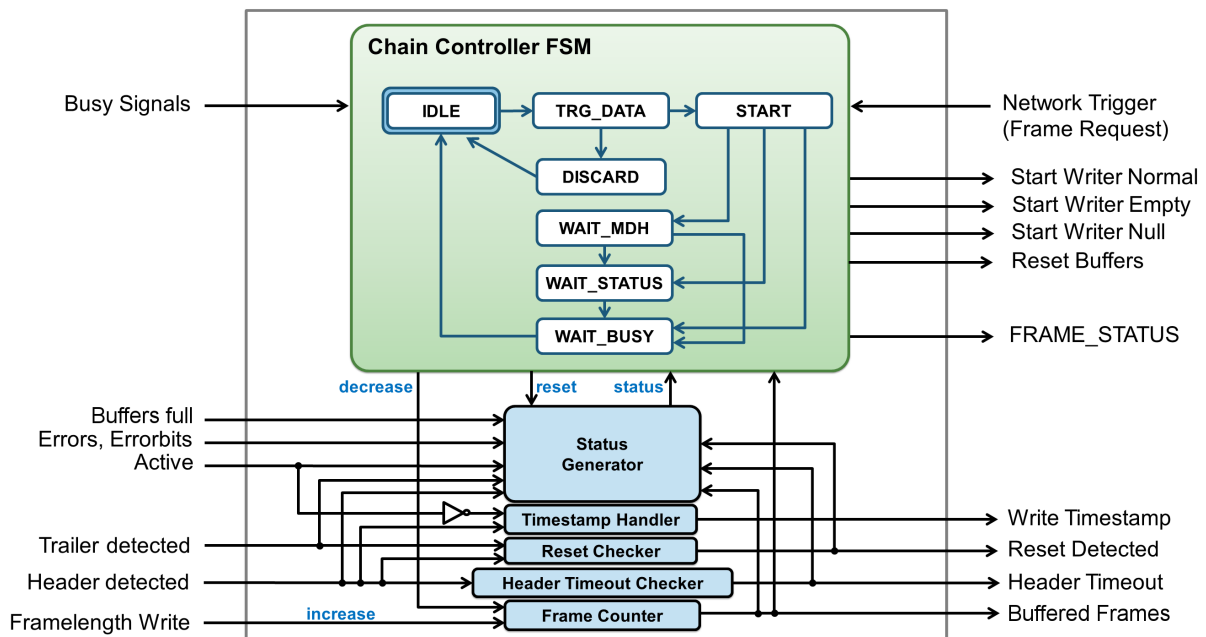


Figure 5.7: The chain controller architecture with TrbNet support. The key role of the module is the full control over the readout chain. Busy signals from all active components are checked and based on the current state of the sensors, the network and the readout chain itself, further course of action is determined.

data checker output is included as well.

The current FSM implementation is specifically prepared for TrbNet. The TrbNet needs to be configured to periodically request frames via the trigger channel (see section 5.4). Once a frame request is signalized, the chain controller decides what to send to the TrbNet: one frame, only the current status, or nothing. In case that one full frame is stored inside the frame buffer, the chain controller starts the formatter in 'Normal' mode (see formatter). Otherwise, the formatter is started in 'Empty' mode. However, if one of the TrbNet buffers¹ is running full signaling a network overload, a certain bit inside the frame request message is set and can be accessed by the chain controller. The bit signalizes that some buffer(s) in the network can not store an additional frame and that some of the MVD data is going to be lost. In such case, where a complete set of data can not be provided by the MVD, all ROCs stop the data acquisition until the affected Endpoint can store further frames. The formatter is started in the 'Null' mode, not outputting any data until the problem is resolved. Additionally, any frames stored by the buffer after the overload are discarded. As long as the bit is set, no additional frames are gathered by the ROCs. After the corresponding Endpoint(s) clear(s) the bit, the data acquisition continues as usual. The readout network provides coherent frames again, as long as the sensors keep running synchronously. The procedure therefore avoids incomplete or mixed data sets in overloaded-network scenarios.

5.2.5 Proposed Extensions for CBMnet

According to the modular readout chain design, only the network specific components require a CBMnet compatible re-implementation, i.e. the chain controller and the formatter. Moreover, their functionalities can mostly be left unchanged, e.g. the data formatting and the status generation. However, the CBM master clock needs to be included into the readout stages. The clock recovery over optical links and the implementation of a CBMnet FPGA module is addressed in a related work [116].

After a careful analysis, the modified readout chain is presented in Fig. 5.8. The changes are only outlined and not implemented yet. The left part including the input stage, the frame buffer and the data checker requires no modification, at all. The right part of the readout chain, however, has been redesigned. There is a new component planned, currently named 'data output', which ensures that the data is transmitted in a 16-bit format synchronous to the recovered CBM clock. The data output module should have full control over the DTM channel (see section 4.3.1)

¹The TrbNet Endpoint component has large buffers of its own, where it temporarily stores all the acquired frames.

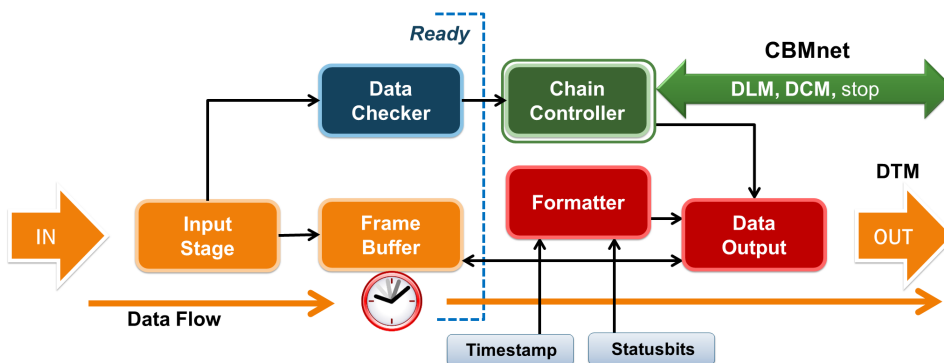


Figure 5.8: The proposed readout chain modifications for CBMnet.

and contain small transmission buffers for better flow control. An implementation of the cluster finder and the FLES encoding can also be realized within this module. In the latter case, the CBM time stamping mechanism needs to be supported as well. The formatter, on the other hand only provides the format information (see Fig. 5.5) on demand and can be bypassed when not needed. The sensor data from the frame buffer is now handled by the data output module.

The chain controller receives incoming CBMnet messages over a high-level interface, presented in Appendix F, and initiates DCM and DLM network transfers synchronous to the CBM clock. It also controls the data output module, checks its state, or resets the frame buffer if necessary. Proper slow control and monitoring procedures need to be included as well. To reduce complexity, the status and error bits could also be combined over a purely combinatorial process into a 32-bit status word which is accessed by the formatter and slow control at any time.

The main difference between the CBMnet- and the TrbNet-based approach is that the chain controller does not need to initiate the data readout for each frame any more. The start signals for the three cases Null, Empty, and Normal are not required. The data flow obtains free-streaming characteristics, where the chain controller merely decelerates the flow on network demand or on sensor failure.

5.3 Readout Controller Prototypes

Several hardware prototypes are used to validate the designed ROC firmware. The hardware is provided by the HADES collaboration. The ROCs are all integrated into multi-purpose boards used for data acquisition within the HADES readout network.

5.3.1 TRB2-Based ROC Prototypes

Within the scope of the HADES Experiment, a versatile Trigger and Readout Board (TRB) has been developed [117, 118, 119]. Besides its applications in all HADES subsystems, it is used for several other experiments and detector prototypes. As explained in [118], the TRB version 2 (TRB2) features following components:

- FPGA – Xilinx Virtex4 XC4VLX40
- CPU – Axis EtraxFS, 200 MHz, Ethernet and SDRAM support
- Optical Link, with TLK2501 SERDES, approx. 2 Gbit/s
- Ethernet connector, 100 Mbit
- SDRAM, 2×128 MB connected to the CPU and the FPGA each
- Flash ROM
- High Precision TDCs, $4 \times$ with 30 ps RMS precision
- LVDS add-on connectors, 15 Gbit/s
- Clock oscillator, 100 MHz

The integrated CPU runs a Linux 2.6 kernel loaded from the flash ROM at startup. One of the SDRAM modules is at the CPU's disposal as well as the Ethernet connector. Therefore, a fully operational embedded Linux system with Ethernet connection is available on-board.

The FPGA is the central TRB2 data processor performing all the high-speed data acquisition related tasks. Therefore, it is connected to all the other components: the optical link, the second SDRAM module, the TDCs and the LVDS connectors. The FPGA operates at a clock rate of 100 MHz provided by the oscillator. In order to gain direct access to the internal registers, the CPU is wired with 36 data lines directly to the FPGA. A photo of the TRB2 and its architecture is shown in Fig. 5.9. An important feature of the TRB2 board is that the fast LVDS connectors on the backside allow more add-on boards to be attached. The connectors provide necessary grounding and power lines, as well as 62 bidirectional data pairs.

One great advantage of the TRB2 board is that it is designed to form only the backbone of a large-scale readout network. It is intended to provide merely the necessary networking, slow control and monitoring abilities. The actual data acquisition is carried out by the add-on boards mounted on the back plane. HADES uses a multitude of such add-ons and there has been even one developed for MIMOSA sensors¹, in a related work [120]. The idea discussed there has been continued and applied to the first readout controller prototype.

The MAPS Add-On

The TRB2 add-on developed for MIMOSA sensors is named the MAPS add-on [120]. It has already contributed to a successful MVD Demonstrator in-beam test in the year 2009. The board was specifically designed to support the analog readout of some former MIMOSA generations, however the PCB is compatible to MIMOSA-26 as well. The main purpose of the add-on was to digitize the data and perform online correlated double sampling with a common mode filter. The MVD Demonstrator helped to qualitatively understand the crucial MIMOSA parameters and sensor characteristics. For example, the spatial resolution of $\leq 5.7 \mu\text{m}$ has been recovered from the data proving the suitability of MIMOSA chips for CBM [120, 121]. In addition to the TRB2 resources, the MAPS add-on contributes with more LVDS connectors, ADCs, SDRAMs, and another Virtex4 FPGA to the readout system. The board is shown in Fig. 5.10.

As an extension for digital sensor outputs, a $10 \times$ RJ45 socket PCB has been developed for the MAPS add-on. Therefore, commercial Ethernet cables can be used as direct connectors to the FPGA. The board is fundamental for the initial readout phases of the MVD and FEE development. Its interface to the sensors via RJ45 sockets can be used for JTAG programming and readout.

Due to the MAPS add-on attached to a TRB2, the first ROC prototype features two FPGAs. However, the input and output lines (I/Os) of the MAPS add-on can not handle more than three MIMOSA-26 sensors. Since there are sufficient resources on one FPGA to handle three readout chains, all the logic is moved to the TRB2 FPGA. The MAPS add-on FPGA serves merely as a data bridge and contains nearly no logic². The ROC firmware for MIMOSA-26 is implemented by setting the design parameters as shown in table D.1 in Appendix D. Besides the readout chains, the FPGA features a full TrbNet Endpoint with the connection to the optical link, as found in the

¹In particular, MIMOSA-18 and MIMOSA-20.

²A former version of the ROC did incorporate the readout chain on the MAPS add-on FPGA, but since the TrbNet related logic requires an optical link, as found on the TRB2, all logic has been moved to the TRB2 FPGA. This allows bypassing some FPGA-to-FPGA communication routines and eliminating additional sources of errors.

HADES setup. Slow control and monitoring can be performed by accessing certain Endpoint registers.

General Purpose Add-On

The MAPS add-on resources are strongly limited by the number of its input and output lines. Only two MIMOSA-26 sensors can be easily accessed, and a third sensor requires additional programming effort. Most of the resources are unused as well. Thus, it is possible to replace the MAPS add-on with a passive PCB providing more I/Os to the TRB2. Revisiting the HADES network, an ideal candidate has been discovered with the TRB2 General Purpose (GP) add-on. The board features 24 LVDS pairs (16 inputs, 8 outputs), 16 TTL lines, and a SCSI cable connector with 31 additional LVDS pairs. A photo of the GP add-on mounted on the back of a TRB2 board is shown in Fig. 5.11.

The board allows a simpler design without any additional FPGAs or further add-on PCBs. It merely contains several LVDS drivers for a TTL-to-LVDS conversion. The 16 LVDS input pairs are perfectly suited to read out four MIMOSA-26 sensors. Therefore, a ROC with four readout chains is integrated together with a TrbNet Endpoint into the FPGA. Again, parameters from table D.1, Appendix D, are used.

5.3.2 TRB3-Based ROC Prototype

A new TRB version has been recently developed at the GSI in Darmstadt, the TRB board version 3 (TRB3) [122, 123, 124]. It is mainly used as FPGA-driven time-to-digital converter, e.g. for time-of-flight and time-over-threshold measurements in HADES and FAIR experiments. Hereby, a time resolution of < 14 ps can be achieved [122, 124]. More importantly, the TRB3 can serve as a stand-alone digital processing and readout board as well. It features the following main components:

- $5 \times$ FPGA, Lattice ECP3-150EA
- Flash ROM for each FPGA
- $8 \times$ optical or Ethernet transceivers, up to ~ 3 Gbit/s
- Add-on connectors, 4×208 -pin, 4×80 -pin
- Hardware trigger interface

The logic density and the number of available block-RAMs of each FPGA is approximately four times larger than their Virtex4-based predecessor. One FPGA is centrally placed and interconnected to all other FPGAs serving as the central on-board controller, a data bridge, or a data extraction point (e.g. TrbNet Hub). The transceivers are also managed by the central FPGA. The peripheral FPGAs are situated symmetrically in the four corners of the PCB. The flash ROM of each FPGA can contain their bit-code allowing rapid programming at startup. Each of the connectors can be used to mount add-on boards on top, or below the TRB3. A photo of the TRB3 board and its internal architecture is presented in Fig. 5.12.

The readout chain firmware is compatible to the TRB3 FPGAs. It will be integrated into the peripheral FPGAs leaving the central FPGA as a data hub. The peripheral FPGAs are hard-wired




Figure 5.9: The TRB2 board and its internal architecture. Source: [118]




Figure 5.10: The MAPS add-on (blue cover) on top of the TRB2 board (green). The add-on architecture is displayed to the right. With the new add-on, conventional RJ45 Ethernet cables can be used to drive FPGA input and output signals. However, the number of I/Os is very limited and the ports are not bidirectional. Source: [120].




Figure 5.11: The General Purpose add-on attached to a TRB2. The LVDS lines are counted as pairs.




Figure 5.12: The next-generation ROC is based on the TRB3 board which features five FPGAs, eight data transceivers, and altogether ~ 20 times more FPGA resources than the TRB2. Each peripheral FPGA can be directly connected to an add-on board. A larger add-on board can be mounted on the back-side as well. Source: [122].

Board	Set-up	LUTs	block-RAMs
TRB2	TrbNet, $4 \times$ ROC	36 %	94 %
TRB3	TrbNet, JTAG	9 %	13 %
TRB3	TrbNet, JTAG, $4 \times$ ROC	16 %	29 %
TRB3	TrbNet, JTAG, $8 \times$ ROC	24 %	49 %
TRB3	TrbNet, JTAG, $12 \times$ ROC	31 %	70 %

Table 5.2: Comparison between the TRB2 and TRB3 FPGA resource consumption. Only one peripheral FPGA of the TRB3 board is applied. The TRB3 FPGAs additionally feature an integrated JTAG controller for sensor programming. Look-Up Tables (LUTs) represent a good measure of applied logic while the block-RAMs represent the number and sizes of applied FIFO buffers. Each additional MIMOSA-26 sensor requires linearly more resources. By decreasing the FIFO sizes, e.g. by changing some integrated VHDL parameters, the block-RAM consumption can be further optimized.

to the 208-pin connectors. Thus, additional TRB3 MAPS add-on boards are planned as a sensor interface. Up to twelve readout chains have been successfully integrated into one peripheral FPGA, allowing up to 48 MIMOSA-26 sensors to be handled by the TRB3, in total. However, due to bandwidth considerations a smaller number (≥ 16) could appear more practicable, in reality. The results of the FPGA resource analysis are presented in table 5.2. Further tests and the development of the MAPS add-on will be addressed in near future.

5.4 Development of a Readout Network Prototype

Due to the reasons explained at the beginning of the chapter, the readout network is based on the HADES DAQ with TrbNet as the network protocol. The initial ROC hardware uses TRB2 boards with the MAPS add-on. However, it is applied only to demonstrate the suitability of the readout system to program and operate the sensors via TrbNet. Laboratory tests with up to 4 sensors

and a radioactive source have verified the most principal design choices. Afterwards, due to the limited resources, the ROC hardware is replaced by the TRB2 board with a GP add-on. In this way, a more realistic, larger network setup can be created which is described within this section¹. The ROC, however, plays only a small role in the entire process. The remaining parts, which are as vital for the readout as the ROC, are constituted by: FEE, JTAG controller, TrbNet controller and TrbNet Hubs.

5.4.1 General Network Operation

As first, a bottom-up network description is presented in the following. The network connects several MIMOSA-26 sensors via Flex-Print Cables (FPCs) to specialized FEE. The FEE contain RJ45 sockets for a simpler data transmission via conventional Ethernet cables. The data reaches the ROCs which are based on the TRB2 board with the GP add-on. However, the ROCs do not contain RJ45 sockets and require therefore patch panels to interface the sensors. The data is gathered in ROCs within their TrbNet Endpoint modules, which prepare the data for the high-speed transmission over optical links. The TrbNet Hubs read out the ROCs, aggregate the data and then transmit it via Ethernet for storage and analysis. All components of the TrbNet section use bidirectional, high-speed optical links for communication over the three TrbNet channels. In order to control the data flow throughout the network, a central TrbNet controller is applied. The controller is named Central Control Unit (CCU) in analogy to the HADES Central Triggering System (CTS, see [113]). Since the original version of TrbNet is applied in this network, such a controller is necessary for correct data flow, as explained further in text. In addition to the readout, the sensors also need to be programmed. Therefore, one additional TRB2 board with a GP add-on is used as the main JTAG controller. The board is termed MVD Acquisition and Interaction Node (MAIN) and contains the JTAG controller implementation, as well as the CCU. The full network scheme is presented in Fig. 5.13. The entire network contains various different PCBs for the ROCs, Hubs, and the FEE which are all interconnected. Therefore, a non-trivial grounding strategy is required to avoid ground loops which may decrease the overall performance. More details are presented in Appendix H.

After the network description, the network operation can be described, as follows. The applied TrbNet protocol is originally developed for detectors relying on low-level triggers. For example, the HADES network distributes triggers to denote which events are of physical interest and need to be stored. In case of the MVD, all data needs to be gathered, thus a trigger is not required. However, a central trigger controller (e.g. the HADES CTS) is an integral part of the TrbNet version applied here and can not be avoided without modifying the TrbNet code. Therefore, periodic **frame requests** are sent by the CCU over the TrbNet trigger channel. After at least one frame request is processed by the ROCs and the reply from all TrbNet Endpoints arrived, the CCU requests the data for readout. This is done by sending a **readout request** over the TrbNet data channel to all the ROC Endpoints, which subsequently forward one frame from their Endpoint buffers to the Hubs. The network is conceptualized with a large focus on global **synchronization**. All ROCs process the same frame request and, afterwards, the same readout request. Since the integration time of MIMOSA-26 is fixed, by starting the sensors simultaneously and by using one common clock for their operation they will continue to operate perfectly synchronized. All counters and sensor clocks in the network are generated from a single, jitterless source, namely

¹ROC tests on a TRB3 board are currently in preparation and will be available in near future.

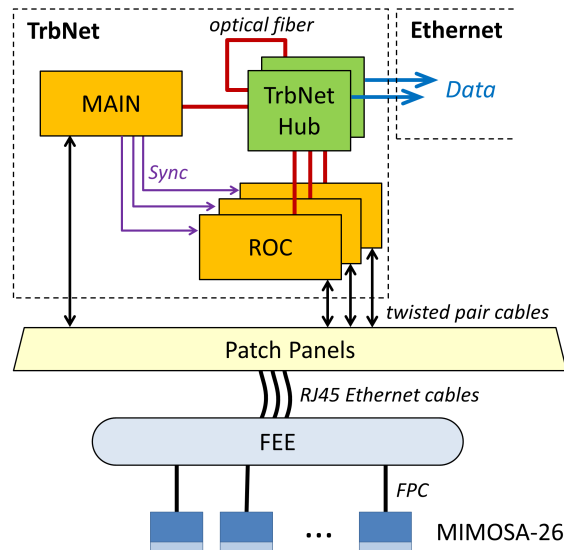


Figure 5.13: The readout network prototype is based on TrbNet. All TrbNet components are connected via optical fibers. The MAIN board acts as a central network controller and additionally programs the sensors. The MAIN and the ROC PCBs are realized on a TRB2 board with a GP add-on. Therefore, they interface the FEE over twisted-pair cables attached to patch panels. Current implementation supports up to 12 MIMOSA-26 sensors, however this can be increased by adding more FEE and ROCs. If the readout bandwidth is exceeded, more Hubs can be added as well.

the clock oscillator of one of the TrbNet Hubs. The clock is distributed to all other components via optical links. The applied TrbNet implementation does not support synchronized message transport, thus additional effort is required to synchronize the ROCs. This is resolved by applying additional LVDS lines carrying synchronization signals from the CCU to the ROCs, as indicated in Fig. 5.13. Naturally, such a network requires modules to detect synchronization and algorithms on how to proceed if the synchronization breaks down, which are currently under development.

5.4.2 MIMOSA-26 Front-End Electronics

As a first step towards the sensor readout, four front-end PCBs are designed to power-up the sensors and provide an adequate JTAG and readout support [120]. They are presented in Fig. 5.14. The JTAG chain is formed by the JTAG queue board (QB), one or more front-end boards (FEB) and one termination board (TB), at the end. The QB is equipped with RJ45 sockets to receive the TDI, TMS and TCK signals, and to send the TDO back to the main controller (see section 3.2.1). More details on the entire JTAG operation can be found in [84]. Apart from the JTAG interface, there are additional MIMOSA-26 specific signal paths present for the *Start*, *Reset* and the 80 MHz *Clock* (SRC) signals transmitted over a RJ45 socket as well. Therefore, the QB handles six sensor input signals and one output.

The signals are transmitted via flex-print cables further to the FEB. Currently, only sensor ladders with two MIMOSA-26 sensors are supported (see Fig. 5.14). The FEB, which is connected directly to a ladder, shows therefore a similar geometry. The board conjoins all the power lines, JTAG and sensor signals on one small two-layer PCB. The final FEB will be placed nearest to the beam pipe and needs therefore to be designed without any active components¹ due to the

¹Currently, the FEB contains only resistors, capacitors and copper lines.




Figure 5.14: The initial FEE prototypes are designed to support the conventional RJ45 Ethernet cables. Four sensors, organized in two ladders, are operated in this example. In principle, arbitrary many sensors can be connected to one large JTAG chain, but each additional FEB increases the signal noise levels on the thin and relatively unshielded flex-print cables.

large expected radiation dose. After the first FEB, the system can concatenate arbitrary many additional FEBs. However, with each further board and with each FPC, the noise in the cables is growing larger. A detailed study on this subject is given in [84] as well. The end of the JTAG chain requires appropriate termination of LVDS signals via the TB.

Sensors are dissipating their power over converter boards (CBs), which also read out the data, measure the temperature and can be used to switch off damaged sensors in the JTAG chain. If a sensor gets damaged, its TDO output could be malfunctioning and the JTAG chain interrupted. All the subsequent sensors would not get programmed correctly. Thus, an LVDS input on the CB has been reserved to receive a signal from the ROC and activate or deactivate a switch allowing the chain to bypass the TDO signal of any sensor in the ladder and reconnecting the non-defective TDI signal instead. The CB uses RJ45 sockets to obtain the switch signals and transmit the raw sensor data. One RJ45 connector is used per sensor and one more for the switches.

As an add-on to the CB, a latch-up protection board has been developed [37]. The latch-up, equivalent to a short circuit, can be detected due to increased power dissipation. In such case, the latch-up protection sets in within 2 μ s and disrupts the power to the sensor preventing it from taking damage. The latch-up sensitivity level can be tuned manually via four potentiometers (two per sensor for digital and analog voltages). The pin assignment of QB and CB is presented in Appendix G. Due to the RJ45 sockets, conventional Ethernet cables can be used. However, certain pins have been inverted to minimize the noise and improve the data quality. For details, refer to Appendix G.

5.4.3 JTAG Controller Implementation

The QB provides merely the appropriate sensor interface, however an implementation of the actual JTAG controller functionalities to program and steer the sensors is additionally required. Within the scope of a related thesis [84], such a JTAG controller firmware for FPGAs has been

developed in VHDL. The code is initially tested on the TRB2 board with a MAPS add-on. A TrbNet Endpoint is integrated into the controller allowing the access via slow control.

After the implementation on the MAPS add-on, the component has been migrated to the GP add-on setup (MAIN board). With the increased number of I/Os, it can handle up to three synchronous JTAG chains. For this occasion, the SCSI output of the GP add-on has been customized and connected to a patch panel in order to support the RJ45-based QB connectors. TRB3 tests have been successfully performed as well.

5.4.4 Central Control Unit Implementation

Current network implementation requires a central controller to steer the readout. Such a component, the CCU, is implemented in VHDL within the scope of this thesis and integrated into the MAIN board. The component additionally performs global network synchronization as well. The sensors can be started synchronously, however the ROCs need to be synchronized at runtime. For this occasion, equally long LVDS lines are used to connect the CCU with the ROCs and transmit two specialized signals. The *ROC_START* signal resets the ROC timestamp counters and starts the readout chains altogether. Since the FPGAs operate over their internal clock oscillators which, in general, differ by few parts per million, they need to be re-synchronized periodically. Therefore, the *ROC_SYNC* signal can be applied to re-align the timestamp counters. Both signals share one line and can be distinguished based on the pulse length. The layout of the CCU FPGA (without the JTAG controller part) is shown in Fig. 5.15, as well as its connection to the ROCs. All CCU components are described in following:

- **HADES TrbNet CTS wrapper:** The top-most wrapper module of the HADES central triggering system is integrated within the CCU for data flow control. The CTS, as used by HADES, detects whether there are some physical or calibration triggers present and then sends a TrbNet trigger message with top priority to all the Endpoints. The wrapper module can therefore be applied for the MVD readout network prototypes to transmit messages on the TrbNet trigger and data channels. The CCU frame request is sent periodically over the trigger channel with top priority, while data channel is used for readout. After initiating a frame request, the module is busy awaiting the response from all ROC Endpoints. Once the TrbNet Endpoints have finished their data buffering, the busy release signal is communicated over the trigger channel and the module exits its busy state allowing the transmission of another frame request. With every frame request, an 8-bit random code is sent as well. It is important that the data readout request for every frame matches the random code sent by the corresponding frame request. Therefore, a FIFO buffer is applied in order to keep track of all the random codes. All frame requests are implemented as TrbNet triggers of type 8 (reference-time-less triggers, see [111]).
- **CCU Sync:** After starting the component via slow control, the module will send a 160 ns long *ROC_START* pulse over its LVDS lines to activate and synchronize the ROCs. Afterwards, it starts counting up to a preset duration, e.g. 11520 clock cycles equivalent to one MIMOSA-26 frame, and sends 90 ns long *ROC_SYNC* pulses in a loop.
- **CCU Frame Request:** This module is responsible for the appropriate trigger channel handling. After a start signal has been provided by the CCU Controller (see below), it will generate a random code and initiate exactly one frame request on the trigger channel.




Figure 5.15: The CCU component and its connection to the ROCs. While the data flow control is performed over TrbNet, the time critical synchronization signals require specific LVDS lines.

- **CCU Readout Request:** This module takes one random code from the FIFO buffer and generates all appropriate data channel signals to send exactly one readout request. The readout request initiates the readout of one frame, per sensor, in all ROC Endpoints which then forward their data to the Hubs for extraction via Ethernet.
- **CCU Controller:** In order to control the correct sequence of events and to keep an overview of the current status, this component controls the CCU Frame Request and Readout Request components. After the CCU Sync component sends its first *ROC_START* signal, the controller starts to operate by initiating frame requests periodically. The following request is sent after the previous has been answered. The component keeps track of how many frame requests are answered in a counter. As long as the counter is ≥ 1 , the Readout Request module is started to read out the data frame-wise. Usually, the time to store a sensor frame in the ROC is greater than the time to read it out, since MIMOSA-26 sensors have a relatively large integration time. Therefore, the frame counter always stays at an acceptable level.

5.4.5 TrbNet Hubs

The TrbNet Hubs are used to interconnect all the network nodes, however they require specific hardware designed to support constant data flow on numerous input channels. The HADES detector uses customized Hubs equipped with two FPGAs and 20 optical link sockets designed as TRB2 add-on PCBs. Up to 12 links can be used for data aggregation from the TrbNet Endpoints, per Hub. The Hub firmware features data merging algorithms and implements data extraction protocols, e.g. UDP¹. Hence, a dedicated link is used to transmit the merged data packets via Ethernet. The peak extraction rate amounts to 50 MB/s per Hub, at the moment. However, more Hubs can be easily integrated into the setup, as shown in Fig. 5.16.

In the current network setup, the TRB2 boards with their corresponding add-ons can be organized within two actively air-cooled stacks (see Fig. 5.16). Up to two Hubs are currently applied increasing the total output bandwidth to 100 MB/s. Each Hub features a clock oscillator with a 100 MHz frequency obtained from the TRB2 board which it is attached to. One of these oscillators is chosen to drive all the timestamp counters and sensor input signals in order to synchronize the network as much as possible.

¹UDP (User Datagram Protocol) is a minimal network protocol used to stream data over Ethernet to a certain computer port.




Figure 5.16: Two stacks of TRB2 boards with one Hub (left-hand side). The setup contains the entire TrbNet network. In order to use two Hubs for readout, they need to be connected according to the scheme to the right. Additionally, the Hubs need to be configured properly via slow control.

5.4.6 Laboratory Tests and Results

The initial test stage [37] consists of a small-size network setup with ROCs and JTAG controller implemented on TRB2 boards with MAPS add-ons. After the sensors were bonded and all the FEE tests concluded, the sensors were programmed and configured via their JTAG interface. Only one ladder featuring two MIMOSA-26 sensors was used. The sensor outputs were first monitored with a logic analyzer. Unless the sensor registers get properly programmed and the sensors receive a start signal, they will not start operating. However, the correct data format header has been observed with the logic analyzer demonstrating correct system behavior of the FEE and the JTAG controller, as shown in Fig. 5.17. The MKD signal marking the start of the frame shows a length of four clock cycles, as denoted in the documentation. The Header pattern was preset to $0x5555$ via JTAG, which is $1010\dots10$ in binary. The data length of 570 indicates a full frame. Hence, the sensor is programmed and successfully started. In addition, the Header packages of both sensors were arriving at the same time, within the logic analyzer time resolution of 2 ns, which indicates an appropriate sensor synchronization by the JTAG controller.

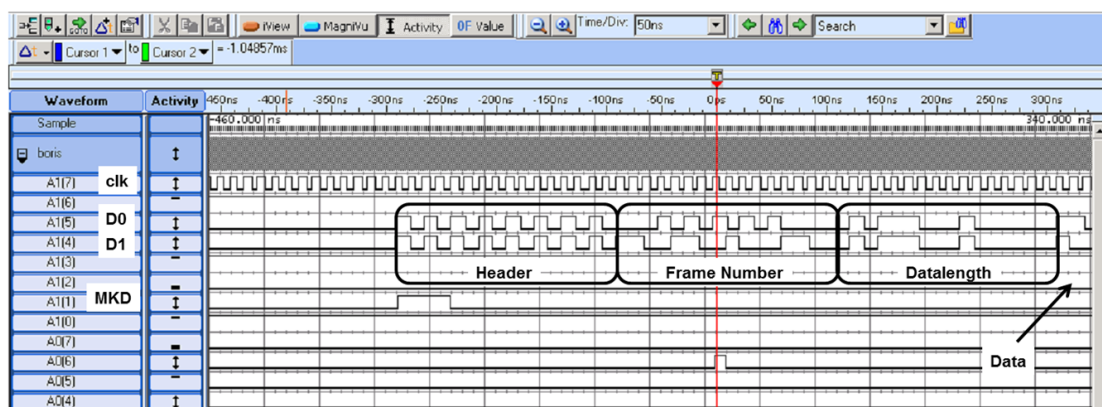


Figure 5.17: The first MIMOSA-26 output indicating correct behavior of the FEE and the JTAG controller. All of the 16-bit packages per channel constituting the MIMOSA-26 header have been recognized.


After a suitable sensor configuration, some tests were conducted to validate the correct ROC operation as well. The event builder from the HADES setup was used to store the TrbNet Hub data on a local computer. An extensive data analysis has shown that the raw sensor data was recovered without errors. In order to confirm this statement and to demonstrate correct sensor behavior, an ^{55}Fe X-ray source has been used for further test procedures. The results are presented in Fig. 5.18. The source was mounted on top of one sensor covering only half of its matrix, and placed into a darkroom. The source hot spot could be unambiguously determined from the visualized sensor data which were obtained via the readout network. Therefore, the presented initial network tests have validated the suitability of the readout network to handle and read out MIMOSA-26 sensors.

The readout chains could obtain all data packets without loss. The sensor reset signal was detected as well. However, no actual sensor data pre-processing algorithms, e.g. a cluster finder, were performed within the ROC FPGAs. The frame was buffered in the frame buffer and after a readout request forwarded to the TrbNet Endpoint. Merely some customized format information were placed ahead of the raw data. With the 100 MHz FPGA clock frequency, a full MIMOSA-26 frame containing 1148 sensor packages could be handled within 12 μs . Therefore, the network had sufficient time to store the frames and distribute all TrbNet messages which have $\approx 3 \mu\text{s}$ network latency. However, with additional data processing routines, a faster sensor integration time and high sensor occupancy, the readout procedure could reach its limit. This needs to be addressed, in future, with the newest clustering algorithms.

After the initial tests, the readout network was upgraded with GP add-ons. Three ROCs, each equipped with four readout chains, were applied. In total, 12 sensors were used for the tests. The CCU and JTAG controller have been combined together on the MAIN board FPGA. The sensors were arranged in three synchronous JTAG chains. Four sensors per chain were controlled by the MAIN board. Since all the 16 LVDS input ports of the ROCs were occupied by the sensors, the *ROC_START* and *ROC_SYNC* signals needed to be transmitted over the TTL ports.

The readout network prototype has passed all the elementary laboratory tests. Data from all sensors could be acquired at any time. The sensors were synchronized and remained synchronous during a two-day stability test. No errors could be produced by additional stress tests. The JTAG chains did not yield any errors, with and without an attached termination board. In order to demonstrate the JTAG chain performance, a large chain with eight sensors was created. Two of the sensors were damaged and have been bypassed with JTAG switches. Four FPCs with a total length of $\approx 2 \text{ m}$ were used to connect the four sensor ladders via FEBs. The large chain, presented in Fig. 5.19, has passed all basic tests without errors. A radioactive β -source was applied to the test setup, with the results presented in Fig. 5.20.

The only negative result obtained during the tests was the inability to synchronize the ROCs over the TTL lines. In case of LVDS lines, the signal voltage level is, in principle, obtained by forming the difference of the voltage levels between the two lines. Therefore, the resulting voltage is independent of common-noise effects and exhibits generally less noise. The TTL cables are single-ended and do not carry a ground line. Hence, they are prone to noise effects. A relatively large common-noise has been measured in the presented network setup, thus the TTL signals have their voltage levels shifted by a considerable amount over time. Therefore, the synchronization signals from the CCU to the ROCs are not transmitted simultaneously resulting in inaccurate timestamps. More details on this issue are shown in section 6.4.3 in the following chapter, which describes an in-beam test of the presented network.



(a) Integrated hits in the sensor matrix of 576×1152 pixels. The source hot spot and the round shape of the source can be clearly identified.

(b) Individual pixel clusters can be found in all analyzed frames.

Figure 5.18: The data obtained from one MIMOSA-26 sensor irradiated with an X-ray source.


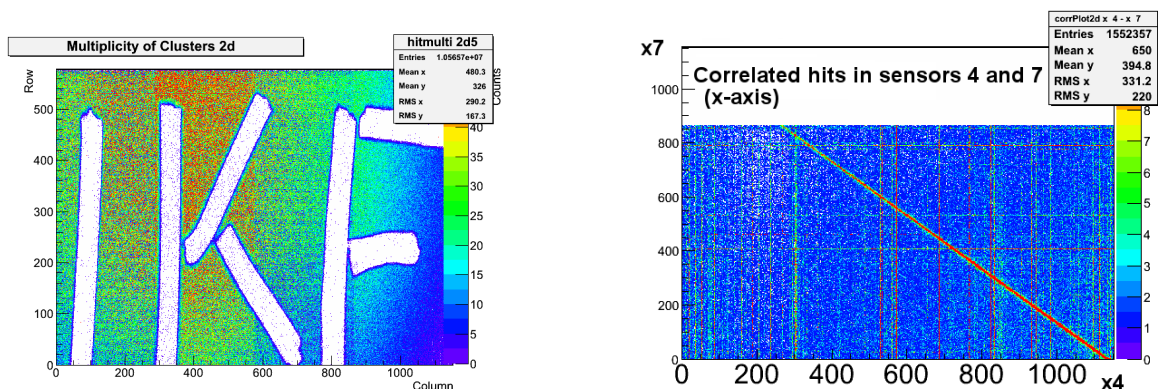


Figure 5.19: The large JTAG chain was formed by one QB, four CBs, four FEBs and eight MIMOSA-26 sensors (enumerated), two of which were placed in a darkroom. The length of FPCs forming the chain was ≈ 2 m, in total.



(a) An IKF logo stick-figure has been placed in front of one sensor. As a result, the covered regions were exposed to less β -particles.

(b) The diagonal line is formed by particle hits detected by both sensors simultaneously and can be used for alignment.

Figure 5.20: The main results obtained with the fine-tuned system exposed to a radioactive β -source. Sensors 4 and 7 were placed in front of each other, as a dual-layer module. One of the sensors was rotated with respect to the other and glued to the back side of the same mechanical support.

Chapter 6

In-Beam Tests and Results

The implemented readout controller and the readout network from section 5.4 have been tested under realistic conditions as part of an experimental setup at the CERN SPS secondary test beam facility. In total, 12 MIMOSA-26 sensors were studied within the scope of the first MVD prototype. This chapter describes the in-beam test and presents all of the data acquisition related results.

Section 6.1 presents the first MVD prototype module, the in-beam test setup, as well as the DAQ architecture. Furthermore, section 6.2 introduces the applied software toolkit forming a small-size detector control system. The slow control and monitoring tools are included as well, followed by a comprehensive list of all implemented security features. The performed measurements are discussed subsequently, in section 6.3. Afterwards, all of the data acquisition related results are presented in section 6.4. In the end, section 6.5 provides a general conclusion concerning the present and future MVD readout system.

6.1 Test Setup at the CERN SPS

The first MVD prototype module, corresponding to a fraction of an MVD station quadrant, has been developed with the goal to prove the principle and perform tracking-related studies with high-energetic sub-atomic particles.

6.1.1 MVD Prototype Core Module

The proposed MVD station presented in section 4.1 comprises double-sided MIMOSA sensor modules which are arranged into ladders and can be operated independently, yet synchronously using a global clock and a common start signal. In order to prove the principle and to determine crucial system parameters, one small-sized prototype module is developed at the University of Frankfurt [125]. Two MIMOSA-26AHR sensor pairs, thinned down to 50 μm thickness, are glued back-to-back on actively cooled, 200 μm thick CVD diamond support structure, as shown in Fig. 6.1. The thickness of the entire core module is within the CBM material budget specifications of $0.3 X_0$ for the first MVD station. Studies concerning mechanical integration, heat dissipation, long-term stability, readout mechanism and vacuum operation of such modules are




Figure 6.1: The image of the MVD prototype module. The schematics of the module (cross-section, side view) is presented to the right. The bonding wires have been additionally encapsulated for their protection. Source: [125, 126].




Figure 6.2: The beam telescope layout with a cross-section through one of the reference planes. One sensor ladder is integrated into each reference station, whereas the prototype contains two. One FEB per ladder is situated inside the casing. The heat-sink is cooled with silicon oil. Additionally, the humidity inside the casing is reduced with nitrogen gas. Source: [125, 126].

currently being addressed. The core module is integrated into a five-station setup and tested with a pion beam, as explained in the following section.

6.1.2 Beam Telescope Setup

The in-beam experiment [125] performed at the CERN SPS secondary test beam facility (target 4, hall 6) focuses on the newly developed MVD prototype core module. In total, 12 sensors organized within five stations are applied, as outlined in Fig. 6.2. Four reference stations, each equipped with two sensors, accompany the prototype module in an arrangement where two of the stations are placed in front of the prototype and two behind in order to detect and track beam particles. The sensors are grouped into ladders. Two MIMOSA-26AHR sensors are glued to a common support structure (either aluminum or CVD diamond) and read out using one FEB per such ladder. Hereby, six sensors in a row serve as a beam telescope, while the other six are intended as spare parts. High energetic π^- particles are tracked through the telescope allowing the determination of the prototype detection efficiency and its spatial resolution based on tem-

perature, beam intensity and different threshold settings.

The main task from the DAQ perspective is to analyze the readout network performance and long-term stability under realistic work conditions in an irradiated environment where sensors and hardware components are subject to single-event effects (e.g. SEU and SEFI, see section 2.3.4), bias voltage fluctuations and latch-ups (see section 2.1.3). The beam intensity can be varied as well, allowing performance-related DAQ studies. The measured spill duration amounts to 9.5 seconds during the test, with a spill period of 45 seconds. Two scintillators are used to detect the particle fluxes, one in front and one behind the telescope setup. They operate in the coincidence mode where hits in one scintillator are correlated in realtime with hits in the other. The coincidence rate is directly proportional to the effective beam intensity which the sensors are exposed to. They are thus used to detect the beam spot. With this method an overall rate of 8 000 - 380 000 coincidences per spill is measured during the experimental runs.

The readout hardware is based on the network prototype from section 5.4. All sensors are organized within three JTAG chains with four sensors per chain. Chains 1 and 3 are used to operate the reference planes while chain 2 is used for the prototype. All chains are synchronized from the start, but operated individually via slow control.

6.1.3 Power Distribution

In order to provide sufficient power to the sensors, two GW Instek PSP - 405 power supplies are set up, one for the reference planes and one for the prototype. The prototype, together with its FEE, consumes unprogrammed 1.8 W and programmed 6 W¹. Due to the symmetric setup, the remaining two sensor chains dissipate the same amount. The sensors are protected against overvoltage with a crowbar circuit [127]. An additional power supply is set up for the TRB2 boards with a 48 V source. The add-ons, e.g. the Hubs, are powered over the TRB2. The TRB2 power supply outlet has been equipped with a remotely controllable on/off switch.

6.1.4 Computer Network

The entire setup is controlled and monitored remotely via Ethernet. One conventional Ethernet cable of 30 m length connects the main network switch in the beam cave with the data acquisition computer (DAQ-PC) situated approx. 20 m away in the counting house. The switch supports Gigabit Ethernet and contains some optical transceivers to obtain the TrbNet data. Each Hub is able to stream the data in a UDP format over a certain optical fiber which is connected to the switch. Therefore, the DAQ-PC can listen to the designated port and receive all the UDP packets. The Gigabit switch also connects some other network switches which are used exclusively for slow control and monitoring.

The DAQ-PC plays a major part during the beam test. The computer has large resources at its disposal: i7 Quad-Core CPU, 32 GB of RAM, 12 TB storage space and two independent Gigabit Ethernet controllers. It performs all the slow control related tasks, stores the data and connects two network domains together. Within the first domain, the DAQ-PC performs slow control, monitoring and data acquisition of the beam cave setup via the Gigabit switch. The second

¹The measurement involves nominal fake hit rates and no beam. Based on sensor occupancy these numbers can vary.

network domain is situated in the counting house and contains additional computers which are part of an internal NFS¹ domain.

The computers used for analysis have immediate access to the data due to the NFS setup. The servers in the network slowly backup the data files in the background without disturbing the main tasks. Hard drive errors are accounted for by mirroring the hard drives in a RAID1² configuration. One additional DAQ-PC with less resources is kept as a backup.

The DAQ-PC requires full TrbNet access. All TRB2 boards are connected to one predetermined Ethernet switch in the beam cave allowing FPGA programming, slow control and monitoring via the DAQ-PC. A designated TRB2 slow control board is used to acquire low-level Ethernet calls from the DAQ-PC, forward all messages further to TrbNet and transmit the results back to the counting house. In addition, a Raspberry Pi mini-PC remotely controls all power supplies. It can power off the TRB2 power outlet in case of emergency and controls the sensor power outlets with individual PERL³ scripts. A further Ethernet switch provides access to a remote laptop computer that is set up to control the telescope cooling system and to monitor several webcams situated in the beam cave.

6.1.5 Summary

Sensors are organized into three JTAG chains with four sensors per chain. They are programmed and controlled over the MAIN board and exhibit therefore nearly perfect synchronization⁴. Each chain is handled by one ROC according to the readout network from section 5.4. The ROCs are connected to one or, optionally, two Hubs via their optical links. The network applies the TrbNet protocol for data transmission. The TRB2 boards additionally rely on their Ethernet connection for slow control and monitoring. One board is hereby acting as the slow control gateway to TrbNet. The Gigabit switch communicates over a 30 m long Ethernet cable with the DAQ-PC, which is used for data acquisition, slow control and monitoring. The scheme of the entire setup is summarized in Fig. 6.3.

6.2 Applied Software Toolkit

Controlling a complex network of various different components demands specific low level tools. The programming and slow-control of TRB2 boards and Hubs is carried out using scripts developed in the scope of the HADES experiment. The Hub additionally merges frames from different sensors which are correlated in time into one event.

6.2.1 Event Builder

The event building relies on tools provided by the HADES data acquisition system. One application listens to a designated computer port for UDP packets, while another one assigns these

¹NFS – Network File System

²RAID – Redundant Array of Independent Disks. In the RAID1 configuration, all files are stored to two or more hard drives simultaneously. Therefore, if one of them gets damaged the data can still get fully recovered.

³PERL – Practical Extraction and Report Language

⁴Apart from the non-measurable relative clock phases and the offset needed for the common start signal to reach the sensors serially within a JTAG chain.

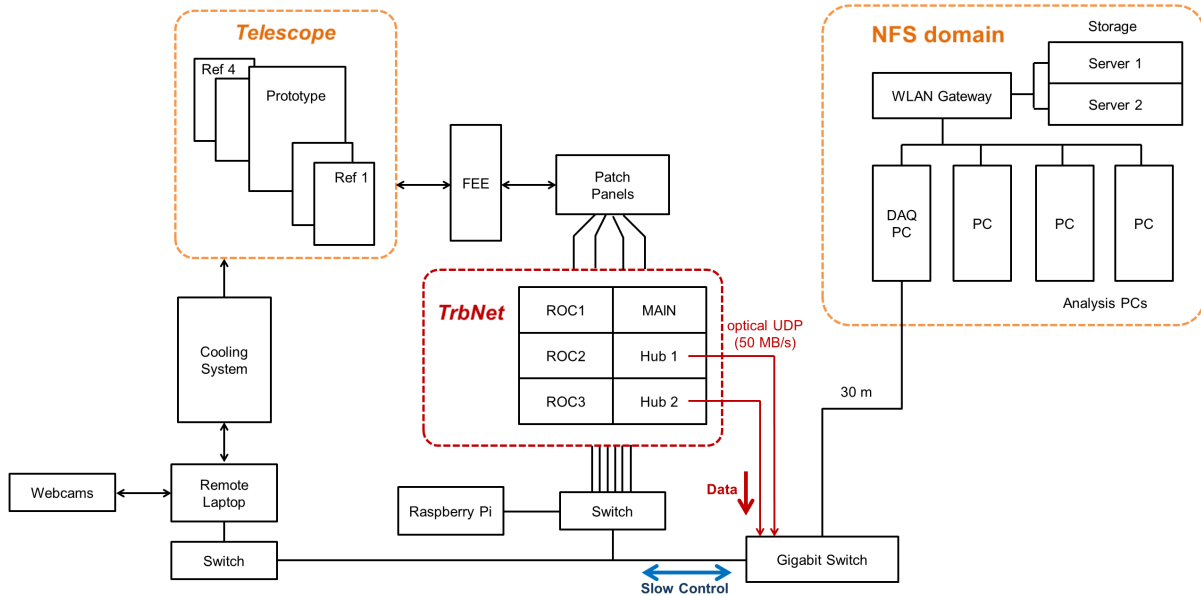


Figure 6.3: The architecture of the presented test setup. Two network domains are used interfaced by the DAQ-PC. The Gigabit switch connects the DAQ-PC to the cave setup and reads out the data over the optical Hub connection. The NFS domain is shown in the top right corner as part of the counting house. The Raspberry Pi mini-PC in the beam cave controls the sensor and FPGA power outlets.

packets to events. Each event comprises 12 frames from the 12 sensors. The sensors are synchronized, therefore the frames can be obtained simultaneously. If one sensor gets damaged or turned off, then an empty frame is still provided by the ROC containing sensor status information. The HADES event builder creates files with a certain structure which is presented in Appendix J. The events are automatically merged by the Hub (more details can be found in [111]). A fraction of the event file is copied to a different location in order to perform quality assessment.

6.2.2 Sensor Threshold Settings

As mentioned in section 3.3, the sensor fake-hit rate (FHR) can be associated with the detection efficiency. Lower comparator thresholds allow better particle detection at the cost of higher pixel noise. Thus, prior to starting the beam test, the FHR needs to be adjusted for maximum sensor performance. Hereby, Fig. 3.10 can serve as a reference. Fake-hit rates of 10^{-5} and 10^{-6} are of particular interest since the sensor is then expected to reach nearly full efficiency. In analogy to the comparators, the pixel matrix can be divided into four blocks as well, with 165 888 pixels per block. Therefore, the rate of 10^{-5} produces approx. 1.66 fake hits on average per frame, per block, and the rate of 10^{-6} leads to 0.166. An example of a sensor output for three different threshold settings is given in Fig. 6.4. According to the calculations at hand, the sensor thresholds can be tuned via two software tools, presented in Appendix I.

6.2.3 Slow Control and Monitoring Utilities

Upon establishing the variety of low-level tools needed for the beam test, a control system has been designed on a higher level of abstraction for faster and easier control. Several scripts and GUIs are developed granting access to the beam setup.

Basic tasks can be performed via the **control GUI**. Here, with one mouse click the user can reprogram all FPGAs, reset the network, start the event builder and perform other related tasks. The sensors are programmed over the **JTAG GUI**, which also allows many debugging options for the individual JTAG chains. All the relevant DAQ settings are kept in a **central configuration file** and can be edited in one place. The file allows changing the sensor thresholds, activating and deactivating individual columns and blocks of the pixel matrix as well as individual sensors, setting the sensor ID, setting the geometry parameters for monitoring tools and configuring the sensor order within the JTAG chains. A PERL script uses the file to update all the other, related scripts.

Monitoring tools are developed for the JTAG chain, the readout network, the sensor pixel matrices and the 3D beam setup display. They are presented in Fig. 6.5. Some additional scripts have been developed for the power supplies as well. The tools apply PERL-based scripts which dynamically generate HTML¹ pages. One Apache HTTP² server is running on the DAQ-PC to provide access to these websites. With this method, any computer in the network can monitor the system with a conventional web browser. The scripts are described in the following:

- **JTAG:** The JTAG scripts are accessing TrbNet in order to acquire all the JTAG related status messages from the MAIN board. They can display the errors for each of the three chains individually. Currently, there is no possibility to actively monitor the sensor registers. The register content can only be accessed during a write after which the sensor has a certain dead time. Therefore, reprogramming the sensor in order to monitor the JTAG registers is not considered efficiently.
- **Network:** The readout network together with the ROCs, Hubs and the MAIN board is monitored by generating one large ASCII³ file containing 247 register values which are acquired twice per second. Several monitoring scripts can access the file in order to create their individual HTML page. The monitoring scripts use customized software libraries to facilitate and modularize the HTML generation. They allow monitoring of the TrbNet status, network load, FPGA states, sensor activity, and other important sensor and network properties.
- **Pixel Matrix:** Few scripts responsible for the monitoring of the sensor pixel matrices are developed as well. The event builder stores a fraction of its current event file into a special folder which can be easily accessed by these utilities. The tools scan the events and reconstruct the pixel matrix hits of all active sensors. All the hits are integrated over time. The generated image files are then displayed in the browser.
- **Beam Setup:** Lastly, a 3D visualization tool has been incorporated into the DAQ to display the current geometry of the entire setup. It shows deactivated sensors, deactivated matrix columns and the current sensor position with respect to the beam, i.e. the sensor coordinate system. The tool uses 'Processing', a javascript library to display 3D images inside HTML5 websites [127]. The central configuration file provides the geometry information.

¹HTML – HyperText Markup Language

²HTTP – HyperText Transfer Protocol

³ASCII – American Standard Code for Information Interchange




Figure 6.4: The number of fake hits in the sensor matrix depends strongly on the comparator threshold settings. The applied threshold decreases from left to right. The data is collected for 30 seconds in all three images. A lower threshold setting (middle) allows more fake-hits to be accumulated over the time. The rightmost image is produced with the minimum possible threshold, in which case the comparators produce a fake hit for each pixel. Due to the sensor limitations, only 6 'states' per SDS bank and 9 'states' per row are stored, as explained in section 3.3. This leads to the unusual structure observed here. After 114 rows, the bandwidth limit is reached and the readout interrupted.




Figure 6.5: The monitoring tools create HTML pages which can be accessed by any computer in the internal network. Top left, the full 3D setup is visualized. Bottom left, two sensor matrices are displayed after collecting data for several minutes. To the right, the basic TrbNet monitoring is shown (DAQ Monitor). At the time, two sensors were deactivated. The sensor frame rate (8680.5 Hz) can be directly observed, which is also the number of frame requests answered per second.

6.2.4 Safety Measures

Such a complex system as a particle detector is prone to many errors. A variety of different components is necessary to obtain the data. They all have to fit together for a complex interplay between signal timings, noise and durability. At the same time, one small error in one system can brake the entire experiment. It is therefore crucial to detect any inconsistency within the various systems early on. Important error indicators do not disappear from the screen and have to be reset manually. If some errors are missed by the human eye, the ROCs write a full status report together with the sensor data for each frame allowing offline error analysis. The components are monitored according to the following list:

- **Sensors:** Failing sensor clock and other errors detected by the ROCs can be quickly found using the network monitors. A detailed 32-bit error report (the `FRAME_STATUS`) is embedded into the data stream as well. In addition, deviations in the event builder data rates indicate sensor errors and can be directly observed.
- **FEE:** The CB and QB boards have integrated power regulators which use red LED¹ light as an error indicator. Therefore, webcams are used to monitor their status.
- **JTAG:** Programming errors are observed via the JTAG monitor tool. Sensor programming is executed in a loop, e.g. with 1000 iterations. If no errors are found in any of the runs, it can be assumed that the sensors are programmed correctly.
- **TrbNet:** The network status can be viewed and checked with various monitoring tools. Deviations in the event builder data rates may indicate a network error as well.
- **Power Dissipation:** Sensor power supplies can be remotely controlled and monitored via PERL scripts. A crowbar circuit provides reliable overvoltage protection. Latch-up protection is not used during this experiment. TRB2 boards have on-board power regulators but their power supply is left unmonitored. In case of emergency the TRB2 power can be remotely switched off.
- **Cooling System:** The cooling temperature can be changed remotely and viewed by pointing one webcam at the cooling system display. But the displayed temperature does not reflect the actual sensor temperature. The sensor temperature could not be obtained during this beam-time, but improvements have been made for future experiments.

6.3 Measurements

The MVD prototype module has been expected to demonstrate a variety of important properties which are crucial for further detector development. Hence, it has undergone numerous tests. This section describes all of the performed tests during the in-beam test.

¹LED – Light-Emitting Diode

6.3.1 Preparations

Before measurement, proper sensor thresholds need to be set. They are different for each sensor, therefore the threshold GUI and the threshold finder are used to tune the fake hit rates to the desired level. Prior to starting the DAQ, a JTAG test is performed. After that, the setup is mechanically adjusted to find the beam spot. The measured coincidences in the scintillators are used as a reference. The initial beam intensity is found to produce approx. 8 000 coincidences per spill.

6.3.2 Performed Tests

The initial beam intensity is kept throughout most of the experiments. The rate is sufficient to analyze the tracking performance and perform the detection efficiency study within several days. Near the end of the beam experiment, a threshold scan is performed, followed by an inclination angle study and a temperature test. On the last day, the DAQ performance test is conducted. Some of the results can be found in [125], or in related works [94, 114]. However, this thesis focuses only on the final test regarding the DAQ performance.

The DAQ system could already demonstrate its stability during the first days of the in-beam test. But in order to proceed further, the system needs to be tested at its limits. Therefore, two TrbNet Hubs are connected in order to increase the readout bandwidth to 100 MB/s. One Hub is responsible for the reference planes (ROCs 1 and 3), while the other handles the prototype (ROC 2). After this short system upgrade, the beam intensity is increased to a maximum. The scintillators could measure an increased flux up to approx. 380 000 coincidences per spill. After obtaining some results with the high data rates, the next DAQ test is performed. Three of the off-beam reference plane sensors are put into their saturation mode by manually decreasing the comparator thresholds. Therefore, these sensors are responsible for an increased data load of ≈ 60 MB/s. After collecting some data with this setup, the network is put into an overloaded state by decreasing also the thresholds of several other sensors, in the end. Altogether, the data rate exceeded the 100 MB/s limit imposed by the Hubs. The results are discussed in the following section.

6.4 Results

This section summarizes the results obtained by the DAQ-related studies.

6.4.1 Initial Results

At first, the JTAG chains have been proven stable. All three chains were independently tested with 100 000 programming cycles. No errors were observed.

After setting the appropriate threshold values to reach a fake hit rate of 10^{-5} , the individual spills could be observed in the data. This served as a first proof that the DAQ and the analysis tools operate correctly. Also, the beam spot could be clearly identified within the pixel matrix of the prototype. Additionally, the sensor data was analyzed properly with the result that no status errors were present. All frames were extracted safely and no data was lost.

As an example, a run of ≈ 5 minute length has been chosen and presented in Fig. 6.6. The first plot (a) shows the number of clusters per frame as registered by one of the prototype sensors. The equidistant peaks in the image correspond to individual spills. The frame number of 2.5 Mil. corresponds to 4.8 minutes run time. The second plot (b) shows the distribution of hits during the entire run in one of the sensors. All neighboring pixel hits have been merged into clusters by the analysis software, and only the cluster center-of-gravity is displayed. Clearly, the beam spot can be observed in the lower-left corner of the matrix. The last image (c) shows that no errors are recorded during the run. A status '1' corresponds to no errors, whereas a '0' indicates one or more errors in the FRAME_STATUS. All sensors except one show a good status during the entire run, thus no frame has been lost. The damaged sensor is known to be broken, probably during the manufacture process. It is not used during the entire experiment.

6.4.2 Stability Analysis


Regarding the robustness, the DAQ system was operated with low beam intensity for four days. All the analyzed data sets up to now show no notable errors¹. No transmission errors are found in any of the examined data sets, however the data rate² of ≈ 6 MB/s is an order of magnitude below the maximum system capacity. On the last day, the beam intensity is increased up to 380 000 coincidences per spill, as indicated by the scintillators. The sensors have detected an order of magnitude more particles with approx. 90 hits per frame at the peak. The spill structure obtained with one of the prototype sensors is shown in Fig. 6.7. The observed data rate has increased to ≈ 25 MB/s per spill, on average. The readout network still showed no errors and all frames could be acquired without any loss.

As an example, a run of approx. 11.1 minutes length has been carefully analyzed with the result that none of the 5.7 Mil. events was lost. The status of all active sensors is denoting no errors, besides one. It appears that one of the prototype module sensors (sensor 8) occasionally showed some errors in 0.008 % of its frames due to reasons which are not understood yet. During the frame readout, it appears that the sensor data is not read out or encoded properly, possibly due to a SEU or SEFI error. According to section 3.3, the raw data can be interpreted either as a row encoding or as a column encoding (a 'state'). The data word 0x0B00 appears very frequently and disturbs the normal data flow. When it is obtained as a row package, the word is invalid since the last 0 denotes zero 'states' in that row. MIMOSA-26 does not encode rows without hits in its data stream. More frequently, the word is output as a 'state' in which case the ROC data checker often recognizes the error due to a preceding state that had a higher column address. Thus, the column addresses are not in the ascending order and the corresponding error-bit is set. This error could be present in more cases where the ROC data checker can not uncover it, e.g. with a valid row encoding followed by 0x0B00 as the first or only 'state'.

Since the network is designed to support higher data rates, several off-beam sensors were modified to increase their output rates. Data taking continued at an average rate of ~ 81.3 MB/s during the spill (not counting the UDP headers). This run was 8.4 minutes long and 4.35 Mil. events were taken. But being at its limit, some network packages could not get routed to the DAQ-PC which has lead to a loss of several events, presented in Fig. 6.8. In total, 221 339

¹A minor coding error in the ROC data checker has falsely marked less than 0.01 % frames with an error bit. The error is corrected after the beam time.

²No UDP headers are included in all data rate calculations.




(a) The analyzed data set contains six spills. The presented sensor occasionally produces up to three fake-hits.

(b) The hits in one prototype sensor are displayed. The beam spot can be identified in the lower-left corner.

(c) The status analysis shows no bad frames, besides the one mechanically damaged sensor which is known.

Figure 6.6: The results of the initial analysis.



(a) Noticeably more clusters are found in the pixel matrix than previously in Fig. 6.6a. Source: [125].

(b) The spill structure, belonging to a different spill, after zooming in. There can be a 50 Hz pattern observed which is common for the CERN SPS accelerator.

Figure 6.7: The spill structure during the increased beam-intensity test.



Figure 6.8: During the high data rate run, 5.1 % of the events were lost. The network bandwidth has been near the limit and the event builder has discarded some events due to missing UDP packets.

events have been lost during the run, which is approx. 5.1 %. All sensors have lost their frames simultaneously. Since no error bits are set by the ROC, it appears that the event builder has discarded some of the events due to incomplete UDP packets, caused by a network overload.

During all the experiments, the data acquisition was running stable for five days with only one short interruption. One of the converter boards needed to be exchanged due to a hardware defect. This could be observed directly on the monitoring screen. One of the sensors did not transmit its data all the time, the clock disappeared occasionally and the frames-per-second rate has decreased. The problem was solved with the replacement of the broken PCB within minutes.

6.4.3 Sensor Synchronization

The clock for all three JTAG chains is driven from a single source and the chains also receive a common start pulse, therefore all frames should be synchronized due to the equal integration time. The ongoing data analysis could confirm this statement [128]. Furthermore, the statement is validated by examining the intrinsic frame numbers given by the sensors and checking the timestamps provided by the ROCs.

There is a certain difficulty introduced due to the fact that the common clock signal is situated on a different board, the TrbNet Hub, as shown in Fig. 6.9. The 100 MHz clock needs to pass through several optical links before it reaches the MAIN board or a ROC. The logic on the MAIN board then uses a PLL to transform the clock to 80 MHz and triples it onto three different output ports for each of the JTAG chains. The FPGAs are known for their deterministic properties. The jitter on the output ports is in the ps range, and the clocks remain phase synchronous as they come from a single source, therefore the sensor synchronization is guaranteed by the setup.

However, the recovered MAIN board clock is not phase-locked with any of the ROCs. The optical link uses serializer/deserializer components operating at GHz frequency to transmit 16-bit TrbNet packets serially over a single optical fiber. The clock is integrated into the bit-stream, alongside the 8b10b¹ encoding. This enables clock recovery on the receiver link, but the recovered clock is not optimal. It depends on the phases between the optical links to each other, particularly their barrel shifters. If the barrel shifters are not aligned, the recovered clocks on the ROCs and the MAIN board will deviate over time. Thus, the *ROC_SYNC* signal is needed to synchronize the timestamps. However, the signal can not be used properly, as described at the end of section 5.4.6, due to already occupied LVDS input lines on the ROCs. Nevertheless, the ROC timestamps can be compared to each other in order to search for effects which could be attributed to inappropriate sensor synchronization.

Several data sets with varying beam duration are analyzed. Events containing 12 frames from different sensors are examined one-by-one. The timestamp is set by the ROC anytime the sensor Header bit pattern is detected. The results considering the longest run (\approx 240 Mil. frames, 7.6 hours) are discussed in following.

The results are presented in Fig. 6.10. The phase difference between ROCs 2 and 3 is not too large, whereas ROC 1 shows a larger deviation. This is a completely random effect which can be eliminated by using the *ROC_SYNC* signal from the MAIN board. It can be concluded that once the phase difference is established, it does not change which is re-expressed in the linear slope of the increasing timestamp deviations. The vast majority of the analyzed frames takes 115.2 μ s to

¹8bit/10bit encoding is used to create data streams with equally many '0'-s and '1'-s for a better voltage balance and clock recovery. Eight bit of data are encoded with a ten-bit code.

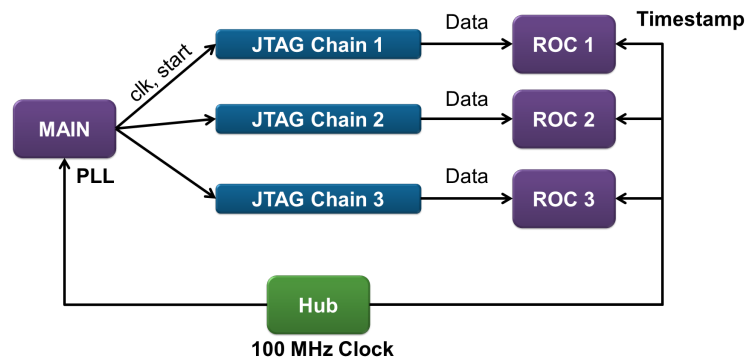


Figure 6.9: The synchronization of the readout network. The common clock is transmitted by the TrbNet Hub over optical links. The logic on the MAIN board uses a PLL to transform the clock to 80 MHz, whereas ROCs use the clock directly to time stamp the incoming frames. The recovered clock on the MAIN board is not phase-locked with any of the ROCs.




Figure 6.10: The deviation of timestamp-generating counters on different ROCs. The timestamp is set every time a frame Header is detected. Since the sensors are synchronized, all readout chains set the timestamps at the same time in each ROC. However, the timestamp-generating counters on different ROCs drift apart due to the suboptimal clock recovery. The timestamps deviate by a total of ≈ 300 ms during the 7.6 hours. However, ROCs 2 and 3 deviate only by ≈ 50 ms in total. The linear slope indicates deterministic system properties as expected.

Sensor	1	2	3	4	5	6	7	8	9	10	12
Time -1	—	—	—	—	5.5 %	5.5 %	5.5 %	5.5 %	6.9 %	6.9 %	6.9 %
Time +1	6.0 %	6.0 %	6.0 %	6.0 %	0.0 %	—	—	—	—	—	—
ROC	ROC 1			ROC 2			ROC 3				

Table 6.1: The timestamp deviations per sensor. Over 93 % of all timestamps are set exactly after 11 520 FPGA clock cycles, i.e. the MIMOSA-26 integration time. The ones that deviate by one clock cycle (corresponding to 10 ns) are displayed here. The damaged sensor 11 is excluded from the table. Again, ROCs 2 and 3 show similar behavior, while ROC 1 distances itself (see Fig. 6.10).

arrive which is exactly 11 520 FPGA clock cycles. Only 5 - 7 % of the timestamps deviate by one clock cycle which can be accounted to the suboptimal clock recovery. No timestamp deviates by more than one FPGA clock cycle. The results of the analysis are shown in table 6.1.

As a conclusion resulting from the evaluation of Fig. 6.10 and table 6.1, the sensors in the setup behave fully synchronized. The timestamps deviate by a constant factor which is understood and which can be extracted from the data. The original MAIN board clock phase is in between the ROC 1 phase and ROC 2 and 3 phases. More than 93 % of all frames indeed arrive after 115.2 μ s, hence the rest must be miscounted by one clock cycle due to the phase deviation. The frame numbers of all sensors are also equal for each event. Only some frame numbers of sensor 2 are showing an error which is not related to sensor synchronization. Therefore, it is discussed in section 6.4.5. Additionally, as can be seen, sensor 5 seems to provide the sensor Header in certain cases one clock cycle after the 115.2 μ s limit, although its phase configuration is adjusted to receive it one clock cycle sooner. This effect is not understood yet and could be attributed to an internal sensor error.

6.4.4 Overload Study

As the final DAQ performance test, the thresholds of the prototype sensors were decreased producing data rates beyond the 100 MB/s network limit. In such case, some errors were detected by the ROCs indicating buffer overflows. However, the event builder continued at a data rate of 99 MB/s, i.e. the network did not stop operating.

In total, 6.3 Mil. events corresponding to 12.2 minutes were recorded at an effective rate of up to 97.8 MB/s per spill (excluding the UDP headers and the discarded data packages). But this time, many disruptions were detected in the data stream and the ROC error bits were set frequently. In most cases the error bits indicate an overflow of TrbNet buffers, but in some cases also other, sensor-related errors are observed, e.g. a loss of the clock pulse. This is believed to be a chain controller error since the errors occur only during this single run. This issue is discussed in section 6.4.5.

The TrbNet buffers were running full during each spill. In such case, an error bit is set inside the TrbNet which is transmitted with the following frame request to all ROCs. Once the ROCs receive the error bit, they stop their internal data readout procedure (see section 5.2.4). Only after the TrbNet buffers have been read out sufficiently, the data taking can continue. Therefore, the data stream gets interrupted, producing the pattern shown in Fig. 6.11. Despite the overloaded network, the ROCs still provide their frames simultaneously. However, one error has been observed. At each interrupt, the frame numbers get shifted by 1. But only some sensors are affected. This leads to the conclusion that the procedure to keep the frame readout synchronous under network overload is not optimal. Even though the principle is proven correct, further optimization strategies and an error analysis need to be performed. Hence, serving as a first approach to this problem it can be concluded that the overload issue is solvable but requires some modifications to clear all errors.

6.4.5 Summary and Evaluation

The presented DAQ studies are summarized in table 6.2. All stability tests are successful without any notable errors, regardless of beam intensity. Only when the data rate increases towards the limit, some Ethernet packets are discarded leading to a loss of a small fraction of events. The




Figure 6.11: An excerpt of the data provided by sensor 1 during the overload phase. A '0' denotes the loss of the frame. The frame discarding algorithm operates well, since the slope is proven to coincide with all the other sensors. Bursts of up to ≈ 50 frames are recorded coherently despite the network overload.

Name	Coincid.	Duration	Events	Data rate	Errors
Common Tests	8 000	4 days	uncounted	≈ 6 MB/s	1 CB broken
Stability	8 000	1 hour 27 min.	45.4 Mil.	≈ 6 MB/s	none
High Intensity	$\leq 380\,000$	11.1 min.	5.7 Mil.	≈ 25 MB/s	none
High Data Rate	$\leq 380\,000$	8.4 min.	4.35 Mil.	81.3 MB/s	5.1 % lost
Synchronization	8 000	7.6 hours	238.5 Mil.	≈ 6 MB/s	none
Overload	$\leq 380\,000$	12.2 min.	6.3 Mil.	97.8 MB/s	Frame shift

Table 6.2: Summary of the performed DAQ tests. The beam intensity is expressed in terms of the scintillator coincidences (Coincid.). High intensity studies involve beam intensities between 330 000 and 380 000 coincidences. Only the maximum effective data rate without the UDP headers is listed, as calculated during the spill. There are few more errors not listed here which are not related to the given study. They are explained further in text.

sensor synchronization is guaranteed by the setup and could be confirmed within the long test of 7.6 hours. However, the overload study gives room for improvement. The resynchronization algorithm is not optimized and leads to a loss of global frame synchronization.

In the following, the errors detected during the studies are discussed. At first there has been an error eliminated which would under certain circumstances falsely mark a frame as bad. The ROC data checker would produce a certain error bit even if the frame is intact due to a coding error in the VHDL code. Though, nearly all events are unaffected and the issue has been resolved after the experiment.

As for the sensor status, two sensors were found to produce some errors. The data format of sensor 8 has shown a particular error where the row-column allocation is interrupted within a frame by a $0x0B00$ package. This error occurred in 0.008 % of the analyzed frames. Furthermore, sensor 2 showed some errors at the output of its frame number. The registered frame number is not in the ascending order, but exhibits a bit-flip in one of the MSB. The frame after that, the number continues in its original order. A high-energetic impinging particle could have modified the sensor- or FPGA-RAM, an internal register or a multiplexer, thus causing the error (SEU or SEFI, see section 2.3.4). This error occurred only five times during the analyzed run of 7.6

hours duration and no other sensor is affected. Further explanations involve a defective cable or improper routing inside the FPGA chip. Both of the presented errors could also indicate first signs of radiation damage in the sensor readout circuits.

Regarding the sensor synchronization, no errors were observed. However, such a trivial setup with only three JTAG chains driven from a single source is not difficult to manage. In future, a synchronized network is required, e.g. CBMnet or synchronous TrbNet to support a larger number of synchronized JTAG chains. The unresolved synchronization error with sensor 5 from table 6.1 will not be pursued any further. Since no other sensor showed similar behavior, it is very likely that the error is caused by the sensor itself. More important is the application of the *ROC_SYNC* pulse in future experiments to synchronize the timestamp allocation between different ROCs. During the presented in-beam test, ROC synchronization could not be performed efficiently due to the absence of a free LVDS input pair on the ROCs.

However, the only dissatisfying test showing undesired behavior is uncovered by the overload study. The current ROC implementation is not sufficient to guarantee global frame synchronization in case of an overloaded network. The network overload is handled correctly, only inappropriate frames are provided to the TrbNet. The frame numbers are not equal, but get shifted over time. This issue requires further investigation and could be attributed to the chain controller. A future update should be able to correct the problem. Additionally, there seems to be another issue related to the chain controller. The complexity of the component is rather large (see section 5.2.4), thus an optimal connection of all of its modules seems not to be provided yet. For example, it appears that if one major error indicating a bad frame is found (e.g. a buffer overflow), further error checks might be omitted leaving an incomplete error report. During the overload study, apart from the buffer overflow, the status bits have also signaled sensor errors which were not present before and after the test. Thus, the component should be re-examined and optimized in future to reach its full efficiency.

Altogether, neglecting the minor issues found in the studies, the in-beam test was very successful. Under normal conditions, no data has been lost and no unusual behavior observed. The DAQ system is able to perform MIMOSA-26 related studies and obtain important, unprecedented MVD prototype information. The hardware-near error checking and the frame-wise output of status bits for each sensor has proven to significantly increase the speed of the offline data analysis. Some improvements are scheduled for the near future that should eliminate all the minor design flaws still present in the code and equipment, e.g. eliminate the frame number shift during the network overload, optimize the chain controller and improve sensor temperature measurement via converter boards.

6.5 Conclusions After the Test

The proposed readout controller has demonstrated its suitability for small-sized MVD setups. The FPGA-driven readout allows early prototyping stages and supports a high level of scalability. One PCB can serve various purposes, e.g. as a readout-, central network- or JTAG-controller. The developed firmware supports a variety of options allowing rapid reconfiguration for different tasks and setups. The VHDL programs can be precompiled for a given use case and loaded within seconds, on demand. All of the readout chain modules fulfill their task, however the chain controller requires further optimization.

The applied readout network prototype can be operated with common HADES-DAQ tools. The network- and most register-addresses are modifiable. New slow control and monitoring routines can be integrated as well. Moreover, additional Hubs can be included to support higher readout rates. The network is therefore highly adaptable and can be used for many different applications.

The sensor programming and slow control over the JTAG interface has been validated as well. Due to its generic nature, it is possible to integrate the JTAG-controller code into any FPGA. Current implementations support a dedicated JTAG programming board and a central controller board (MAIN). The sensor synchronization is of prime importance, however all tests have shown that with the proposed method sensors do not lose their synchronization over a longer period of time. Thus, if the network can be synchronized and driven with a common master clock, it is possible to integrate the JTAG controller on each ROC individually. Since CBMnet will be able to provide such a feature, early tests with the TRB3 board have been performed in order to demonstrate the compatibility of the JTAG controller with the readout chains, as indicated in section 5.3.2. Currently, both modules can be integrated on one TRB3 FPGA and accessed individually via slow control.

The CBMnet interface is not implemented at the moment, but will be included in near future. The FLES interface needs to be developed as well. One important aspect was not covered during the in-beam test. The sensor raw data was not pre-processed. After reading the data out of the frame buffer, it has been directly forwarded to the TrbNet Endpoint. No cluster finding and no FLES encoding algorithms were applied. Their implementation is still necessary in future which is expected to increase the invested FPGA processing time per sensor package. The time to process each sensor packet by the FPGA, t_{Packet} from section 5.1.1, can be used as a benchmark.

Since the readout chain and the JTAG controller are both compatible with the TRB3 board, a ROC version is currently being developed to merge them together on the peripheral TRB3 FPGAs. Such next-generation ROC would finally provide adequate resources for the SIS-100 experiments. A number of 16 - 48 sensors can be handled by the TRB3 resources, which requires 7 - 19 TRB3 boards for the final setup. However, the uneven assignment of sensors to ROCs requires additional add-on cards. The data rates of 22 Gbit/s can be handled by 55 optical links, each operating at 400 Mbit/s data extraction rate, as applied at the beam time. The central TRB3 FPGA can be used as a data Hub providing access to four such links, hence a total of 14 TRB3 boards could already be sufficient for the readout. Under the assumption that the integrated frame buffer is sufficiently large to handle the highest beam intensity fluctuations, the proposed network would suffice to perform all of the envisaged SIS-100 experiments.

Chapter 7

Summary

In the scope of this thesis, a readout system for the CBM-MVD detector has been proposed and prototyped. The readout is based on the readout controller PCB (ROC) featuring FPGA chips and optical transceivers. Several models of the board from the HADES experiment have been considered. The TRB3 board, showing the best results, has been selected as a possible candidate for the initial CBM experiments at FAIR. In order to demonstrate the proof of principle, a previous model, the TRB2 board, has been applied for the initial prototyping stages.

The MVD is based on the monolithic active pixel sensor technology, therefore, the current state-of-the-art represented by the MIMOSA sensor family has been introduced and extensively analyzed. Subsequently, a suitable FPGA firmware supporting all present and some of the upcoming MIMOSA generations has been developed. The core part of the designed FPGA module is represented by the readout chain. One readout chain can handle one sensor and runs independently from others. The raw sensor data is transferred to the FPGA clock domain, deserialized, decoded and buffered. The buffer can be used to alleviate the effect of beam intensity fluctuations. In parallel, the data is checked for syntactic and semantic errors, allowing early failure detection and controlled frame rejection strategies. All errors are stored in a 32-bit status word and can be accessed either via slow control, or from the data stream where it is written during a customized formatting process. By cascading several readout chains in parallel, the firmware can be applied to read out an arbitrary number of sensors.

In order to test and validate the design, an MVD readout network is developed using the synergy with the HADES experiment. The network applies TrbNet for data transmission and slow control. The unmodified TrbNet components are accompanied by additional wrapper modules to enable free-streaming data acquisition. The sensor data is, by design, organized into frames which are provided at fixed time intervals. An optical network was created with the TRB2 boards and TrbNet Hubs to successfully handle the sensor frames. The final result was tested at the CERN SPS secondary test beam facility. The test involved 12 sensors organized within three TRB2 boards and programmed via JTAG on a separate TRB2. The sensors were synchronized by using one common clock and by transmitting a simultaneous start signal.

The in-beam test has focused on the MVD prototype core module. The module represents a micro-tracking device with a low material budget. A beam telescope with five stations has been created. MIMOSA-26AHR sensors were applied for the prototype and the reference planes. A dedicated control system is developed for the experiment providing sufficient monitoring and

slow control functionalities.

The network is proven very stable during all conventional tests within 4-5 days of nearly uninterrupted beam, regardless of beam intensity. All test results indicate perfect sensor synchronization. By increasing the data rate to the network limit, some frames were automatically discarded, according to the expectations. After putting the network into a constantly overloaded state, a data throttling algorithm was tested. Even though the principle has been confirmed, the algorithm requires further studies. Apart from several minor issues, unprecedented results have been achieved during the in-beam test. The network could be operated dead-time free and the data quality has been proven excellent which has contributed to specify crucial prototype parameters.

The ROC firmware has passed all of the main tests and is suitable for further MVD studies. In order to increase its computing resources, the ROC is foreseen to continue its operation on the TRB3 platform. However, this requires a suitable add-on board to interface the sensors, as well as the corresponding front-end electronics, which are currently under development. If the sensor architecture for SIS-100 does not change considerably, the ROC firmware can be integrated (with minor changes) into the final MVD setup as well. Slow control and monitoring can still be continued via TrbNet, in parallel, in order to take profit from the large debugging and control utilities at hand.

Zusammenfassung

Motivation

Das GSI¹ Helmholtzzentrum in Darmstadt wird zurzeit um weitere Teilchenbeschleuniger, Speicherringe und eine Vielzahl an Experimenten erweitert. Die neue Einrichtung, unter dem Namen FAIR², wird unter Anderem ermöglichen, Materie in einen Zustand zu versetzen, wie er kurz nach dem Urknall vorzufinden war. In dieser heißen und dichten Phase, dem sogenannten Quark-Gluon-Plasma, werden die Nukleonen in ihre elementaren Bestandteile aufgelöst. Hierbei kommt es zu Phasenübergängen, die bei genauer Betrachtung Aufschluss über die tiefgründigsten Mechanismen der Natur geben könnten. Um diesen Zustand insbesondere in der Umgebung maximaler baryonischen Dichten systematisch zu charakterisieren, wird das Compressed Baryonic Matter (CBM) Experiment aufgebaut. Dieses Fixed-Target Experiment wird einen hochauflösenden Spektrometer einsetzen, um Teilchen aus Schwerionen Kollisionen zu messen, die einen Aufschluss über die hochdichte Phase der Reaktion geben könnten.

Besonders interessant sind Teilchen mit Charm-Quark Inhalt. Sie eignen sich dazu, Rückschlüsse auf den Zustand der Materie in der frühen Phase der Schwerionenkollision zu ziehen. Da solche Teilchen im Energiebereich von CBM noch nie untersucht wurden, gelten sie als besonders geeignet um die verschiedenen physikalischen Theorien und Modelle zum Verlauf der Schwerionenkollision und dem eventuellen Auftreten eines Phasenübergangs zu testen. Allerdings zerfallen diese Teilchen innerhalb von wenigen 100 μm und müssen deshalb aus den Eigenschaften ihrer Zerfallsprodukte rekonstruiert werden. Die Zerfallsprodukte sind zusätzlich von dem gewaltigen Untergrund aus dem primären Kollisionsvertex, dem Punkt der Schwerionenkollision, schwer zu unterscheiden. Um Open-Charms Teilchen, wie die D-Mesonen dennoch zu erkennen, verwendet man die Methode zur Rekonstruktion Sekundärer Vertices mit einem geeigneten Mikro-Vertex Detektor (MVD). Der CBM-MVD wird CMOS-Pixel Sensoren einsetzen, angeordnet in vier planaren Stationen, um den Zerfallsvertex von D-Mesonen durch Spurrekonstruktion zu vermessen und somit den Teilchenhintergrund effizient zu unterdrücken.

Da die Wirkungsquerschnitte für die Produktion von Charm-Quarks bei CBM-Energien sehr niedrig sind, müssen die entsprechenden Experimente mit hohen Kollisionsraten durchgeführt werden. Demzufolge müssen alle wesentlichen Komponenten strahlungshart sein und über eine gute Zeitauflösung verfügen. Um die Experimente mit höchsten Kollisionsraten zu unterstützen, wird eine Auslesezeit von 30 μs für die ersten und weniger als 10 μs für die finalen Experimente betragen. Darüber hinaus müssen Massnahmen ergriffen werden, um die anfallenden Messdaten in Echtzeit aus den Detektorsystemen auszulesen und an ein geeignetes Datenverarbeitungssystem weiterzuleiten. Die Teilchen müssen mit einer Effizienz deutlich oberhalb von 99.5 % und einer XY-Ortsauflösung von unter 5 μm pro Station gemessen werden. Zur präzisen Vermessung

¹GSI – Gesellschaft für Schwerionenforschung

²FAIR – Facility for Antiproton and Ion Research

der sekundären Zerfallsvertizes ist es ausserdem von besonderer Bedeutung, die ersten Detektorstationen unmittelbar hinter dem Target so dünn wie möglich auszulegen und im Vakuum zu betreiben, um Vielfachstreuung zu minimieren und so die Selektivität dieser Methode zu steigern. Das Materialbudget sollte 0.3 % für die Erste und 0.5 % X_0 für die darauffolgenden Stationen nicht überschreiten. Dies begrenzt stark den Querschnitt und die Bandbreite der eingesetzten Kabel und erfordert eine genau Abschätzung der Sensordatenraten. Eine ionisierende Strahlung von 3 MRad und nicht-ionisierende Strahlung von 10^{14} n_{eq.}/cm² sollten von den Sensoren und der Front-End Elektronik toleriert werden.

Diese Voraussetzungen lassen sich nur mit einer ausgereiften Pixelsensor Technologie in Einklang bringen. Die MIMOSA¹-Familie des IPHC² Institus aus Strassbourg gilt als eines der Vorreiter dieses Gebiets und wurde unter Anderem auch wegen der Kosteffektivität kommerzieller CMOS-Prozesse für den CBM-MVD ausgewählt. Das Ziel der vorliegenden Arbeit ist es, ein performantes Auslesesystem für den CBM-MVD zu entwickeln und Dieses mithilfe erster Studien mit Daten und Simulationen zu charakterisieren.

Anforderungen an das MVD Auslesesystem

Zunächst gibt es Anforderungen die allein durch die Sensortechnologie vorgegeben sind. Die MIMOSA Sensoren verarbeiten die Teilchentreffer in der Pixelmatrix zeilenweise mit einer fest vorgeschriebenen Auslesezeit. Dabei wird der Datenstrom nullunterdrückt, was bedeutet dass nur Informationen bezüglich der getroffenen Pixel ausgegeben werden. Hierbei kommt ein spezifisches Datenformat, mit einer fest vorgegebenen Wortbreite zum Einsatz. Zudem erfolgt die Ausgabe seriell, also bitweise, auf wenigen Ausgabekanälen.

Um die Trefferinformationen vorab zu extrahieren, die Pakete weiter zu komprimieren und in ein CBM-spezifisches Datenformat einzukodieren, müssen die Wörter zunächst deserialisiert und anschliessend dekodiert werden. Dies erfordert zusätzliche Datenverarbeitungsschritte innerhalb des Auslesesystems. Allerdings steht das endgültige Datenformat noch nicht fest. Es wird von der Geometrie und Leistung des finalen MIMOSA Sensors abhängen, welcher noch in Planung ist. Demnach wird es erforderlich sein, ein Auslesesystem vorzubereiten, welcher sich den Anforderungen und der Leistung des finalen Sensors anpassen lässt.

Des Weiteren muss das gesamte lokale Auslesesystem an die restlichen CBM Bestandteile angepasst werden. Aufgrund ausserordentlich hoher Reaktionsraten werden Datenraten im Bereich von 1 TByte/s im gesamten CBM Detektor anfallen. Davon sind aber nicht alle Ereignisse vom physikalischen Interesse und ein Grossteil kann verworfen werden. Demzufolge benötigt man einen Trigger-Mechanismus, um die interessanten Ereignisse herauszufiltern. Der Grossteil der Observablen, die zum Triggern herangezogen werden können, können erst nach einer vollständigen Spurenrekonstruktion berechnet werden. Dies erfordert die Implementierung eines aufwendigen Tracking-Systems, dass einen bedeutenden Teil der vollständigen Daten in Echtzeit verarbeiten muss. Erst nachdem dieser Verarbeitungsschritt abgeschlossen ist, kann eine Trigger-Entscheidung erfolgen und die Datenrate reduziert werden. Hierbei ist es angestrebt, nach Datenreduktion einen Datenstrom von 1 GByte/s auf die Massenspeicher des CBM-Experimentes zu senden.

¹MIMOSA – Minimum Ionizing particle MOS Active sensor

²IPHC – Institut Pluridisciplinaire Hubert Curien

Allerdings bringt die Komplexität solch eines Auswahlmechanismus ganz neue Vorgaben mit sich. Die Trigger-Entscheidung steht den Front-End Modulen nicht zur Verfügung. Demnach müssen alle Daten frei laufend und mit einem Zeitstempel versehen an die höheren Auslestufen weitergeleitet werden. Die Vergabe der Zeitstempel erfordert eine globale Synchronisierung aller Auslesekomponenten. Dies wird bewerkstelligt durch ein synchrones Auslesenetzwerk, welches zurzeit in Entwicklung ist. Dem MVD wird eine synchrone Hauptclock mit extrem kleinen Jitter-Werten zur Verfügung gestellt, die verwendet werden muss um die Daten innerhalb des MVDs mit möglichst genauen Zeitstempeln zu versehen. Ausserdem müssen die notwendigen CBM Schnittstellen und die geplanten Datenformate implementiert werden.

Durchgeführte Simulationen und Ergebnisse

Um die Leistungsfähigkeit des MVD Auslesesystems genauer abzuschätzen, wurden im Rahmen dieser Arbeit mehrere Simulationen durchgeführt. Es wurden die zwei anspruchsvollsten Reaktionen gewählt, die während der Anfangsphase der FAIR Experimente am SIS-100 Synchrotron von Bedeutung sein werden. Die Au-Au Reaktionen werden die höchsten Teilchenmultiplizitäten generieren. Sie wurden mit einer Strahlenergie von 10 AGeV und Reaktionsraten von 10 - 100 kHz simuliert. Weiterhin sind p-Au Experimente geplant, die wesentlich weniger Teilchen pro Reaktion erzeugen und daher mit viel höheren Reaktionsraten durchgeführt werden können. Sie wurden mit 30 GeV Strahlenergie und 1 - 10 MHz Reaktionsraten simuliert.

Alle Reaktionen sind mit dem UrQMD¹ Model erzeugt und die resultierenden Teilchen mit GEANT3² durch ein Detektormodel propagiert worden, welches in das CBMroot-Framework integriert ist. Das Framework ermöglicht es, die jeweiligen Treffer in den MVD-Stationen bis auf die Pixelebene genau zu bestimmen und somit die Trefferdichten, Datenraten und Sensorlimitierungen zu analysieren. Für den MVD wurde ein virtuelles Sensormodell verwendet, welches auf den neuesten Forschungsergebnissen der MIMOSA-Familie basiert. Es wurde eine Pixelgrösse von $22 \times 33 \mu\text{m}^2$ angenommen. Die Pixel sind verteilt über eine sensitive Fläche von 3 cm^2 . Die Auslesezeit betrug $30 \mu\text{s}$ in den Simulationen. Die Daten wurden über 6 Auslesekanäle mit einer Bandbreite von 320 Mbit/s pro Kanal ausgelesen. Ausserdem wurde ein vereinfachtes CBM-Detektor Model verwendet, bestehend ausschliesslich aus dem Dipol-Magneten, dem Target, dem Strahlrohr und dem MVD. Für den MVD wurde allerdings ein realistischer Detektoraufbau eingesetzt, inklusive den Sensoren, dem Supportmaterial, dem Klebstoff und der Verkabelung. Dies ist wichtig für die möglichst genaue Abschätzung der Sekundärteilchen, welche in Materialien erzeugt werden können. Insbesondere wird bei Au-Au Experimenten eine Vielzahl an δ -Elektronen im Target erzeugt, welche aufgrund des Magneten auf bestimmte Bereiche des MVD Detektors gelenkt werden. Diese Zonen werden sehr hohen Trefferbelastungen ausgesetzt sein und stellen somit die höchsten Anforderungen an den Detektor. Darüber hinaus definieren sie in vielerlei Hinsicht die Anforderungen an die Sensoren, insbesondere im Hinblick auf die Bandbreite, die Auslesezeit, die Pixelgeometrie und die Strahlungshärte.

Ein weiterer Faktor, welcher die Sensorbelastung beeinflusst und bei den Simulationen ebenfalls eine entscheidende Rolle spielt, ist die erwartete Schwankung der Strahlintensität und deren zeitliche Struktur. Es ist gegenwärtig zu früh, um genaue Prognosen für den SIS-100 Synchrotron zu treffen. Deshalb wurden die Erfahrungen am laufenden GSI Synchrotron SIS-18 herangezogen.

¹UrQMD – Ultra-Relativistic Quantum Molecular Dynamics

²GEANT – GEometry ANd Tracking

gen. Am SIS-18 wurde festgestellt, dass die Strahlintensität auf einer Zeitskala von 30 μ s die mittlere Intensität um bis zu einem Faktor 3 übersteigen kann. Die Experten erwarten einen ähnlichen Verlauf für SIS-100. Die Sensoren werden dadurch kurzzeitig einer deutlich erhöhten Reaktionsrate ausgesetzt. Auch dies wurde mithilfe der Simulationen eingehend studiert.

Die Resultate, basierend auf dem verwendeten Sensormodell für den MVD-Sensor am SIS-100, zeigen, dass der Sensor samt Auslese erst bei Strahlintensitätsschwankungen, wie sie am SIS-18 beobachtet wurden, an Grenzen stösst. Sie unterstreichen so die Notwendigkeit, diese Schwankungen auf ein Minimum zu reduzieren. Es konnte gezeigt werden, dass vorwiegend drei Sensoren die Bandbreitenlimitierung überschreiten. Sie befinden sich alle in der δ -Elektronen Zone. Aus diesem Grund sollten die Au-Au Experimente, bei den berücksichtigten Strahlintensitätsschwankungen, die 50 kHz Reaktionsrate nicht überschreiten. Insgesamt wurde eine durchschnittliche Datenrate von bis zu 22 Gbit/s im MVD produziert. Allerdings konnte beobachtet werden wie sich die Gesamtrate aufgrund von Intensitätsschwankungen vorübergehend verdoppelt. Vereinzelt Sensoren die sich vorwiegend in der δ -Elektronen Zone befanden, haben ihre Ausgabe nahezu verdreifacht. Daher sollte das Auslesesystem für eine eingehende Datenmenge von 50 Gbit/s ausgelegt werden.

Aufgrund der sehr unterschiedlichen Belegungsdichten der Sensoren innerhalb des MVD erzeugen die einzelnen Sensoren sehr unterschiedliche Datenraten. Gleichzeitig ist das Auslesesystem der Sensoren im Wesentlichen durch die Datenbandbreite der eingesetzten FPGA-Platinen begrenzt. Es ist daher naheliegend, die Zahl der von der individuellen Platine ausgelesenen Sensoren an die Datenraten der jeweiligen Sensoren anzupassen. Da die meisten Sensoren eher wenige Daten produzieren, können durch eine entsprechende Optimierung erhebliche Hardwarekosten eingespart und die Komplexität des Systems reduziert werden. Dieser Argumentation folgend müssen die Platinen hinreichend flexibel gestaltet werden, um mit einer veränderlichen Zahl von Sensoren zurecht zu kommen.

Implementierung des Auslesenetzwerks

Um frühstmöglich mit dem Aufbau eines Auslesesystems zu beginnen, werden rekonfigurierbare FPGA¹ Mikrochips eingesetzt. Sie können je nach Bedarf ihre logischen Funktionen ändern und somit verhältnismässig einfach an unterschiedliche Sensortypen und Experimente angepasst werden. Die Realisierung des MVD Auslesesystems kann prinzipiell in drei Teile gegliedert werden. Als Erstes muss eine MIMOSA-Kompatible FPGA Firmware implementiert und eine geeignete Mehrzweck-Platine für die Auslese gefunden werden. Als Zweites sollte anschliessend ein minimales Auslesenetzwerk aufgebaut und getestet werden, bestehend aus mehreren MIMOSA Sensoren, FPGA-Platinen und weiteren Hardware-Modulen zur Steuerung, Sensorprogrammierung und Überwachung. Als Drittes muss das minimale System schlussendlich zum finalen Auslesenetzwerk erweitert werden, was aufgrund der geforderten Skalierbarkeit durchführbar sein sollte. Diese Arbeit hat sich mit den ersten zwei Themen befasst. Es wurde ein Grundgerüst erstellt und ausgiebig getestet, mit dem es möglich ist, das finale Auslesenetzwerk vor der Inbetriebnahme des CBM Experiments am SIS-100 aufzubauen.

¹FPGA – Field Programmable Gate Array

Die FPGA-Firmware muss eine verschiedene Anzahl an Sensoren pro FPGA unterstützen, wobei sie flexibel an verschiedene MIMOSA-Generationen angepasst werden sollte. Des Weiteren muss sie erweiterbar bezüglich neuer Datenverarbeitungsschritte sein, um künftige Datenkompressionsalgorithmen und Ausleseprozeduren zu unterstützen. Aus diesen Gründen wurde ein modularer Ansatz gewählt. Die ankommenden Sensordaten werden zunächst in der Eingangsstufe vorverarbeitet. Dadurch lassen sich die Datenpakete nach der Deserialisierung wiedergewinnen, die anschliessend in einem FIFO-Puffer¹ gespeichert werden. Die Wortbreite und etliche weitere Sensorparameter können vor der Firmware-Synthese angepasst werden, um die Module auch für künftige MIMOSA-Generationen kompatibel zu machen. Der Puffer kann dazu genutzt werden, die kurzzeitigen Strahlintensitätsschwankungen auszugleichen. Seine Grösse ist ebenfalls flexibel einstellbar und kann, in Anbetracht der zur Verfügung stehenden Hardware-Ressourcen, an die zeitliche Struktur der Schwankung angepasst werden.

Die implementierte Firmware ist in der Lage, den Datenfluss auf semantische und syntaktische Fehler zu überprüfen. Weiterhin wurden netzwerkabhängige Module entworfen, welche die Schnittstelle zum verwendeten Auslesenetzwerk in Anspruch nehmen. Damit lassen sich alle Basismodule für die unterschiedlichsten Netzwerke verwenden, ohne dass man in ihren Code eingreift und lediglich die Netzwerkschnittstelle austauscht. Alle Firmware-Module sind in einem Hauptmodul, der Readout-Chain (RC), zusammengefasst. Jede RC ist in der Lage genau einen MIMOSA Sensor auszulesen. Sie verfügt über ihr eigenes Steuerungsmodul, welcher den Datenfluss steuert und mit dem Auslesenetzwerk kommuniziert. Daher lassen sich RCs unabhängig voneinander betreiben. Dies wurde ausgenutzt um eine beliebige Zahl an Sensoren, welche mit einem Designparameter gesetzt ist, parallel auszulesen. Dadurch wird die Firmware skalierbar. Sie ist ausserdem auf einer Vielzahl an FPGA Typen und Modellen einsetzbar.

Nach der Implementierung und erfolgreichen Tests der Firmware, wurde nach einer geeigneten Hardware-Plattform gesucht, um mit dem Netzwerkaufbau fortzufahren. Um Zeit und Kosten zu sparen, wurden Hardware-Elemente aus verwandten Experimenten berücksichtigt. Eine Mehrzweckplatine bestückt mit FPGA Chips und optischen Links, welche im Rahmen des HADES Experiments entstanden ist, wurde schliesslich für den ROC verwendet. Das Trigger- und Readout-Board (TRB) des HADES Auslesenetzwerks verfügt über ausreichend Ressourcen und lässt sich mit zusätzlichen Add-On Karten erweitern, um eine Vielzahl an Aufgabenbereichen abzudecken. Es wird mittlerweile bei zahlreichen Experimenten eingesetzt, inklusive CBM und PANDA² am FAIR. Es wurden bisher drei Generationen der Platine erstellt, wobei die Versionen 2 und 3 im Rahmen dieser Arbeit verwendet wurden.

Die Firmware lässt sich auf beiden Platinen einbauen. Der auf TRBv2 basierte ROC ermöglicht die Integration von vier RCs pro Platine und ist damit für vier Sensoren ausgelegt. Das TRBv3 verfügt jedoch über einen Faktor von ~ 20 mehr an logischen Ressourcen und besitzt ausserdem vier optische Links mit jeweils einer Datenrate von 3 Gbit/s. Dementsprechend ist eine Implementierung des ROCs auf der TRBv3 Plattform für die nahe Zukunft vorgesehen. Es wurden bereits bis zu 12 RCs auf einem der fünf TRBv3 FPGAs implementiert, wobei aber im Rahmen dieser Arbeit nur wenige getestet werden konnten. Alle beschriebenen Netzwerktests wurden im Folgenden mit TRBv2 Platinen durchgeführt.

Die Sensordaten erreichen die ROCs über spezialisierte Front-End Karten, die parallel zu dieser Arbeit entworfen wurden. Sie sind speziell für MIMOSA-26 und MIMOSA-26AHR Sen-

¹FIFO – First In, First Out

²PANDA – antiProton ANnihilation at DArmstadt

soren ausgelegt. Die MIMOSA-26AHR Generation stellte während der Entstehung der vorliegenden Arbeit den, für unsere Anwendung, modernsten Stand der Technik dar. Die Sensoren erfüllen den Grossteil der Anforderungen an den final MVD Sensor und wurden im Rahmen eines ersten MVD-Prototypen, welcher im darauf folgenden Abschnitt beschrieben wird, ausgiebig getestet. Aus diesem Grund wurden die MIMOSA-26AHR Sensoren zum Aufbau des minimalen Netzwerks verwendet.

Schnelle Sensorauslese erfordert eine entsprechend schnelle und sichere Datenübertragung, mit dedizierten Werkzeugen zur Überwachung und Kontrolle aller Netzwerkkomponenten. Hierfür wurde das TrbNet Ausleseprotokoll vom HADES-Detektor ebenfalls übernommen. Das Protokoll ist vollständig getestet und in zahlreichen Strahlzeiten erfolgreich eingesetzt worden. Damit stellt es einen idealen Kandidaten für die ersten Netzwerkprototypen dar. Es erfordert allerdings ein zentrales Steuerungssystem. Zu diesem Zweck wurde ein weiteres FPGA-Firmware Modul implementiert, die Central Control Unit (CCU), mit der Aufgabe, die TrbNet-basierte Auslese zu steuern. Die CCU sendet Steuerungsbefehle an die ROCs, welche die vorverarbeiteten Daten über optische Links an die nächste Auslestufe, bestehend aus TrbNet Hubs, liefern. Die Hubs sind notwendig, um die Daten von allen ROCs im Netzwerk zu bündeln, in Events anzuordnen und per Ethernet zu einem nahe gelegenen Computer für die Analyse zur Verfügung zu stellen. In der anfänglichen Konfiguration erreichen die Hubs Datentransferraten von 400 Mbit/s pro Auslese-Link. Das Netzwerk ist dargestellt in Abbildung 8.

Die Sensoren wurden in dem TRB2-basierten Netzwerk mit einem gemeinsamen Takt angetrieben. Im Rahmen einer weiteren zugehörigen Arbeit wurde die Sensorprogrammierung über JTAG¹ implementiert. Alle Sensoren bekommen hierbei ein gemeinsames Startsignal nach der Programmierung. Die Sensoren sind identisch und verfügen über einen gleichschnellen Auslesezyklus, wodurch sie nach der Programmierung synchron arbeiten. Die ROCs wurden über zusätzliche LVDS² Leitungen ebenfalls global synchronisiert. Dadurch lassen sich die Sensordaten mit einem Zeitstempel versehen, um mögliche Abweichungen zur idealen Synchronisierung zu entdecken.

Insgesamt wurde ein Testsystem mit 12 MIMOSA-26AHR Sensoren aufgebaut. Dünne Flex-

¹JTAG – Joint Test Action Group

²LVDS – Low Voltage Differential Signaling

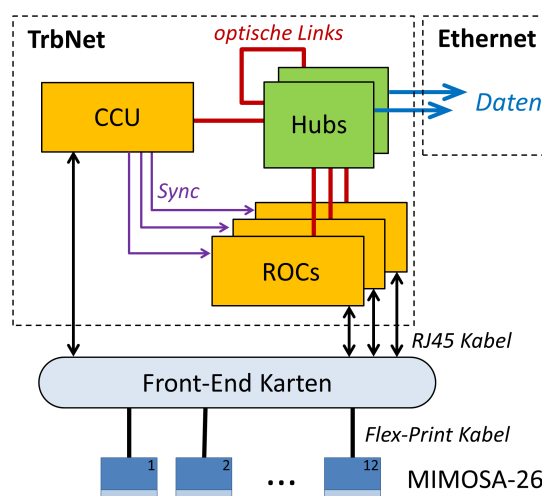


Abbildung 8: Das minimale Auslesenetzwerk.

Print Kabel haben die Sensoren mit der Front-End Elektronik verbunden. Die ausgelesenen Daten gelangen über LVDS-Leitungen zu den ROCs. Es wurden insgesamt drei TRBv2 Platinen hierfür verwendet, wobei jeder ROC vier Sensoren verarbeitet hat. Zwei TrbNet Hubs wurden eingesetzt, um die Daten mit bis zu 800 Mbit/s über das Ethernet zu extrahieren. Die CCU und die JTAG Programmierung sind über eine weitere TRBv2-Platine ermöglicht worden. Die so entstandene Netzwerkarchitektur stellt ein minimales Abbild des HADES-Netzwerks dar. Nach zahlreichen Labortests wurde es schliesslich in einer Strahlzeit eingesetzt und validiert.

Strahlzeit Tests und Ergebnisse

Um möglichst frühzeitig mit den notwendigen MVD Tests anzufangen, wurde ein Prototypmodul, basierend auf dem MIMOSA-26AHR Sensor, entworfen. Zwei Sensoren, angebracht jeweils auf der Vorder- und Rück-Seite einer dünnen CVD¹-Diamant Unterlagestruktur, wurden eingesetzt um ein Mikro-Tracking Modul zu realisieren, welcher die Materialbudget-Voraussetzungen für die erste MVD Station erfüllt. Die Eigenschaften des Prototyps bezüglich der Ortsauflösung und Effizienz der Teilchendetektion wurden in einer Strahlzeit am CERN² getestet. Es konnten eine Effizienz von über 99 % und eine Ortsauflösung von $\approx 3.6 \mu\text{m}$ nachgewiesen werden, womit die Anforderungen an einen MVD für CBM im Hinblick auf diese beiden Parameter erfüllt wurden.

Neben dem Prototypen, mussten vier weitere Referenzstationen für die Tests eingesetzt werden, die ebenfalls aus MIMOSA-26AHR Sensoren aufgebaut waren. Im Rahmen dieser Strahlzeit wurde ein vollwertiges Kontrollsystem implementiert, mit dem es möglich war, die Sensoren zu programmieren und alle Netzwerkkomponenten zu steuern und zu überwachen.

Das Auslesesystem hat alle Stabilitätstests erfolgreich bestanden. Im getesteten Zeitraum von fünf Tagen konnten die Daten nahezu ununterbrochen mit einer effektiven Datenrate von 6 – 25 MB/s gesammelt werden. Ein kleinerer Hardware-Defekt konnte direkt mit der Überwachungssoftware beobachtet werden und wurde innerhalb von wenigen Minuten behoben. Die Sensorsynchronisierung konnte durch die Zeitstempel und die rekonstruierten Teilchenspuren in allen analysierten Datensätzen bestätigt werden. Das Netzwerkverhalten nahe dem Bandbreitenlimit konnte ebenfalls studiert werden. Sobald die Ausgabe den Limit von 100 MB/s erreicht hat, wurden kontrollierte Totzeiten ausgelöst, in denen die Daten ohne Verlust von Synchronisierung fortwährend ausgelesen werden konnten.

Schlussfolgerung und Ausblick

Das vorgestellte, auf TRBv2 basierte Auslesesystem ermöglicht eine verlustfreie Auslese von mehreren, synchronisierten MIMOSA Sensoren. Aufgrund seiner klar-unterteilbaren Architektur ist es hoch-skalierbar. Weitere Sensoren können mit zusätzlichen ROCs verarbeitet werden. Ausserdem können weitere Hubs die Ausgabebandbreite beliebig erhöhen, ohne dass sich die Latenzzeiten nennenswert ändern. Deshalb wäre das System theoretisch auch in der Lage, die finalen SIS-100 Experimente zu unterstützen. Allerdings ist die Zahl der erforderlichen Hardware-Platinen viel zu hoch. Deshalb soll langfristig die leistungsfähigere TRBv3-Plattform zum Einsatz kommen.

¹CVD – Chemical Vapor Deposition

²CERN – Conseil Européen pour la Recherche Nucléaire

Vorläufige Tests haben demonstriert, dass die TRBv3-Ressourcen eine wesentlich höhere Anzahl an Readout-Chains pro Platine erlauben. Bis zu 48 ROCs wurden auf dem neuen ROC integriert. Aufgrund fehlender Add-On Karten konnten jedoch nur wenige davon getestet werden, was sich in baldiger Zukunft ändern wird. Unter der Annahme, dass die optischen Links der TRBv3-Platinen über dieselbe Leistung verfügen wie bei der Strahlzeit, könnte man bereits mit 14 solchen ROCs die höchsten simulierten Datenraten von 22 Gbit/s auslesen. Der in die FPGA-Firmware integrierte FIFO-Puffer könnte hierbei verwendet werden, die Strahlintensitätsschwankungen auszugleichen. Dadurch würden die 50 Gbit/s Worst-Case Eingangsdaten auf einen zeitlichen Mittelwert von 22 Gbit/s ausgeglichen werden. Dies setzt aber eine genaue Vermessung der zeitlichen Struktur des Strahls voraus. Falls der FIFO-Puffer zu klein ausfallen sollte, würde die Zahl der eingesetzten TRBv3-Platinen auf ~ 30 wachsen. Um die Auslesebandbreite am effizientesten auszunutzen, müssten die Sensoren ungleichmässig an die ROCs verteilt werden. Aufgrund der anisotropen Trefferverteilung ist es hilfreich, die Vielzahl der am wenigsten beanspruchten Sensoren von einer kleinen Anzahl an ROCs auszulesen. Damit wäre es möglich die durchschnittliche Ausgangsdatenrate für alle ROCs über die Zeit konstant zu halten. Dies erfordert allerdings eine erhöhte Zahl an Eingangskanälen bei diversen ROCs. Eine Möglichkeit dies zu realisieren wäre, die TRBv3 ROCs mit zusätzlichen Add-On Karten auszustatten.

Die globale Synchronisierung der ~ 300 geplanten Sensoren erfordert ebenfalls leichte Modifikationen des minimalen Testnetzwerks. Im finalen Experiment können die Sensoren nicht alle effizient aus einer Quelle, der CCU, gestartet und angetrieben werden. Dies würde zu viele Ausgangsleitungen benötigen und die Signallaufzeiten wären ebenfalls nicht optimal. Demnach sollte die Sensor-Programmierung und -Steuerung in die ROCs verlagert werden. Das setzt aber voraus, dass auf allen ROCs die globale Hauptclock jitterfrei wiederhergestellt werden kann. Demnach wird es erforderlich sein, ein synchrones Netzwerkprotokoll zu verwenden, wie z.B. das CBMnet, oder das erweiterte, synchrone TrbNet, welche zurzeit auf den TRBv3-Platinen entwickelt werden. Das modulare Firmware-Design erleichtert dabei erheblich die Modifikationen.

Das MVD Auslesesystem steht bereit, in das globale Ausleseschema von CBM integriert zu werden. Die ROC-Firmware ist kompatibel zu weiteren, für CBM relevanten Netzwerkprotokollen, die sich zurzeit noch in Entwicklung befinden. Nach genauer Spezifikation lassen sie sich nahezu mühelos in die Firmware integrieren. Sobald die finale MIMOSA-Generation zur Verfügung steht, wäre man in der Lage, ein grossflächiges Netzwerk basierend auf den TRBv3-Platinen aufzubauen, welches laut den Simulationen über genug Leistung für die anspruchsvollsten SIS-100 Experimente verfügt. Dies setzt allerdings die Implementierung geeigneter Front-End und Add-On Karten voraus, welche zurzeit in einer zugehörigen, anderen Arbeit entwickelt werden. Selbst bei den grössten Strahlintensitätsschwankungen, wie sie bei SIS-18 auftraten, wäre das System in der Lage ohne Unterbrechung die MIMOSA Daten vorverarbeitet an die höheren Auslestufen weiterzuleiten.

Appendix A

Correlated Double Sampling with Clamping

Assuming that two electronic circuits are coupled by a capacitor in accordance with Fig. A.1, as is the case within the MIMOSA-26 pixel. Then, the noise in the first circuit, expressed by U_{in} , directly influences the voltage provided to the second circuit (U_{out}). In the most common case, the noise contains a constant offset voltage (noise pedestal) and a time dependent part. Moreover, considering MIMOSA-26, the noise can arbitrarily change from pixel to pixel making the readout highly inaccurate.

However, by closing the switch after the capacitance for a short duration, the output voltage U_{out} gets shifted to the value of U_{clamp} , regardless of any noise which was previously present in the first circuit. By opening the switch again, the output potential stays approximately at the same value. Merely the temporal noise is now affecting the second circuit. The noise pedestal has been shifted to a well known value of U_{clamp} . This allows for the removal of the noise pedestal in all MIMOSA-26 pixels simultaneously. Upon subtracting U_{clamp} in a subsequent readout step, the clear output signal can be recovered¹.




Figure A.1: By applying the CDS method with clamping, the random noise pedestal which is present in circuit 1 can be removed to obtain a clearer output signal.

¹The signal is still affected by temporal and fixed-pattern noise.

Appendix B

Reduced MIMOSA-26 Bonding Scheme

MIMOSA-26 sensors do not require all pads to be bonded. Moreover, all pads are doubled in order to account for mechanical damage and errors during production. The following tables summarize all of the pads which were bonded for the MIMOSA-26 tests. There, GNDA and VDDA refer to the analog bias voltage, GND and VDD to the digital bias voltage, GND_Signal to the ground sensing line provided by the sensor, and Data to all other signals, e.g. JTAG, clamping, etc. If reference voltage monitoring is not needed, pads 3 – 17 can also be excluded.

Pad	Description	GNDA	VDDA	GND	VDD	GND_Signal	Data
1	Temperature						X
3	VDiscriRef1A						X
5	VDiscriRef1B						X
7	VDiscriRef1C						X
9	VDiscriRef1D						X
11	VDiscriRef2A						X
13	VDiscriRef2B						X
15	VDiscriRef2C						X
17	VDiscriRef2D						X
23	vdda		X				
24	vdda		X				
26	gnda	X					
27	gnda	X					
34	TMS						X
36	TDI						X
38	TCK						X
40	TDO						X
46	gnd			X			
49	vdd				X		
56	gnd			X			
58	CLKL_p						X
59	CLKL_n						X
61	vdd				X		

Table B.1: The reduced MIMOSA-26 bonding scheme part 1, pads 1 - 62.

APPENDIX B. REDUCED MIMOSA-26 BONDING SCHEME

Pad	Description	GND_A	VDD_A	GND	VDD	GND_Signal	Data
65	START						X
67	RSTB						X
69	gnd			X			
71	vdd				X		
76	vdd_latch				X		
79	gnd_latch			X			
82	v_clp						X
83	v_clp						X
87	gnda	X					
89	gnda	X					
91	gnda	X					
93	gnda	X					
96	vdda		X				
98	vdda		X				
100	vdda		X				
102	vdda		X				
108	gnd_mem			X			
112	vdd_mem				X		
120	DO1_n						X
121	DO1_p						X
123	gnd					X	
125	DO0_n						X
126	DO0_p						X
128	vdd				X		
133	gnd					X	
135	CLKD_n						X
136	CLKD_p						X
143	vdd				X		
176	vdd_latch				X		
180	gnd_latch			X			
188	gnda	X					
189	gnda	X					
192	vdda		X				
193	vdda		X				

Table B.2: The reduced MIMOSA-26 bonding scheme part 2, pads 63 - 234.

Appendix C

Proposal of the Microslice-Container Encoding for the MVD

The timeslice containers are produced within the FLES, however the microslice containers (MCs) should be managed by the subdetectors. Hence, the MVD is responsible for its own data format. After the generic MC header, all the relevant particle hits in form of clusters should be provided by the MVD. One possibility considers a cluster encoding based on the *center-of-gravity*. Thereby, the cluster shape is analyzed within the ROC readout chain, and the most-central pixel is selected. Its (x, y) coordinate is then transmitted, as provided by the corresponding sensor. A value of 16 bits for x and y are reserved, respectively. However, this number can be reduced in future, when the number of pixels per sensor is known. Additional n bits containing the cluster shape are foreseen. The code is obtained from a look-up table, which features 2^n most probable shapes. Thus for $n = 10$, the 1024 most frequently occurring shapes can be specified, regardless of the cluster size. However, if a cluster shape can not be found within the table, the ROC creates an exception and sends all the pixels one-by-one.

The minimum MC data packet size should therefore amount to $16 + 16 + 10 = 42$ bits, considering x , y and n . Additionally, the unique sensor ID needs to be provided, as well. Thus, a minimum size of 64 bits per cluster is proposed which should suffice for all the additional information. A designated bit within the package marks the exception. If an exception occurs, an additional packet needs to be transmitted containing two values. The first value denotes the cluster window in the $n \times m$ format, where n and m are 16 bit integers. The second entry contains the actual $n \cdot m$ binary values of pixels within that window. The proposed encoding is shown in Fig. C.1.

An important MVD specification which differs from other CBM subsystems is the unusually low time resolution which depends on the sensor integration time. In following, an integration time of $30 \mu\text{s}$ is assumed. The MVD data is read out frame-wise, with a constant latency of one frame, i.e. $30 \mu\text{s}$. Under the assumption that an MC is requested every $1 \mu\text{s}$ by the FLES, all subsystems need to buffer their MCs in the FLIBs until the MVD provides its data. However, once the MVD sensors provide their frame data, the readout time of each cluster can be determined and the cluster packed into the corresponding MC. Hence, the MC allocation is resolved on the cluster level and is compatible with the $1 \mu\text{s}$ time resolution, since a pixel-row is processed within 100 - 200 ns (MIMOSA-26 and MIMOSA-28 sensors). Thus, the MVD MCs will be delayed by a factor of 30, but then all of them will reach the FLIBs and the interval building




Figure C.1: A proposal of an MC datagram for the MVD. Four packets are produced by three clusters in this example. Common cluster shapes are encoded within 64-bit packets, where x and y denote pixel coordinates of the center-of-gravity, 'shape' the cluster shape-code, 'misc' additional information and e an exception. In case of an exception (third cluster), the cluster shape is uncommon and requires an additional data package with variable size to encode its shape.

can proceed. However, the clusters have additionally a large time uncertainty in the order of the integration time. After a row is processed by the rolling shutter operation, a new particle hit can occur at any time. If a particle hits the row again immediately after it is read out, the next readout is scheduled after $30 \mu s$ of integration time. This needs to be handled properly by FLES during the event reconstruction.

Appendix D

Readout Controller Firmware Parameters

Following list summarizes all the VHDL design parameters that can be used to reconfigure the readout chain:

- **Number of sensors (*c_MAPS_QUANTITY*)** - Sets the number of parallel readout chains from Fig. 5.2. Each chain is responsible for one sensor.
- **Number of sensor output channels (*c_MAPS_CHANNELS*)** - The readout chain can be prepared for a different number of sensor output channels with this parameter.
- **Sensor package size (*c_MAPS_DATASIZE*)** - The internal sensor word size affects nearly all readout chain components. Common values are 16 (SUZE-1) and 30 (SUZE-2).
- **Frame readout time (*c_MAPS_FRAMETIME*)** - This constant, which denotes the sensor integration time, is used for timeout counters. If the frame is deserialized longer than its actual integration time, an error is signaled and the frame acquisition is stopped to avoid deadlocks.
- **Maximum frame packages (*c_MAPS_DATAPACKS*)** - Sets the maximum number of packages delivered by the sensor in one frame per channel (e.g. 574 for MIMOSA-26). If this number is exceeded during readout, it signals an error and resets a state-machine.
- **Header bit pattern (*c_HEADER*)** - The unique bit pattern which marks the start of a new frame, as set during the sensor programming. The number is expressed per output channel in sensor word size format. For MIMOSA-26, a 16 bit word 0x5555 in hex can be used.
- **Trailer bit pattern (*c_TRAILER*)** - Same as the Header, only the Trailer marks the end of all the data in one frame. In case of MIMOSA-26, 0x8001 in hex is used.
- **Underlying FPGA architecture (*c_ARCHITECTURE*)** - Many components use IP cores which are generated during the development on one particular FPGA architecture. Changing this constant enables using the architecture-specific IP cores, e.g. FIFO buffers. Currently only Virtex4 and ECP3 types are supported.

An example for MIMOSA-26 is given in table D.1.

APPENDIX D. READOUT CONTROLLER FIRMWARE PARAMETERS

Design Constant	Value
c_MAPS_QUANTITY	2 – 4
c_MAPS_CHANNELS	2
c_MAPS_DATASIZE	16
c_MAPS_FRAMETIME	11520
c_MAPS_DATAPACKS	574
c_HEADER	0x5555
c_TRAILER	0x8001
c_ARCHITECTURE	_virtex4

Table D.1: Following readout chain settings are compatible with MIMOSA-26 sensors.

Appendix E

Readout Controller Error- and Status-Bits

E.1 Data Checker

Bit	Type	Description
0	Error	Data handler error (a timeouts occurred)
1	Error	Received package is not a Header when expected
2	Error	Frame number is not in the ascending order
3	Error	Datalengths are not same on both channels
4	Error	Datalength is larger than 570 per channel
5	Error	Data counter is 0, but the 'state' counter is not 0
6	Error	Data counter is not 0 on Trailer package
7	Error	Data counter turned 0 during normal package readout
8	Error	Number of 'states' is not between 1 and 9
9	Error	Matrix row address is larger than 575
10	Error	Overflow bit is set, but less than 9 states are present
11	Error	Matrix column address is larger than 1151
12	Error	Row address inconsistent (row is lower than the one before)
13	Error	Column address inconsistent (column is lower than the one before)
14	Error	State counter is not 1 when expected
15	Error	Wrong column address on channel 2
16 - 29	None	Reserved
30	Status	Trailer detected, going back to IDLE state
31	Status	Frame extraction successful

Table E.1: Error bits set by the data checker for the MIMOSA-26 sensor. The sensor features 576 rows, 1152 columns and a maximum of 570 data packages per channel.

E.2 Chain Controller

Bit	Type	Description
0	Status	Sensor was all the time active during readout
1	Status	Sensor is currently active
2	Status	The TrbNet buffers are OK
3	Status	Frame complete (data checker status bit 31)
4	Error	Data handler error occurred (data checker 0)
5	Error	Data handler - dataready timeout error
6	Error	Data handler - Trailer timeout error
7	Error	Reset detected (two Header packages and no Trailer between them)
8	Error	Received package is not a Header when expected (data checker 1)
9	Error	Frame number is not in the ascending order (data checker 2)
10	Error	Datalengths are not same on both channels (data checker 3)
11	Error	Datalength is larger than 570 per channel (data checker 4)
12	Error	Data counter is 0, but the 'state' counter is not 0 (data checker 5)
13	Error	Data counter is not 0 on Trailer package (data checker 6)
14	Error	Data counter turned 0 during normal package readout (data checker 7)
15	Error	Number of states is not between 1 and 9 (data checker 8)
16	Error	Matrix row address is larger than 575 (data checker 9)
17	Error	Overflow bit is set, but less than 9 states are present (data checker 10)
18	Error	Matrix column address is larger than 1151 (data checker 11)
19	Error	Row address inconsistent (data checker 12)
20	Error	Column address inconsistent (data checker 13)
21	Error	State counter is not 1 when expected (data checker 14)
22	Error	Wrong column address on channel 2 (data checker 15)
23	Error	Formatter datalength error - FIFO empty
24	Error	BUFFERS_STOP error
25	Error	Header timeout error
26	Error	Frame counter overflow (internal error)
27	Error	Frame counter underflow (internal error)
28	Status	Reserved (set to 1)
29	Status	Reserved (set to 1)
30	Status	Trailer detected, data checker is IDLE (data checker 30)
31	Status	Status handler has prepared a status

Table E.2: The FRAME_STATUS bits delivered by the chain controller. They are written together with the data in the format header. The word 0xF00000F marks a successful frame extraction without errors.

Appendix F

CBMnet User Interface

The CBMnet user interface is shown in Fig. F.1. All messages and signals need to be provided with the recovered clock (link_clk). Thus, a FIFO buffer may be used to synchronize the user logic with the CBMnet module if they are needed to run asynchronously. The user module may send and receive DCM, DTM and DLM. In case of a DLM, only the DLM type is communicated as a 4-bit code. The DCM and DTM are transferred as 16-bit packets. First data packet requires a 'valid' bit set (e.g. data2send_start), and the last packet requires the 'stop' bit (data2send_end). This mechanism is ideally suited for free-streaming systems. In case that no further messages can be processed, e.g. due to full buffers, the 'stop' signal in backward direction can be set (data2send_stop). This signal affects only following messages, after the current message transfer is finished. All further encoding and link control is performed automatically within the CBMnet module.

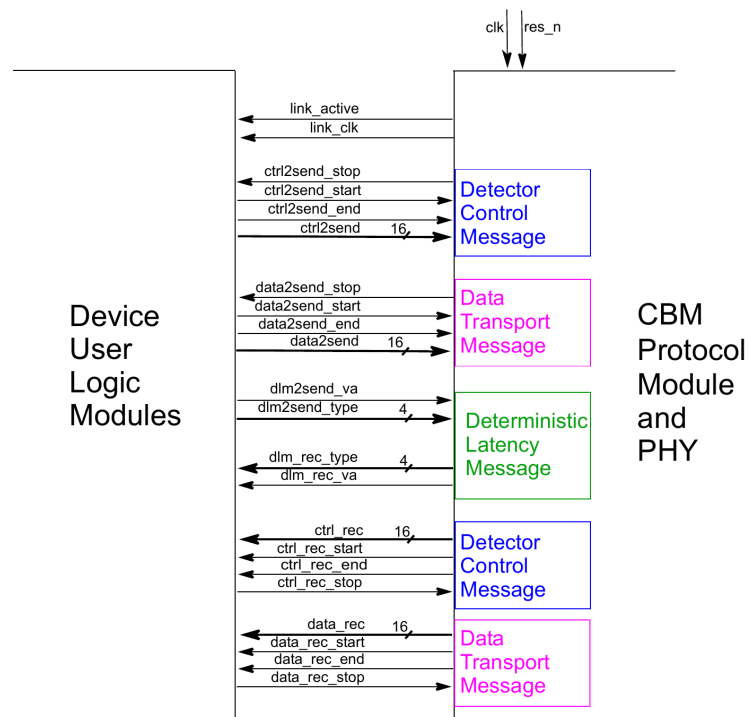


Figure F.1: The CBMnet user interface. Source: [45].

Appendix G

MIMOSA-26 Front-End Electronics Connectors

The LVDS output of the RJ45 connectors on the converter board (CB) and JTAG queue board (QB) is not set according to the Ethernet standard. Therefore, some data lines might experience more noise during the transmission when using conventional Ethernet cables. In order to optimize the data quality, pins 4 and 6 on one of the Ethernet cable connectors need to be interchanged. Only such specialized cables are used in combination with the FEE prototypes. The modifications are outlined in Fig. G.1. The pin assignments of CB and QB are shown in Fig. G.2 and Fig. G.3, respectively.




Figure G.1: One connector of the applied Ethernet cables needs to be modified by interchanging pins 4 and 6.




Figure G.2: The pin assignment of the converter board.




Figure G.3: The pin assignment of the JTAG queue board.

Appendix H

Grounding Scheme of the Readout Network Prototype

The TRB2 boards require a 48 V input voltage which they transform over internal DC/DC converters to their needs. The DC/DC converters generate a ground level which is then, by design, connected with the TRB2 shielding. Over the shielding, the ground reaches the add-on boards. Thus, all Ethernet and SCSI cables connected to the add-ons carry the TRB2 ground on their shielding, as well. Particularly, the SCSI cable of the GP add-on carries the ground to the patch panels which then transmit it over the RJ45 sockets to the FEE. The cable is necessary to provide sufficient LVDS pairs for sensor programming and can not be circumvented. Nevertheless, by connecting the shielding on the FEE with the ground formed by the plus-pole of their power supply, the entire setup would be driven with only one floating ground. However, the FEE require power supplies of their own (with a 5 V input). Their ground level is therefore different than the one generated by the TRB2 DC/DC converters. Hence, both grounds need to be connected together. This is performed by connecting integrated TRB2 ground pins with the plus-pole of the FEE power supplies, as well. The ground loops in the network are herewith minimized. Additionally, by connecting the plus-pole to the mass the entire setup can be fixed to mass in one place. This leads to a grounding scheme developed by the IKF electronics department shown in Fig. H.1.

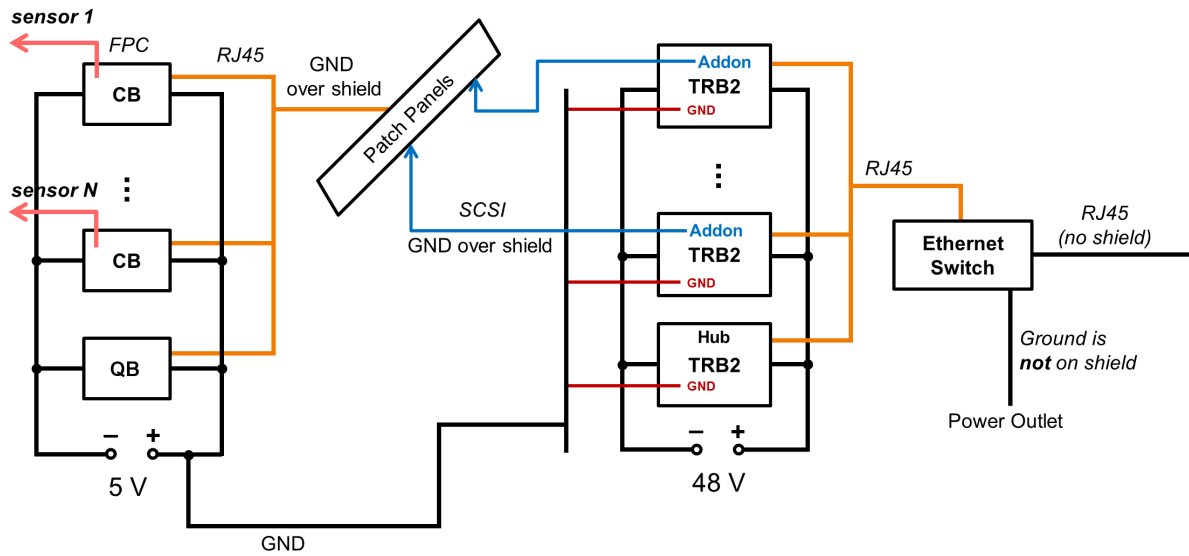


Figure H.1: The grounding scheme of the prototype readout network. The basic idea is to connect the FPGA boards and the FEE (CB/QB) with a common ground to avoid major ground loops affecting the sensor performance. The TRB2 ground is carried over to the FEE via the SCSI and Ethernet cable shielding. Therefore by connecting the FEE ground with the shielding and by attaching the FEE power supply ground (plus-pole) to the TRB2 grounding pins, the entire setup can be fixed. All the colored paths (blue, orange and red) carry the same ground level. The Ethernet switch is decoupled from all external ground sources.

Appendix I

Software Tools for Sensor Threshold Settings

This section describes the customized software tools developed explicitly for sensor threshold fine-tuning.

I.1 Threshold GUI

The first tool, named the threshold GUI, allows the user to modify the threshold settings of a particular sensor. V_{Ref2} is the pedestal corresponding to V_{Ref2} from section 3.2.3 and will be automatically subtracted from V_{Ref1} of the corresponding comparator block (i.e. from VREF1A to VREF1D) during the CDS procedure. Then, the frame data can be acquired for a user-specified duration with the result displayed on the screen. The pixel matrix is visualized with all the hits recorded during the specified time. Beneath each discriminator block, the normalized distribution denoting the number of fired pixels per frame is given. Without the beam or a source, this number is equal to the sensor fake-hit rate. This allows fine-tuning the thresholds to the desired FHR and also detecting some pixel matrix errors, e.g. dead pixels or dead comparators. The tool is shown in Fig. I.1.

I.2 Threshold Finder

The second tool is programmed in PERL and uses a binary search algorithm to automatically find the optimal threshold values for the four comparator blocks. The desired fake-hit rate per block is given as an input parameter. The tool then starts adapting the threshold settings.

It begins with the comparator block A and afterwards applies the same algorithm for the other blocks, as well. The tool operates on some given range of possible values which is initially 0-255. It takes the middle of the search interval as a reference value in each iteration, e.g. 127 for the first iteration. Then, some data is taken for approximately 0.1 seconds, which yields ≈ 870 frames with current threshold settings. The tool normalizes the average number of hits for the discriminator block under study and checks subsequently whether the calculated value is greater or smaller than the user-specified input parameter. If the numbers differ, the tool updates the boundaries of its search interval accordingly and repeats all the steps in the following iteration. The algorithm automatically finds the approximate threshold within eight steps per

discriminator block. However, if the matrix contains errors, e.g. dead pixels, the algorithm will not work accurately. In some cases, the threshold finder can not find the optimal value due to some improper interval formation, in which case the optimal threshold is by 1 higher or lower than the result. Therefore, it is always advised to use the threshold GUI after the threshold finder to cross-check the result. In practice, the threshold finder needs less than two minutes per sensor and finds, in general, quickly a suitable solution.

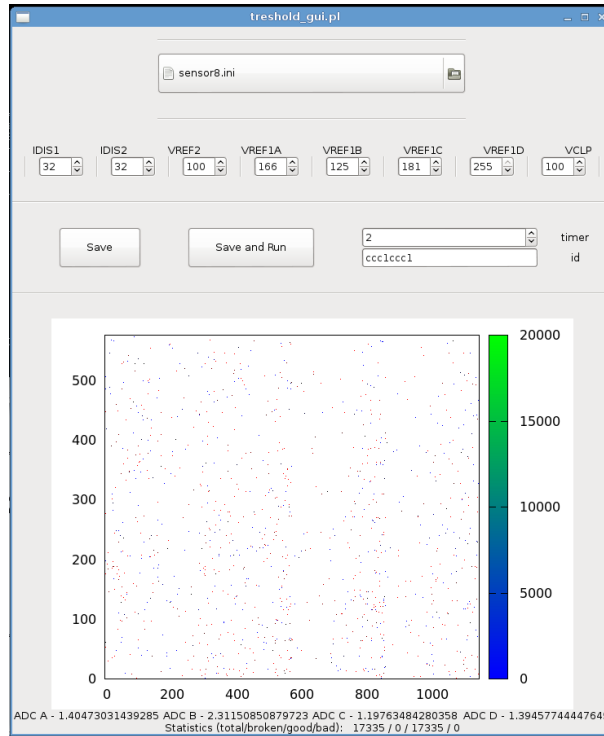


Figure I.1: The threshold GUI allows setting individual sensor thresholds and examining the matrix for errors.

Appendix J

HADES Event File Structure

One event stored in the file contains one or more subevents depending on the number of Trb-Net Hubs used for the readout. Each subevent contains all the sensor frames merged by the corresponding Hub. The data blocks contain unprocessed sensor data, since no pre-processing algorithms were used during the beam time. In the applied prototype setup, 12 frames provided at the same time by all 12 sensors are recorded per event. Additionally, the padding at the end of the data stream corrects the package alignment to 64 bit.




Figure J.1: The structure of the HADES event file. Source: [111].

Bibliography

- [1] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory*. Westview Press, 1995
- [2] B. Friman, C. Höehne, J. Knoll, S. Leupold, J. Randrup, R. Rapp, P. Senger, et al., *The CBM Physics Book*. Springer Press, 2011
- [3] L. McLerran, “RHIC physics: The Quark gluon plasma and the color glass condensate: Four lectures”, [arXiv:hep-ph/0311028](https://arxiv.org/abs/hep-ph/0311028)
- [4] A. Rossi, “Heavy-flavour and quarkonia in heavy-ion collisions”, [arXiv:1308.2973](https://arxiv.org/abs/1308.2973)
- [5] E. Bratkovskaya, W. Cassing, V. Konchakovski, O. Linnyk, V. Ozvenchuk, et al., “The QGP dynamics in relativistic heavy-ion collisions”, *Journal of Physics: Conference Series Vol. 455* (2013), [arXiv:1304.4115](https://arxiv.org/abs/1304.4115)
- [6] ALICE Collaboration, B. Abelev, et al., “Suppression of high transverse momentum D mesons in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Journal of High Energy Physics Vol. 2012, Issue 9* (2012), [arXiv:1203.2160](https://arxiv.org/abs/1203.2160)
- [7] STAR Collaboration, J. Adams, et al., “Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR Collaboration’s critical assessment of the evidence from RHIC collisions”, *Nuclear Physics A Vol. 757* (2005) pp. 102–183, [arXiv:nucl-ex/0501009](https://arxiv.org/abs/nucl-ex/0501009)
- [8] PHENIX Collaboration, K. Adcox, et al., “Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration”, *Nuclear Physics A Vol. 757* (2005) pp. 184–283, [arXiv:nucl-ex/0410003](https://arxiv.org/abs/nucl-ex/0410003)
- [9] PHOBOS Collaboration, B. Back, et al., “The PHOBOS perspective on discoveries at RHIC”, *Nuclear Physics A Vol. 757* (2005) pp. 28–101, [arXiv:nucl-ex/0410022](https://arxiv.org/abs/nucl-ex/0410022)
- [10] BRAHMS Collaboration, I. Arsene, et al., “Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment”, *Nuclear Physics A Vol. 757* (2005) pp. 1–27, [arXiv:nucl-ex/0410020](https://arxiv.org/abs/nucl-ex/0410020)
- [11] I. Selyuzhenkov, “Recent experimental results from the relativistic heavy-ion collisions at LHC and RHIC”, [arXiv:1109.1654](https://arxiv.org/abs/1109.1654)
- [12] B. Mohanty, “Exploring the QCD phase diagram through high energy nuclear collisions: An overview”, [arXiv:1308.3328](https://arxiv.org/abs/1308.3328)

- [13] W. Florkowski, *Phenomenology of Ultra-Relativistic Heavy-Ion Collisions*. World Scientific Publishing, 2010
- [14] T. Matsui and H. Satz, “ J/ψ suppression by quark-gluon plasma formation”, *Physics Letters B Vol. 178, Issue 4* (1986) pp. 416–422
- [15] M. Tannenbaum, “Highlights from BNL-RHIC-2012”, [arXiv:1302.1833](https://arxiv.org/abs/1302.1833)
- [16] M. Spousta, “Jet Quenching at LHC”, *Modern Physics Letters A Vol. 28* (2013) , [arXiv:1305.6400](https://arxiv.org/abs/1305.6400)
- [17] K. Fukushima and T. Hatsuda, “The phase diagram of dense QCD”, *Reports on Progress in Physics, Vol. 74* (2011)
- [18] D. Teaney, “Viscous Hydrodynamics and the Quark Gluon Plasma”, [arXiv:0905.2433](https://arxiv.org/abs/0905.2433)
- [19] P. Petreczky, “Lattice QCD at non-zero temperature”, *Journal of Physics G Vol. 39* (2012) , [arXiv:1203.5320](https://arxiv.org/abs/1203.5320)
- [20] J. Randrup and J. Cleymans, “Optimal conditions for exploring high-density baryonic matter”, [arXiv:hep-ph/0607065](https://arxiv.org/abs/hep-ph/0607065)
- [21] M. Stephanov, K. Rajagopal, and E. Shuryak, “Signatures of the tricritical point in QCD”, *Physical Review Letters Vol. 81* (1998) pp. 4816–4819, [arXiv:hep-ph/9806219](https://arxiv.org/abs/hep-ph/9806219)
- [22] FAIR Homepage. <http://www.gsi.de/fair>
- [23] FAIR Green Paper, October, 2009. <http://www.fair-center.de/de/fuer-nutzer/publikationen/fair-publikationen.html>
- [24] FAIR - Baseline Technical Report, Volume 1, September, 2006. <http://www.fair-center.de/de/fuer-nutzer/publikationen/fair-publikationen.html>
- [25] HADES Collaboration, G. Agakishiev, et al., “The High-Acceptance Dielectron Spectrometer HADES”, *The European Physical Journal A Vol. 41, Issue 2* (2009) pp. 243–277, [arXiv:0902.3478](https://arxiv.org/abs/0902.3478)
- [26] HADES Homepage. <http://www-hades.gsi.de>
- [27] K. Lapidus et al., “The HADES-at-FAIR project”, *Physics of Atomic Nuclei Vol. 75, Issue 5* (2012) pp. 589–593
- [28] O. Linnyk, E. L. Bratkovskaya, and W. Cassing, “Open and hidden charm in proton-nucleus and heavy-ion collisions”, *International Journal of Modern Physics E, Vol. 17, Issue 8* (2008)
- [29] J. Heuser, “The Compressed Baryonic Matter Experiment at FAIR”, *Nuclear Physics A Vol. 904-905* (2013) pp. 941c–944c. The Quark Matter 2012 Proceedings of the XXIII International Conference on Ultrarelativistic Nucleus-Nucleus Collisions

- [30] M. Gazdzicki, “NA49/NA61: results and plans on beam energy and system size scan at the CERN SPS”, *Journal of Physics G Vol. 38* (2011), [arXiv:1107.2345](#)
- [31] M. Gazdzicki, M. Gorenstein, and P. Seyboth, “Onset of deconfinement in nucleus-nucleus collisions: Review for pedestrians and experts”, *Acta Physica Polonica B Vol. 42* (2011) pp. 307–351, [arXiv:1006.1765](#)
- [32] P. Senger, “CBM/FAIR capabilities for charm and dilepton studies”, 2009. <https://www-alt.gsi.de/documents/DOC-2009-Nov-34.html>
- [33] J. Eschke, “Compressed baryonic matter experiment at FAIR”, *EPJ Web of Conferences Vol. 20* (2012)
- [34] V. Friese, “Computational Challenges for the CBM Experiment”, in *Mathematical Modeling and Computational Science*, vol. 7125 of *Lecture Notes in Computer Science*, pp. 17–27. Springer Press, 2012
- [35] P. Senger and V. Friese, “Nuclear Matter Physics at SIS-100”, 2012. <https://www-alt.gsi.de/documents/DOC-2011-Aug-29.html>
- [36] V. Friese, C. Sturm, et al., “CBM Progress Report”, 2010. <https://www-alt.gsi.de/documents/DOC-2011-Mar-235.html>
- [37] V. Friese, C. Sturm, et al., “CBM Progress Report”, 2011. <https://www-alt.gsi.de/documents/DOC-2012-Mar-33.html>
- [38] V. Friese, C. Sturm, et al., “CBM Progress Report”, 2012. <https://www-alt.gsi.de/documents/DOC-2013-Mar-49.html>
- [39] V. Friese, C. Sturm, et al., “CBM Progress Report”, 2013. <https://www-alt.gsi.de/documents/FOLDER-9871395905824.html>
- [40] A. Lebedev, C. Höehne, I. Kisel, and G. Ososkov, “Track reconstruction algorithms for the CBM experiment at FAIR”, *Journal of Physics: Conference Series Vol. 219* (2010)
- [41] I. Kisel, “Event reconstruction in the CBM experiment”, *Nuclear Instruments and Methods in Physics Research Section A Vol. 566* (2006) pp. 85–88. Proceedings of the 1st Workshop on Tracking in High Multiplicity Environments
- [42] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nuclear Instruments and Methods in Physics Research Section A Vol. 262* (1987) pp. 444–450
- [43] J. de Cuveland and V. Lindenstruth, “A First-level Event Selector for the CBM Experiment at FAIR”, *Journal of Physics: Conference Series Vol. 331* (2011)
- [44] F. Lemke, D. Slogsnat, N. Burkhardt, and U. Bruening, “A Unified DAQ Interconnection Network With Precise Time Synchronization”, *IEEE Transactions on Nuclear Science Vol. 57, Issue 2* (2010) pp. 412–418
- [45] F. Lemke, *Unified Synchronized Data Acquisition Networks*. PhD thesis, University of Mannheim, 2012

- [46] n-XYTER - Read-out ASIC for High Resolution Time and Amplitude Measurements. <http://cbm-wiki.gsi.de/cgi-bin/view/Public/PublicNxyter>
- [47] M. Deveaux, *Development of fast and radiation hard Monolithic Active Pixel Sensors (MAPS) optimized for open charm meson detection with the CBM - vertex detector*. PhD thesis, University Frankfurt, 2007
- [48] M. Deveaux et al., “Challenges with decay vertex detection in CBM using an ultra-thin pixel detector system linked with the silicon tracker”, *Proceedings of 17th International Workshop on Vertex detectors* (2008) , [arXiv:0906.1301](https://arxiv.org/abs/0906.1301)
- [49] M. Deveaux et al., “Status of the Micro Vertex Detector of the Compressed Baryonic Matter Experiment”, *Proceedings of XLVIII International Winter Meeting on Nuclear Physics, Bormio (Italy)* (2010) . <https://www-alt.gsi.de/documents/DOC-2011-Jan-5.html>
- [50] Particle Data Group. <http://pdg.lbl.gov/>
- [51] C. Dritsa, *Design of the Micro Vertex Detector of the CBM experiment: Development of a detector response model and feasibility studies of open charm measurement*. PhD thesis, University of Frankfurt, 2011
- [52] N. Weste and K. Eshraghian, *Principles of CMOS VLSI design*. Addison-Wesley, 1994
- [53] K. Waldschmidt, “Computer Technology”, 2007. Lecture Notes
- [54] G. Lutz, *Semiconductor Radiation Detectors*. Springer Press, 2001
- [55] G. Moore, “Cramming More Components Onto Integrated Circuits”, *Proceedings of the IEEE Vol. 86, Issue 1* (1998) pp. 82–85
- [56] L. Rossi, P. Fischer, T. Rohe, and N. Wermes, *Pixel Detectors - From Fundamentals to Applications*. Springer Press, 2006
- [57] A. Holmes-Siedle and L. Adams, *Handbook of radiation effects*. Oxford University Press, 2007
- [58] D. Doering, “Eine Ausheilstudie an bestrahlten Monolithic Active Pixel Sensoren”, Master’s thesis, University of Frankfurt, 2010
- [59] L. Braga, S. Domingues, M. Rocha, L. Sá, F. Campos, F. Santos, A. Mesquita, M. Silva, and J. Swart, “Layout techniques for radiation hardening of standard CMOS active pixel sensors”, *Analog Integrated Circuits and Signal Processing Vol 57* (2008) pp. 129–139
- [60] H. Hughes and J. Benedetto, “Radiation effects and hardening of MOS technology: devices and circuits”, *IEEE Transactions on Nuclear Science Vol. 50, Issue 3* (2003) pp. 500–521
- [61] D. Doering, M. Deveaux, M. Domachowski, C. Dritsa, I. Fröhlich, M. Koziel, C. Muentz, S. Ottersbach, F. Wagner, and J. Stroth, “Annealing studies on X-ray and neutron irradiated CMOS Monolithic Active Pixel Sensors”, *Nuclear Instruments and Methods in Physics Research Section A Vol. 658, Issue 1* (2011) pp. 133–136. RESMDD 2010

- [62] D. Doering, M. Deveaux, M. Domachowski, I. Fröhlich, M. Koziel, C. Müntz, P. Scharrer, and J. Stroth, “Pitch dependence of the tolerance of CMOS monolithic active pixel sensors to non-ionizing radiation”, *Nuclear Instruments and Methods in Physics Research Section A* (2013)
- [63] D. Doering, J. Baudot, M. Deveaux, B. Linnik, M. Goffe, S. Senyukov, S. Strothauer, J. Stroth, and M. Winter, “Noise performance and ionizing radiation tolerance of cmos monolithic active pixel sensors using the 0.18 μm cmos process”, *Journal of Instrumentation* Vol. 9 (2014)
- [64] M. Winter, “Development of Swift, High Resolution, Pixel Sensor Systems for a High Precision Vertex Detector suited to the ILC Running Conditions”, 2009.
http://www.iphc.cnrs.fr/IMG/desy_mimosa_report_prc09.pdf
- [65] G. Deptuch et al., “Development of monolithic active pixel sensors for charged particle tracking”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 511 (2003) pp. 240–249
- [66] G. Deptuch et al., “Monolithic Active Pixel Sensors adapted to future vertex detector requirements”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 535, Issues 1-2 (2004) pp. 366–369
- [67] Ch. Hu-Guo et al., “CMOS pixel sensor development: a fast read-out architecture with integrated zero suppression”, *Journal of Instrumentation* Vol. 4 (2009)
- [68] R. De Masi et al., “CMOS Pixel Sensors for High Precision Beam Telescopes and Vertex Detectors”, *11th ICATPP Conference on Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications* (2009)
- [69] C. Hu-Guo et al., “First reticule size MAPS with digital output and integrated zero suppression for the EUDET-JRA1 beam telescope”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 623 (2010) pp. 480–482
- [70] A. Dorokhov et al., “Improved radiation tolerance of MAPS using a depleted epitaxial layer”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 624, Issue 2 (2010) pp. 432–436
- [71] M. Deveaux et al., “Radiation tolerance of a column parallel CMOS sensor with high resistivity epitaxial layer”, *Journal of Instrumentation* Vol. 6 (2011)
- [72] Y. Değerli, “Design of fundamental building blocks for fast binary readout CMOS sensors used in high-energy physics experiments”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 602 (2009) pp. 461–466
- [73] E. Pozzati, M. Manghisoni, L. Ratti, V. Re, V. Speziali, and G. Traversi, “MAPS in 130 nm triple well CMOS technology for HEP applications”, *Nuclear Instruments and Methods in Physics Research Section A* Vol. 473 (2007) pp. 83–85
- [74] G. Traversi, A. Bulgheroni, M. Caccia, M. Jastrzab, M. Manghisoni, E. Pozzati, L. Ratti, and V. Re, “First generation of deep n-well CMOS MAPS with in-pixel sparsification for

- the ILC vertex detector”, *Nuclear Instruments and Methods in Physics Research Section A Vol. 604* (2009) pp. 390–392. Proceedings of the 8th International Conference on Position Sensitive Detectors (PSD8)
- [75] K. Waldschmidt, “Embedded Systems”, 2006. Lecture Notes
- [76] P. Clote and E. Kranakis, *Boolean Functions and Computation Models*. Springer Press, 2002
- [77] Lattice Semiconductors, “LatticeECP3 Family Handbook”, 2012.
<http://www.latticesemi.com>
- [78] H. Quinn, P. Graham, K. Morgan, J. Krone, M. Caffrey, and M. Wirthlin, “An Introduction to Radiation-Induced Failure Modes and Related Mitigation Methods for Xilinx SRAM FPGAs”, 2007
- [79] N. Abel, *Design and Implementation of an Object-Oriented Framework for Dynamic Partial Reconfiguration*. PhD thesis, University Heidelberg, 2011
- [80] Austria Microsystems, “Processes & Runs”. Presentation at the CMP annual users meeting 2012, Paris
- [81] W. Dulinski, “Novel packaging methods for ultra-thin monolithic sensors ladders construction”. Presentation at the FEE-2011, Bergamo University
- [82] List of MIMOSA Chips, IPHC website.
<http://www.iphc.cnrs.fr/List-of-MIMOSA-chips.html>
- [83] IPHC Strassbourg, “MIMOSA-26 User Manual”, 2011. Sensor Documentation v1.5
- [84] B. Neumann, “Entwicklung einer FPGA-Basierten JTAG-Ansteuerung für die Sensoren des CBM-MVD”, Master’s thesis, University Frankfurt, 2013
- [85] O. Torheim, *Design and implementation of fast and sparsified readout for Monolithic Active Pixel Sensors*. PhD thesis, University Bergen, 2010
- [86] IPHC Strassbourg, “MIMOSA-28 User Manual”, 2011. Sensor Documentation v1.1
- [87] A. Himmi, G. Doziere, O. Torheim, C. Hu-Guo, and M. Winter, “A Zero Suppression Micro-Circuit for Binary Readout CMOS Monolithic Sensors”, *Proceedings of the TWEPP 2009 workshop, Paris* .
<http://www.iphc.cnrs.fr/Proceedings.html>
- [88] M. Winter et al., “Development of CMOS Pixel Sensors fully adapted to the ILD Vertex Detector Requirements”, [arXiv:1203.3750v1](https://arxiv.org/abs/1203.3750v1)
- [89] M. Wiebusch, “An Automated System to Calibrate the MVD Prototype Detection Thresholds”, Bachelor’s thesis, University Frankfurt, 2012
- [90] F. Morel et al., “MISTRAL & ASTRAL: two CMOS Pixel Sensor architectures suited to the Inner Tracking System of the ALICE experiment”, *Journal of Instrumentation Vol. 9, Issue 01* (2014)

- [91] TowerJazz Homepage. <http://www.jazzsemi.com/>
- [92] CBM-STS Collaboration, “Technical Design Report for the CBM - Silicon Tracking System”, 2013. <http://www.fair-center.eu/for-users/experiments/cbm/cbm-documents.html>
- [93] C. Trageser, “Systematische Untersuchung zur Auswirkung der Detektorgeometrie auf die Spurrekonstruktionseffizienz und Stoßparameterauflösung des CBM Mikro-Vertex-Detektor”, Master’s thesis, University Frankfurt, 2012
- [94] T. Tischler. PhD thesis, in preparation, University of Frankfurt
- [95] T. Galatyuk, *Di-electron spectroscopy in HADES and CBM: from $p + p$ and $n + p$ collisions at GSI to $Au + Au$ collisions at FAIR*. PhD thesis, University Frankfurt, 2009
- [96] E. Krebs, “Employing the CBM Micro Vertex Detector for background rejection in dilepton analyses”. Presentation at the DPG Spring Meeting 2014 in Frankfurt, <https://www-alt.gsi.de/documents/DOC-2014-Aug-96.html>
- [97] M. Bleicher et al., “Relativistic hadron-hadron collisions in the ultra-relativistic quantum molecular dynamics model”, 1999
- [98] UrQMD homepage. <http://urqmd.org/>
- [99] FairRoot homepage. <http://fairroot.gsi.de/>
- [100] C. Drita and M. Deveaux, “A detector response model for CMOS Monolithic Active Pixel Sensors”, *Proceedings of Science, Bormio (Italy)* (2010)
- [101] P. Forck et al., “Beam Diagnostics Developments for Current Operation of SIS18 and HEBT”, *GSI Scientific Report* (2005)
- [102] S. Seddiki, *Contribution to the development of the Micro-Vertex Detector of the CBM experiment and feasibility study of open charm elliptic flow measurements*. PhD thesis, Strasbourg and Frankfurt, 2012
- [103] C. Trageser, “Simulation der Multiplizitätsverteilung auf den Detektorstationen des MVD am CBM Experiment”, Bachelor’s thesis, University Frankfurt, 2008
- [104] F. Lemke, S. Schenk, and U. Bruening, “The Hierarchical CBM Network Structure and the CBMnet V2.0 Protocol”. Presentation at the DPG Spring Meeting 2012 in Mainz, <https://www-alt.gsi.de/documents/DOC-2012-Apr-5.html>
- [105] J. de Cuveland, “Data Management on the FLES”. Presentation at the CBM Workshop on Online Data Processing, GSI 2014
- [106] D. Hutter and J. de Cuveland, “CBM Readout and Online Processing - Overview and Recent Developments”, 2013. Presentation at the 3rd HIC for FAIR Detector Systems Networking Workshop

- [107] M. Lipinski, T. Wlostowski, J. Serrano, and P. Alvarez, “White rabbit: a PTP application for robust sub-nanosecond synchronization”, in *2011 International IEEE Symposium on Precision Clock Synchronization for Measurement Control and Communication (ISPCS)*, pp. 25–30. 2011
- [108] P. Jansweijer, H. Peek, and E. de Wolf, “White Rabbit: Sub-nanosecond timing over Ethernet”, *Nuclear Instruments and Methods in Physics Research Section A, Vol. 725 (2013)* pp. 187–190
- [109] W. F. J. Müller, “Data Preparation”. Presentation at the CBM Workshop on Online Data Processing, GSI 2014
- [110] GBT Project Homepage.
<https://espace.cern.ch/GBT-Project/default.aspx>
- [111] J. Michel, *Development and Implementation of a New Trigger and Data Acquisition System for the HADES Detector*. PhD thesis, University of Frankfurt, 2012
- [112] J. Michel, “Development of a Realtime Network Protocol for HADES and FAIR Experiments”, Master’s thesis, University of Frankfurt, 2008
- [113] J. Michel, M. Böhmer, I. Fröhlich, G. Korcyl, L. Maier, M. Palka, J. Stroth, M. Traxler, and S. Yurevich, “The HADES DAQ System: Trigger and Readout Board Network”, *IEEE Transactions on Nuclear Science Vol. 58, Issue 4 (2011)* pp. 1745–1750
- [114] Q. Li, 2013. PhD thesis, in preparation, University of Frankfurt
- [115] Q. Li, “An efficient data protocol for encoding preprocessed clusters of CMOS Monolithic Active Pixel Sensors”. Presentation at the 20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013), Amsterdam 2013
- [116] M. Penschuck, 2014. Master’s thesis, in preparation, University of Frankfurt
- [117] I. Fröhlich et al., “The TRB for HADES and FAIR experiments at GSI”,
[arXiv:0810.4723 \[nucl-ex\]](https://arxiv.org/abs/0810.4723)
- [118] I. Fröhlich et al., “A General Purpose Trigger and Readout Board for HADES and FAIR-Experiments”, *IEEE Transactions on Nuclear Science Vol. 55, Issue 1 (2008)* pp. 59–66
- [119] M. Palka et al., “The new data acquisition system for the HADES experiment”,
IEEE Nuclear Science Symposium Conference Record (2008) pp. 1398–1404
- [120] C. Schrader, *A Readout System for the Micro-Vertex-Detector Demonstrator for the CBM experiment at FAIR*. PhD thesis, Goethe-Universität Frankfurt, 2011
- [121] S. Amar-Youcef, *Design and performance studies of the Micro-Vertex-Detector for the CBM experiment at FAIR*. PhD thesis, Goethe-Universität Frankfurt, 2011

- [122] M. Traxler, E. Bayer, M. Kajetanowicz, G. Korcyl, L. Maier, J. Michel, M. Palka, and C. Ugur, “A compact system for high precision time measurements (< 14 ps RMS) and integrated data acquisition for a large number of channels”, *Journal of Instrumentation* *Vol. 6* (2011)
- [123] C. Ugur, W. Koenig, J. Michel, M. Palka, and M. Traxler, “Field programmable gate array based data digitisation with commercial elements”, *Journal of Instrumentation* *Vol. 8* (2013)
- [124] C. Ugur, E. Bayer, N. Kurz, and M. Traxler, “Implementation of a high resolution Time-to-Digital Converter in a Field Programmable Gate Array”, *Proceedings of Science, Bormio (Italy)* (2012)
- [125] M. Koziel et al., “The Prototype of the Micro Vertex Detector of the CBM Experiment”, *Nuclear Instruments and Methods in Physics Research Section A* (2013)
- [126] T. Tischler, personal communication, 2013
- [127] M. Wiebusch, personal communication, 2013
- [128] S. Amar-Youcef, personal communication, 2013

LEBENS LAUF

PERSÖNLICHE ANGABEN

Name Borislav Milanović
E-Mail b.milanovic@gsi.de
Geburtsdatum 17. November 1979
Geburtsort Indjija, Serbien
Familienstand Ledig

STUDIUM

09/2001 – 11/2009 Goethe-Universität Frankfurt am Main
 Studium der Informatik und Physik

 Schwerpunkt: Eingebettete Systeme (Prof. U. Brinkschulte)
 Thema der Diplomarbeit: "Development of a Real-Time
 General Purpose Online Monitoring System for HADES
 and FAIR Experiments"
 Abschluss: Diplom Informatiker (Note: sehr gut)

01/2010 – 12/2014 Goethe-Universität Frankfurt am Main
 Promotion im Fachbereich Kernphysik

 Schwerpunkt: Detektorauslese, Simulationen (Prof. J. Stroth)
 Thema der Dissertation: "Development of the Read-Out
 Controller for the CBM Micro-Vertex Detector"
 Abschluss: Doctor of Philosophy (Disputation in Vorbereitung)

PRAKTISCHE ERFAHRUNGEN

04/2007 – 04/2008 Goethe-Universität Frankfurt am Main,
 Institut für Graphische Datenverarbeitung
 Programmierer

05/2008 – 11/2008 Fortis Bank Netherland, Frankfurt am Main
 IT-Praktikant

06/2013 – 03/2015 Goethe-Universität Frankfurt am Main,
 Institut für Kernphysik
 Wissenschaftlicher Mitarbeiter

SPRACHKENTNISSE

Deutsch, Englisch, Kroatisch: sehr gut
Französisch, Italienisch: basis