





The Phase I

Mu3e Experiment







The Phase I Mu3e Experiment

A. Blondel, A. Bravar, F. Cadoux, A. Damyanova, Y. Favre, D. Ferrere, G. Iacobucci, M. Pohl
*Département de physique nucléaire et corpusculaire,
Université de Genève*

H. Augustin, S. Dittmeier, D. Gottschalk, A. Herkert, L. Huth, M. Kiehn, F. Meier Aeschbacher,
A.-K. Perrevoort, A. Schöning, K. Stumpf, D. Wiedner, B. Windelband
*Physikalisches Institut,
Universität Heidelberg*

H. S. Chen, P. Eckert, H.-C. Schultz-Coulon, W. Shen
*Kirchoff Institut für Physik,
Universität Heidelberg*

I. Perić
*Institut für Prozessdatenverarbeitung und Elektronik,
Karlsruhe Institut für Technologie*

N. Berger, U. Hartenstein, Q.H. Huang, A. Kozlinskiy, D. vom Bruch, F. Wauters
*Institut für Kernphysik,
Johannes-Gutenberg Universität Mainz*

F. Berg, M. Hildebrandt, Z. Hodge, P.-R. Kettle, A. Knecht, A. Mtchedlishvili, A. Papa, S. Ritt,
G. Rutar, A. Stoykov
*Paul Scherrer Institut,
Villigen*

S. Corrodi, G. Dissertori, C. Grab, R. Wallny
*Institute for Particle Physics,
Eidgenössische Technische Hochschule Zürich*

R. Gredig, P. Robmann, U. Straumann
*Physik-Institut,
Universität Zürich*

January 2016



TO DO LIST

- Nik and André: Rewrite executive summary once rest of document is finished 10
- Heiko: Add energy calibration plots by Jan 54
- Antoaneta, Sandro: Round Fibre Production, glue studies, picture 99
- Antoaneta, Sandro: Efficiency, cross-talk (Ti, no Ti), time resolution 100
- Sandro: anything? remove this? 117
- Patrick: Say once what CC bins are . . . 124
- Patrick: Describe tile RO cooling, cabling, mechanics 131
- Dirk: Describe cage inside magnet and how everything attaches to cage . . . 141
- Dirk, Felix, Peter-Raymond: Describe procedure for accessing detector 141
- Dirk, Stefan: Describe power supplies, cabling and power delivery to subsystems 145
- Stefan: Describe generation and distribution of global clock and reset 146
- Stefan, Simon, Nik, Patrick: Describe clock interfaces to sub-detectors 146
- Alex: Not ready chapters marked with red. 187
- Alex: Describe track reconstruction . . . 187
- Heiko: Describe Pixel tuning procedures 207
- Patrick: Describe Tile Calibration 212
- Patrick: Can we use cosmics? - Nik . . . 212
- all: Bring theses and acknowledgements up to date 237
- all: Add Ph.D. funding sources like IMPRS, HGSFP etc. 238



CONTENTS

Executive Summary	10	II Beamline	26
I Introduction	12	4 Muon Beam	28
1 Muon decay to three electrons	13	4.1 Beam Requirements	28
1.1 Kinematics	13	4.2 The Compact Muon Beam Line (CMBL) Phase I Design . . .	28
1.2 Modelling of the signal	13	4.3 CMBL Test Commissioning	30
1.3 Detector Acceptance	14	4.4 Area Layout & Infrastructure . . .	32
1.4 Backgrounds	15	5 Magnet	36
1.4.1 Internal Conversions	15	5.1 Requirements	36
1.4.2 Michel Decays	16	5.2 Magnet Construction	36
1.4.3 Radiative Muon Decays	16	5.2.1 Magnetic bottle effects	37
1.4.4 Bhabha Scattering	16	6 Stopping Target	39
1.4.5 Pion decays	16	6.1 Baseline Design	39
1.4.6 Summary of Background Sources	16	6.2 Alternative Geometries	39
2 Requirements for Mu3e	17	6.3 Production	40
2.1 Goals of the Experiment	17	6.4 Support	42
2.2 Challenges for the Experiment . . .	17	III Pixel detector	43
2.2.1 Backgrounds	17	7 The Mu3e Pixel Sensor	45
2.2.2 Geometric acceptance	18	7.1 HV-Maps Sensor	45
2.2.3 Rate capability	18	7.2 Sensor specification	46
2.2.4 Momentum resolution	18	7.2.1 Mechanical dimensions	46
2.2.5 Vertex resolution	18	7.2.2 Features and performance . . .	46
2.2.6 Timing resolution	18	7.2.3 Readout and flex-print inte- gration	47
3 Experimental Concept	20	7.3 Path towards the Full Sensor . . .	47
3.1 Momentum Measurement with Re- curlers	21	7.3.1 Future MUPIX Improvements	47
3.2 Baseline Design	21	7.4 The MUPIX Prototypes	48
3.3 Building up the Experiment	23	7.4.1 MUPIX 1 and 2	48
3.4 The Phase IA Experiment	23	7.4.2 MUPIX 3 to MUPIX 7	48
3.5 The Phase IB Experiment	25	7.5 The MUPIX Telescope	52
3.6 The Phase II Experiment	25	7.6 Characterization of the Prototypes	52

7.6.1	Laboratory results	53	12 Fibre Detector Prototypes	99
7.6.2	Testbeam results	56	12.1 Round Fibres	99
8	Pixel Mechanics and Supply	67	12.1.1 Round Fibre Setup	99
8.1	Mechanical design of the four pixel layers	68	12.1.2 Round Fibre Results	100
8.1.1	Layer 1	68	12.2 Square Fibres	100
8.1.2	Layer 2	70	12.2.1 Square Fibre Setup	100
8.1.3	Layer 3	70	12.2.2 Prototype V4.1 Performances	101
8.1.4	Layer 4	71	12.2.3 Large Prototype Performances	101
8.2	Flex-print assembly	71	12.2.4 Square Fibre Conclusions	104
8.2.1	Design considerations	71	12.3 Standalone Monte Carlo Simulation	105
8.2.2	Flex-print studies	73	12.3.1 Square Fibre	106
9	Pixel Module Production and Testing	76	12.3.2 Round Fibre	108
9.1	Manufacturing steps	76	13 Fibre Readout	109
9.2	Part tracking	80	13.1 Requirements	109
9.3	Quality control	80	13.2 Baseline Design	109
9.3.1	Common quality tests	81	13.2.1 STiC3.1 Measurements	110
9.3.2	Wafer Test	81	13.2.2 Additional Connections	110
9.3.3	PCB testing	81	13.2.3 Readout Board	111
9.3.4	Flex-print testing	81	13.3 Alternatives	112
9.3.5	Final module testing	81	13.3.1 DRS	113
10	Pixel Readout	83	13.3.2 Others	113
10.1	Clocking	83	14 Fibre Mechanics and Cooling	115
10.1.1	PLL and External Clocks	83	14.1 Cooling	115
10.1.2	Clocking in the digital part	83	14.2 Mechanics	115
10.2	Readout Link	84	14.2.1 Example Solution	115
10.2.1	Link Encoding	84	14.3 SiPM array	117
10.2.2	Synchronization and Alignment	84	V Tile Detector	119
10.2.3	Link Performance	85	15 Tile Detector	121
10.3	Pixel and Column Logic	85	15.1 Detector Design	121
10.4	Readout State Machine	86	15.2 Simulation	122
10.4.1	Readout Sequence	86	15.3 16-Channel Prototype	122
10.4.2	Counter debug mode	89	15.3.1 Detection Efficiency	125
10.4.3	Run Start	89	15.3.2 Time Resolution	125
IV	Fibre Detector	90	16 Tile Scintillator	127
11	Scintillating Fibre Detector	92	16.1 Material	127
11.1	Motivation	92	16.2 Radiation Hardness	128
11.2	Requirements	93	16.3 Manufacturing	128
11.3	Detector Design	94	17 Tile Silicon Photomultipliers	129
11.3.1	Read-out Schemes	96	17.1 SiPM Choice	129
11.3.2	Silicon Photomultipliers	96	17.2 SiPM Radiation Hardness	130
11.3.3	RO Hardware	97	17.3 SiPM Quality Assurance	130
11.4	Calibration & Monitoring	98		

18 Tile Readout	131	VII Data Acquisition and Event Filter	151
18.1 STiC Chip	131	24 Data Acquisition	153
18.2 DRS Sampling Chip	131	24.1 Overview	153
18.3 Readout Scheme	132	24.2 Occupancy	153
18.3.1 Mechanical Support	132	24.3 Front-end requirements	155
18.3.2 Cabling	132	24.3.1 Pixel detector	155
18.3.3 Cooling	132	24.3.2 Timing detector	158
VI Services	133	24.4 Front-end FPGA boards	160
19 Cooling	135	24.4.1 Front-end hardware	160
19.1 Experimental Helium Cooling Tests	135	24.4.2 Front-end firmware	160
19.2 Water Cooling	137	24.5 Read-out links	164
19.3 CFD Simulations of Phase I	137	24.6 Switching cards	165
19.4 Possible Design Optimizations	140	24.6.1 Switching board firmware	165
19.5 Stability against Flow-induced Vibrations	140	24.7 Event filter interface	165
20 Mechanical Integration	141	24.7.1 PCIe Interface	166
20.1 Support of Detector Stations	141	24.8 Data transmission to the GPU	167
20.1.1 Support of Pixel Vertex Layers	141	24.8.1 Event filter firmware	167
20.1.2 Support of the Central Station	141	25 Run Control, Data Collection and Storage	169
20.1.3 Support of the Recurl Stations	143	25.1 The MIDAS system	169
20.2 Detector cage and rail system	143	25.2 Run start/stop synchronisation	169
20.3 Access to the Mu3e Detector	143	VIII Software	171
21 Power Distribution and Cabling	145	26 Software Development and Tools	173
21.1 Power distribution to the front end ASICs	145	26.1 Development Process	173
26.1.1 Organsiation	173	26.1.1 Organsiation	173
26.1.2 Versioning	173	26.1.2 Versioning	173
26.2 Tools and Libraries	174	26.2 Tools and Libraries	174
26.2.1 Operating System	174	26.2.1 Operating System	174
26.2.2 Programming Language	174	26.2.2 Programming Language	174
26.2.3 Build System	174	26.2.3 Build System	174
26.2.4 External Libraries	174	26.2.4 External Libraries	174
26.2.5 Mu3e Libraries	175	26.2.5 Mu3e Libraries	175
22 Clock Distribution	146	27 Simulation	176
22.1 Clock Distribution scheme	146	27.1 Detector geometry	176
22.2 Reference clock generation	146	27.1.1 Beam delivery	176
22.3 Clock distribution through the switching boards	146	27.1.2 Target	176
22.4 Clock distribution on the front-end boards	147	27.1.3 Pixel detector	176
22.5 Clock distribution to the ASICs	147	27.1.4 Scintillating fibres	176
22.5.1 to the MuPix	147	27.1.5 Tile detector	179
22.5.2 to the (Fiber Detector) MuSTiC	147	27.2 Magnetic field	179
23 Slow Control	148	27.3 Physics Processes	179
23.1 Pixel Configuration and Control	150	27.3.1 Multiple Coulomb scattering	179
		27.3.2 Muon decays	182
		27.4 Time structure	184
		27.5 Detector response	184
		27.5.1 Pixel detector response	184

27.5.2 Fibre detector response . . .	185	32 Tile Calibration	212
27.5.3 Tile detector response . . .	185	33 Detector Alignment	213
28 Reconstruction	187	33.1 Effects of Misalignment	213
28.1 Introduction	187	33.2 Position Monitoring System	213
28.2 MS Triplet fit	188	33.3 Track-Based Alignment	217
28.2.1 Linearization	188	33.4 Target Alignment	217
28.3 MS Track fit	189	X Simulated Performance	219
28.3.1 Pixel size	189	34 Sensitivity Study	220
28.4 Track reconstruction	190	34.1 Signal Performance	220
28.4.1 Detector and Simulation . .	191	34.1.1 Vertex Fit	220
28.4.2 Triplet seeds	191	34.1.2 Mass and Momentum Re- construction	222
28.4.3 Short tracks	192	34.1.3 Signal Efficiency	224
28.4.4 Double hits	193	34.2 Backgrounds	224
28.4.5 Track intersections	194	34.2.1 Internal Conversion Back- ground	224
28.4.6 Short tracks performance .	195	34.2.2 Accidental Background . .	225
28.4.7 Long tracks	196	A Appendix	237
28.4.8 Energy loss	197	A.1 Mu3e theses	237
28.5 Vertex reconstruction	198	A.2 Acknowledgements	238
28.6 Fibre linking	199	Bibliography	239
28.7 Tile linking	200	Drawings	246
28.8 Summary	201	Mechanical drawing of the layer 1 pixel module	246
29 Online Event Selection	202	Mechanical drawing of the layer 2 pixel module	247
29.1 Selection Cuts	202	Mechanical drawing of the layer 3 pixel module	248
29.2 Multiple Scattering Fit	202	Mechanical drawing of the layer 4 pixel module	249
29.3 Vertex Selection	204	Overview of all readout links in phase Ia	250
29.4 Current Status	204	Overview of all readout links in phase Ib	251
IX Calibration and Alignment	205	Mapping of pixel chip addresses	252
30 Pixel Calibration	207	Common firmware modules in the read- out chain	253
30.1 MuPix Pixel Tuning	207		
30.2 MuPix Tuning Procedure	207		
30.2.1 Tuning Parameters	207		
30.2.2 Tuning Procedure	208		
31 Fibre Calibration	210		
31.1 MuSTiC Thresholds	210		
31.2 Inter Channel Calibration	210		
31.3 Monitoring Tools	211		



EXECUTIVE SUMMARY

Nik and André: Rewrite executive summary once rest of document is finished

We propose an experiment (*Mu3e*) to search for the lepton flavour violating (LFV) decay $\mu^+ \rightarrow e^+e^-e^+$. We aim for an ultimate sensitivity of one in 10^{16} μ -decays, four orders of magnitude better than previous searches. This sensitivity is made possible by exploiting modern silicon pixel detectors providing high spatial resolution and hodoscopes using scintillating fibres and tiles providing precise timing information at high particle rates.

Existing beamlines available at PSI providing rates of order 10^8 muons per second allow to test for the decay $\mu^+ \rightarrow e^+e^-e^+$ in one of 10^{15} muon decays. In a first phase of the experiment, we plan to make use of this and establish the experimental technique whilst at the same time pushing the sensitivity by three orders of magnitude.

The installation of a new muon beamline at the spallation neutron source is currently under discussion at PSI. Such a *High Intensity Muon Beam* (HiMB) will provide intensities in excess of 10^9 muons per second, which in turn are required to reach the aimed sensitivity of $B(\mu^+ \rightarrow e^+e^-e^+) \sim 10^{-16}$.

The proposed experiment is highly complementary to other LFV searches for physics beyond the standard model, i.e. direct searches performed at the Large Hadron Collider (LHC) and indirect searches in the decay of taus and muons, such as the decay $\mu^+ \rightarrow e^+\gamma$, which is the subject of the *MEG* experiment currently in operation at PSI. The proposed experiment for the search $\mu^+ \rightarrow e^+e^-e^+$ will test lepton flavour violating models of physics beyond the Standard Model with unprecedented sensitivity.

This sensitivity is experimentally achieved by a novel experimental design exploiting silicon pixel

detectors based on High Voltage Monolithic Active Pixel Sensors (HV-MAPS). This technology provides high granularity, important for precision tracking and vertexing, and allows one to significantly reduce the material budget by thinning down the sensors and by integrating the hit digitisation and readout circuitry in the sensor itself. The detector geometry is optimized to reach the highest possible momentum resolution in a multiple Coulomb scattering environment, which is needed to suppress the dominating background from the radiative muon decay with internal conversion, $\mu \rightarrow eee\nu\bar{\nu}$. The time information of the decay electrons¹, obtained from the pixel detector is further improved by a time-of-flight system consisting of a scintillating fiber hodoscope and tiles with Silicon Photo-Multipliers (SiPM) for light detection. By combining both detector systems accidental background can be reduced below the aimed sensitivity of $B(\mu^+ \rightarrow e^+e^-e^+) \sim 10^{-16}$.

We will complete the sensor development and start constructing the detector in 2013, in order to be ready for first exploratory data taking at an existing beam line with a first minimal detector setup in 2015. A detector capable of taking data rates of order 10^8 muons per second and capable of reaching a sensitivity of $B(\mu^+ \rightarrow e^+e^-e^+) \sim 10^{-15}$ will be available in 2016. This *Phase I* detector is the main focus of this proposal.

In *Phase II*, beyond 2017, the experiment will reach the ultimate sensitivity by exploiting a possible new high intensity muon beamline with an intensity of $> 2 \cdot 10^9$ muons per second. In the absence of a signal, LFV muon decays can then


¹Here and in the following, the term “electron” denotes generically both decay electrons and positrons.



THE PHASE I MU3E EXPERIMENT

be excluded for $B(\mu^+ \rightarrow e^+e^-e^+) < 10^{-16}$ at 90 % confidence level.

PART I



INTRODUCTION



THE DECAY $\mu \rightarrow eee$

The Mu3e experiment is designed to detect the decay $\mu \rightarrow eee$ with a high efficiency and at the same time suppress backgrounds mimicking the signal by many orders of magnitude. For the physics motivation and theory predictions, please consult the Mu3e letter of intent [1] and research proposal [2]. The present chapter describes the kinematics of the signal and the studied background sources.

1.1 Kinematics

The decay $\mu \rightarrow eee$ is assumed to proceed promptly. For discriminating signal and background, energy and momentum conservation can be exploited. The vectorial sum of all decay particle momenta \vec{p}_i should vanish:

$$|\vec{p}_{tot}| = \left| \sum \vec{p}_i \right| = 0 \quad (1.1)$$

and the total energy has to be equal to the muon mass:

$$E_{tot} = \sum E_i = m_\mu. \quad (1.2)$$

The energies of the decay electrons (positrons) range from the electron mass up to half the muon mass, which is about 53 MeV. All decay particles must lie in a plane. Therefore, the decay is described by two independent variables in addition to three global rotation angles describing the orientation in space.

1.2 Modelling of the signal

The decay dynamics for the $\mu \rightarrow eee$ signal are dependent on the unknown lepton flavour violating mechanism. We typically assume a phase-space distribution for the signal electrons in our simulations, if not otherwise stated. In order to study effects of different decay dynamics, we utilise the general parametrized Lagrangian proposed by Kuno and Okada [3]:

$$\begin{aligned} L_{\mu \rightarrow eee} = & \frac{4G_F}{2} [m_\mu A_R \bar{\mu}_R \sigma^{\mu\nu} e_L F_{\mu\nu} \\ & + m_\mu A_L \bar{\mu}_L \sigma^{\mu\nu} e_R F_{\mu\nu} \\ & + g_1 (\bar{\mu}_R e_L) (\bar{e}_R e_L) \\ & + g_2 (\bar{\mu}_L e_R) (\bar{e}_L e_R) \\ & + g_3 (\bar{\mu}_R \gamma^\mu e_R) (\bar{e}_R \gamma_\mu e_R) \\ & + g_4 (\bar{\mu}_L \gamma^\mu e_L) (\bar{e}_L \gamma_\mu e_L) \\ & + g_5 (\bar{\mu}_R \gamma^\mu e_R) (\bar{e}_L \gamma_\mu e_L) \\ & + g_6 (\bar{\mu}_L \gamma^\mu e_L) (\bar{e}_R \gamma_\mu e_R) + H.c.] \end{aligned} \quad (1.3)$$

The form factors $A_{R,L}$ describe tensor type (dipole) couplings, mostly acquiring contributions from the photon penguin diagram, whereas the scalar-type ($g_{1,2}$) and vector-type ($g_3 - g_6$) form factors can be regarded as four fermion contact interactions, to which the tree diagram contributes in leading order. We generate different signal models by varying the relative strengths of the $A_{R,L}$ and $g_3 - g_6$ parameters.




Figure 1.1: Energy distribution of the highest energy positron in the decay $\mu^+ \rightarrow e^+e^-e^+$ for different effective LFV models. The black line corresponds to pure dipole and the red and blue line to pure four-fermion contact interaction models (no penguin contribution) whereas the other lines correspond to a mixture of dipole and vector interactions. Based on [3].




Figure 1.2: Acceptance of the lowest energy decay positron in the decay $\mu^+ \rightarrow e^+e^-e^+$ for different effective LFV models as function of energy. The black line corresponds to pure dipole and the red and blue line to pure four-fermion contact interaction models (no penguin contribution) whereas the other lines correspond to a mixture of dipole and vector interactions. Based on [3].

1.3 Detector Acceptance

The acceptance of the proposed $\mu \rightarrow eee$ experiment is determined by its geometrical acceptance and energy coverage. For various coupling as-

sumptions about the LFV amplitude, the energy spectrum of the decay particle with the highest energy E_{max} is shown in Figure 1.1, and the acceptance for the decay particle with the lowest energy E_{min} in Figure 1.2, respectively. In order to




Figure 1.3: Integrated branching fraction of the decay $\mu \rightarrow eee\nu\nu$ for which the three decay electrons lie above a certain visible mass: for all internal conversion decays (blue line) and those with all three decay particles in the detector acceptance ($E > 10$ MeV, $|\cos\theta| < 0.8$; red line). The matrix element was taken from [4].




Figure 1.4: Energy spectrum of all electrons from internal conversion decays.

achieve a high acceptance, the detector must be able to reconstruct tracks with momenta ranging from half the muon mass down to a few MeV with large solid angle coverage.

1.4 Backgrounds

The final sensitivity of the proposed experiment depends on the ability to reduce backgrounds from various sources. Two categories of backgrounds are considered; irreducible backgrounds, such as $\mu^+ \rightarrow e^+e^-e^+\nu\bar{\nu}$, whose suppression strongly depends on the resolution of the detector, and accidental backgrounds that scale linearly or with the square of the beam intensity.

In the following sections, the main background sources considered are discussed.




Figure 1.5: Energy spectrum of the electron or positron with minimum energy from internal conversion decays.




Figure 1.6: Invariant masses of the two possible e^+e^- combinations for internal conversion decays.




Figure 1.7: Invariant masses of the two possible e^+e^- combinations for internal conversion decays with a visible mass above 90 MeV and the electrons and positrons in the detector acceptance ($E > 10$ MeV, $|\cos\theta| < 0.8$).

1.4.1 INTERNAL CONVERSIONS

The decay $\mu \rightarrow eee\nu\nu$ occurs with a branching fraction of $3.4 \cdot 10^{-5}$ [5]. It can be distinguished from the $\mu \rightarrow eee$ process by making use of energy and momentum conservation to reconstruct



the undetected neutrinos; in order to separate the $\mu \rightarrow eee$ events from $\mu \rightarrow eee\nu\nu$ events, the total momentum in the event is required to be zero and the energy equal to the muon rest energy. The branching fraction as a function of the energy cut of the $\mu \rightarrow eee\nu\nu$ process [4] is shown in Figure 1.3. Figures. 1.4 and 1.5 show the energy spectrum of all and the lowest energy electron from internal conversion decays, Figs. 1.6 and 1.7 the invariant masses of e^+e^- combinations calculated with the matrix element from [4]. This process is the most serious background for the $\mu \rightarrow eee$ search and can only be resolved by a very good energy resolution.

1.4.2 MICHEL DECAYS

Using a beam of positive muons, one of the main processes contributing to accidental background is that of the ordinary Michel decay $\mu^+ \rightarrow e^+\nu\bar{\nu}$. This process does not produce a negatively charged particle (electron), which is one of the main characteristics of the $\mu^+ \rightarrow e^+e^-e^+$ decay. Therefore, it can only contribute as potential background on its own if a track is wrongly reconstructed, or in combination with other processes that “naturally” provide negatively charged tracks (electrons), such as radiative decays with internal or external photon conversions or Bhabha scattering.

1.4.3 RADIATIVE MUON DECAYS

The process $\mu^+ \rightarrow e^+\gamma\nu\bar{\nu}$ (branching fraction $1.4 \cdot 10^{-2}$ for photon energies above 10 MeV [6]) can deliver an oppositely charged electron if the photon converts either in the target region or in the detector. Contributions from conversions outside of the target are greatly suppressed if a vertex constraint is applied and by minimizing the material in both the target and detector. Photon conversion in the target generates an event topology similar to the radiative decay with internal conversion $\mu \rightarrow eee\nu\nu$, which is discussed above.

Due to the missing energy from the neutrinos, this process mainly contributes to the acciden-

tal background in combination with an ordinary muon decay.

1.4.4 BHABHA SCATTERING

Positrons from the ordinary muon decay or beam-positrons can undergo Bhabha scattering with electrons in the target material, leading to an electron-positron pair from a common vertex. In combination with a positron from an ordinary muon decay, the electron-positron pair can mimic a signal decay. In addition, Bhabha scattering is the main source of electrons for combinatorial background involving two Michel decays. The amount of Bhabha scattering can (and has to be) reduced by minimizing the amount and the average atomic number Z of the material in the target.

1.4.5 PION DECAYS

Certain decays of pions, especially $\pi \rightarrow eee\nu$ (branching fraction $3.2 \cdot 10^{-9}$ [7]) and $\pi \rightarrow \mu\gamma\nu$ (branching fraction $2.0 \cdot 10^{-4}$ [8]) with subsequent photon conversion are indistinguishable from signal events if the momenta of the final state particles fit the muon mass hypothesis; a low pion contamination of the primary beam (estimated to be in the order of 10^{-12} for the high intensity beam-line), the small branching fraction and the small slice of the momentum is assumed to lead to negligible rates in the kinematic region of interest.

1.4.6 SUMMARY OF BACKGROUND SOURCES

First simulation studies have been performed to calculate the different background contributions. For the phase Ia experiment (see chapter 3 for the phases of the experiment), rates are low enough that accidental backgrounds do not play a large role and $\mu \rightarrow eee\nu\nu$ is the main worry, also due to the limited tracking performance. In the phase Ib experiment, the situation is reversed; $\mu \rightarrow eee\nu\nu$ can be well suppressed by the improved tracking and the main background is an electron positron pair from Bhabha scattering combined with a positron from a Michel decay.



REQUIREMENTS FOR MU3E

2.1 Goals of the Experiment

The goal of the Mu3e experiment is to observe the process $\mu \rightarrow eee$ if its branching fraction is larger than 10^{-16} or otherwise to exclude a branching fraction of $> 10^{-16}$ at the 90% certainty level. In order to achieve these goals, $> 8 \cdot 10^{16}$ muon decays have to be observed¹ and any background mimicking the signal process has to be suppressed to below the 10^{-16} level. The additional requirement of achieving these goals within a reasonable measurement time of one year of data taking dictates a muon stopping rate of $2 \cdot 10^9$ Hz and a high geometrical acceptance and efficiency of the experiment.

The currently best source of low-energy muons in the world, the $\pi E5$ beam line at PSI, provides muon rates up to $1 \cdot 10^8$ Hz. Higher intensities are possible and currently under study in the *high intensity muon beam* (HiMB) project, will however not be available before 2020. In order to establish the novel technologies for Mu3e, set first competitive limits and prepare for the very high intensity running, we plan to run a phase I experiment at $\pi E5$, making best possible use of the available rates. The present document describes the technical design of this phase I detector. Both the phase II experiment for HiMB and a phase 0 technology demonstration to be run at the earliest time possible will be described in separate documents.

The phase I experiment in turn is split into phase Ia with the central tracking system only and a rate capability of about $1 \cdot 10^7$ Hz and phase Ib

including additional tracking and fibre and tile timing systems for running at $1 \cdot 10^8$ Hz.

The requirements for the phase I experiment outlined below are based on a sensitivity of $5 \cdot 10^{-15}$ on the branching fraction, which would require $> 3 \cdot 10^{15}$ stopped muons² or $3 \cdot 10^7$ s (360 days) of run time at $1 \cdot 10^8$ Hz stopping rate.

2.2 Challenges for the Experiment

2.2.1 BACKGROUNDS

There are two kinds of backgrounds: Overlays of different processes producing three tracks resembling a $\mu \rightarrow eee$ decay (*accidental background*) and radiative decays with internal conversion (*internal conversion background*) with a small energy fraction carried away by the neutrinos. Accidental background has to be suppressed via vertexing, timing and momentum measurement, whereas momentum measurement is the only handle on internal conversion.

Accidental background requires an electron track combined with two positron tracks. The main source of electron tracks is Bhabha scattering of muon decay positrons with electrons in the stopping target material. This background can be reduced by minimizing the material in the stopping target, which also helps against electrons from photon conversion or Compton scattering.

In order to suppress fake electron tracks from the recurling part of a positron track, a reliable

¹Assuming a total efficiency of 20%.

² $N_{\text{required}} = 2.3/(s \cdot \epsilon)$ for a sensitivity s and a total efficiency $\epsilon \approx 15\%$ (phase I)




Figure 2.1: Contamination of the signal region (one sided cut) with internal conversion events as a function of momentum sum resolution.

determination of the direction of motion of particles is necessary, which can be achieved by a combination of accurate curvature measurements combined with repeated timing measurements.

2.2.2 GEOMETRIC ACCEPTANCE

For a three-body decay with a priori unknown kinematics such as $\mu \rightarrow eee$, the acceptance has to be as high as possible in order to test new physics in all regions of phase space. There are two kinds of acceptance losses, losses of tracks downstream or upstream, where beam entry and exit prevent instrumentation, and losses of low transverse momentum tracks, which do not transverse a sufficient number of detector planes, and are not reconstructed.

2.2.3 RATE CAPABILITY

The phase I Mu3e detector should be capable of running with $1 \cdot 10^8$ Hz of muon decays. This poses challenges for the detectors, the data acquisition and the readout.

2.2.4 MOMENTUM RESOLUTION

The momentum resolution directly determines to what level internal conversion background can be suppressed and thus to which level the experiment can be run background free. In order to reach a sensitivity of $5 \cdot 10^{-15}$ with a 2σ cut on the reconstructed muon mass, the average momentum resolution has to be better than 1.0 MeV, see Figure 2.1.

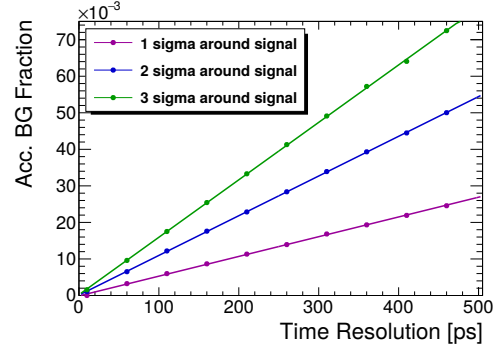


Figure 2.2: Survival fraction of accidental coincidences of a Michel decay and a e^+e^- pair as a function of the combined resolution of the timing detectors.




Figure 2.3: Survival fraction of accidental coincidences of two Michel decays and an e^- as a function of the combined resolution of the timing detectors.

2.2.5 VERTEX RESOLUTION

Keeping apart vertices from different muon decays is a key tool in suppressing accidental background. The vertex resolution is essentially determined by the amount of multiple scattering (and thus material) in the innermost detector layer and the stopping target as well as the average distance between the vertex and the first detector layer.

2.2.6 TIMING RESOLUTION

Good timing is essential for reducing combinatorial background at high rates and also facilitates event reconstruction.

The combinatorial background has a component scaling linearly with the rate (e^+e^- pair plus



a Michel positron) and a component quadratic in the rate (electron plus two Michel positrons). The suppression of these components by timing measurements is also linear and quadratic in the timing resolution. Simulation studies have shown that the linear part is dominating at rates at least up to $2 \cdot 10^9$ muon stops per second.

The requirement of a reduction of accidental background by at least two orders of magnitude

puts very tight demands on the resolution of the timing detectors. This can be seen in figures 2.2 and 2.3, which show the suppression for the two different accidental background components (Michel positron plus a e^+e^- pair, two Michel positrons plus an electron) as a function of the overall time resolution.



EXPERIMENTAL CONCEPT

The phase I Mu3e detector is aimed at the background free measurement or exclusion of the decay $\mu \rightarrow eee$ at the level of $5 \cdot 10^{-15}$. As discussed in more detail in the preceding chapter 2, these goals require to run at high muon decay rates, an excellent momentum resolution in order to suppress background from the internal conversion decay $\mu \rightarrow eee\nu\nu$ and good vertex and timing resolution in order to efficiently suppress combinatorial background.

We intend to measure the momenta of the muon decay electrons in a solenoidal magnetic field using a silicon pixel tracker. At the electron energies of interest, multiple Coulomb scattering in detector material is the dominating factor affecting momentum resolution. Minimizing this material in the active detector parts is thus of utmost importance.

The proposed detector consists of an ultra thin silicon pixel tracker, made possible by the High-Voltage Monolithic Active Pixel (HV-MAPS) technology (see chapter 7). Just four radial layers around a fixed target in a solenoidal magnetic field allow for precise momentum and vertex determination. Two timing detector systems guarantee good combinatorial background suppression and high rate capabilities.

The Mu3e experiment is ultimately designed to have a sensitivity four orders of magnitude better than the current limit on $\mu \rightarrow eee$ (10^{-12}), so it is reasonable to plan for a staged detector design, with each stage roughly corresponding to an order of magnitude improvement (and an order of magnitude in muon rate).




Figure 3.1: Tracking in the spatial resolution dominated regime




Figure 3.2: Tracking in the scattering dominated regime



3.1 Momentum Measurement with Recurlers

Due to the low momenta of the electrons from muon decay, multiple scattering is the dominating effect on momentum measurement. With our fine-grained pixel detector, we are thus in a regime where scattering effects dominate over sensor resolution effects, see Figs. 3.1 and 3.2. Thus adding additional measurement points does not necessarily improve the precision.

The precision of a momentum measurement depends on the amount of track deflection Ω in the magnetic field B and the multiple scattering angle Θ_{MS} , see Figure 3.3; to first order:

$$\frac{\sigma_p}{p} \propto \frac{\Theta_{MS}}{\Omega}. \quad (3.1)$$

So in order to have a high momentum precision, a large lever arm is needed. This can be achieved by moving tracking stations to large radii, which however compromises the acceptance for low momentum particles. In the case of muon decays, all track momenta are below 53 MeV and all tracks will thus curl back towards the magnet axis if the magnet bore is sufficiently large. After exactly half a turn, effects of multiple scattering on momentum measurement cancel in first order, see Figure 3.4. To exploit this feature, we optimized the experimental design specifically for the measurement of re-curling tracks, leading to a narrow, long tube layout.

Measuring the momentum from bending outside of the tracker also allows us to place timing detectors inside, without strongly affecting the resolution.

3.2 Baseline Design

The proposed Mu3e detector is based on two double layers of HV-MAPS around a hollow double cone target, see Figures 3.5 and 3.6. The outer two pixel sensor layers are extended upstream and downstream to provide precise momentum measurements in an extended region with the help of re-curling electrons. The silicon detector layers (described in detail in chapter 7) are supplemented by two timing systems, a scintillating fibre tracker in the central part (see chapter 11) and scintillating tiles (chapter 15) inside the re-curl layers. Precise timing of all tracks is neces-




Figure 3.3: Multiple scattering as seen in the plane transverse to the magnetic field direction.




Figure 3.4: Multiple scattering for a semi-circular trajectory.




Figure 3.5: Schematic view of the experiment cut along the beam axis in the phase Ib configuration.




Figure 3.6: Schematic view of the experiment cut transverse to the beam axis. Note that the fibres are not drawn to scale.



sary for event building and to suppress accidental combinatorial background.

The entire detector is built in a cylindrical shape around a beam pipe, with a total length of approximately 1.2 m, inside a 1 T solenoid magnet with 1 m inside diameter and 3 m total length (chapter 5). In the longitudinal direction the detector is sub-divided into three stations, the central detector with target, inner silicon double layer, fibre tracker and outer silicon double layer, and forward and backward recurl stations with two silicon recurl layers surrounding a tile timing detector. In order to separate tracks coming from different muon decays, the target has a large surface with 10 cm length and 3.8 cm diameter. The target shape is a hollow double cone, see chapter 6. Around the target the two inner silicon pixel layers, also referred to as the vertex layers, cover a length of 12 cm. The innermost layer will have 8, the second one 10 sides of 2 cm each, corresponding to an average radius of 2.1 cm and 2.5 cm, respectively. The inner silicon layers are supported by two half cylinders made from 25 μm thin Kapton foil mounted on plastic end pieces. All silicon sensors are thinned to 50 μm and mounted on thin Kapton flexprints, resulting in a radiation length of $X/X_0 \approx 0.1\%$ per layer. The detector will be cooled with gaseous helium.

The hit information from the silicon sensors is timestamped at 62.5 MHz providing a time resolution of better than 12.5 ns.

The fibre tracker sits inside the silicon pixel layer at around 6 cm, providing timing information for decay positrons and electrons. It is composed from three layers of 250 μm thick round or squared 30 cm long scintillating fibres, see Figure 3.6. The fibre tracker is read out by silicon photo multiplier arrays and provides timing information better than 500 ps accuracy.

The silicon pixel layers three and four are just outside the fibre tracker at a mean radius of 7.6 cm and 8.9 cm. The active area has a cylindrical shape of 36 cm length. The layer three has 24 sides, layer four 28 sides of 1.9 cm width each. Both outer layers are constructed as modules of 4 sides, six modules for layer three and seven modules for layer four.

Copies of silicon pixel layer three and four are also used in the recurl stations. One recurl station each is covering the upstream and downstream regions. These recurl stations add further precision and acceptance to the momentum measurement

of the electrons, see section 3.1. While the silicon layer design is (almost) identical to the central part, the timing detector in the recurl region can be much thicker compared to the fibre tracker, as the particles can and should be stopped here. This is done by using scintillating tiles of about $6.5 \times 6.0 \times 5 \text{ mm}^3$ size. These tiles provide a much better time resolution than the thin fibre tracker in the center. Following the dimensions of the recurl silicon layers, the tile stations have an active length of 36 cm and a cylindrical shape with a radius of ≈ 6 cm. All central detector components are mounted on spokes providing a light stiff support. The recurl silicon layers and tiles are mounted on the beam pipe support. In addition to the silicon and scintillating tile sensors the beam pipe support also carries the services and the PCBs equipped with the front-end electronics (chapter 24). Signal and power connection to the silicon layers is provided by flex prints which are also part of the mechanical support of the silicon sensors.

3.3 Building up the Experiment

One of the advantages of the design concept presented is its modularity. Even with a partial detector, physics runs can be taken. In an early commissioning phase at smaller muon stopping rates, the detector will run with the central silicon detector only (see Figure 3.7). The silicon detectors of the recurl stations are essentially copies of the central outer silicon detector; after a successful commissioning of the latter, they can be produced and added to the experiment as they become available together with the connected tile detectors. The fibre tracker can also be added later, since it is only needed to resolve combinatorial background at higher event rates and track multiplicities. The loss of momentum resolution due to multiple scattering at the additional material of the fibre tracker will be fully compensated by the improved momentum measurement with re-curlers. The configuration with two recurl stations (Figure 3.8) defines a medium-size setup, well suited for phase IB running.

3.4 The Phase IA Experiment

The phase IA of the Mu3e experiment will start with a minimum configuration with the target re-



Figure 3.7: Minimum detector configuration for early commissioning with central silicon only (phase IA).




Figure 3.8: Detector with one set of recurl stations for physics runs and tile detector commissioning (phase IB).




Figure 3.9: Possible final detector with two sets of recurl stations, smaller target and more segmented inner layers for high intensity physics runs (phase II).



gion surrounded by double layers of inner and outer silicon pixel detectors, see Figure 3.7. This configuration defines the minimal configuration, as it allows to determine the momentum, the vertex position and the time of the decay precisely enough to produce very competitive physics results with a sensitivity down to $\mathcal{O}(10^{-14})$. It is foreseen to run in the first year in this configuration at a muon stopping rate on target at around $1 - 2 \cdot 10^7$ Hz. The average spacing in time between decays will be much larger than the time resolution of the pixel detector and combinatorial background can be additionally suppressed with the help of the vertex reconstruction. The precision of the momentum resolution will be somewhat limited, as most tracks do not recur within the instrumented volume.

3.5 The Phase IB Experiment

In the phase IB the detector will be complemented by the first pair of recur stations, the corresponding tile detectors and the fibre tracker, see Figure 3.8. Adding the recur stations will significantly enhance the momentum resolution and thus improve the suppression of internal conversion background. The insertion of the fibre tracker and the tile detector stations gives a much better time resolution in comparison to the silicon pixel only and will greatly help to determine the charge of tracks recurling into the target region. The fibre tracker will deliver a time resolution better than 500 ps, while the tile detector will have < 100 ps resolution for the tracks passing the recur stations. The high time resolution will allow running at the highest possible rate at the $\pi E5$ muon beam line at PSI of $\approx 1 \cdot 10^8$ Hz. The sensitivity reach in

this phase of the experiment of $\mathcal{O}(10^{-15})$ will be limited by the available muon decay rate.

3.6 The Phase II Experiment

A new high intensity muon beam line currently under study at PSI and delivering $> 2 \cdot 10^9$ Hz muon stops is crucial for the phase II of the proposed experiment. To fully exploit the new beam facility the limited detector acceptance at phase IB will be further enhanced by adding a second pair of recur and tile detector stations, see Figure 3.9. These extra stations will allow to measure precisely the momentum of all particles in the acceptance of the inner tracking detector. At the same time the extra tile detector stations with their high time resolution and small occupancy will help to fight the increased combinatorics at very high decay rates. The larger initial muon rate allows for a more restrictive collimation of the beam and thus a smaller (and potentially longer) target region leading to a much improved vertex resolution. The HV-MAPS technology potentially has a time resolution of $\mathcal{O}(1 \text{ ns})$ if an adequate time-walk correction is implemented - this would allow to further reduce combinatorial background without adding material. Advanced wafer post-processing technologies and chip-to-chip bonding could obviate the need for parts of the flexprints, further reducing the multiple scattering. The combined performance of the enhanced detector setup together with the high stopping rate will allow to search for the $\mu \rightarrow eee$ decay with a sensitivity of $B(\mu \rightarrow eee) \leq 10^{-16}$. Whilst we always keep this ultimate goal in mind, the rest of this document is concerned with the phase I detector for existing beamlines.

PART II



BEAMLINER



MILESTONES FOR THE BEAMLINE

1. Realistic compact optical design that allows the Mu3e superconducting solenoid to be placed in the front-part of the $\pi E5$ area, sharing similar features also required by the MEG II experiment and so enabling a minimal change-over between experiments. Furthermore, a maximum transmission of 28 MeV/c surface muons, with minimal beam-related background and minimal beam emittance is required, without compromising the physics goals of either experiment.
2. Test commissioning of the beam line up to the injection-point of the solenoid with a rate close to $1 \cdot 10^8 \mu/s$. This is considered as “proof-of-principle”.
3. Implementation of improvement found necessary during the test commissioning and the maximizing of the transmission efficiency.
4. Once the solenoid is implemented into the beam line, final optimization of injection and the stopping regions.

The current status of the beam line is such that the first two milestones have been achieved and the resulting modifications from the test commissioning are now being assessed for implementation



MUON BEAM

4.1 Beam Requirements

A coincidence type experiment such as Mu3e, with a stage I aimed sensitivity of $\mathcal{O}(10^{-15})$ not only requires running at the intensity frontier, but also, as in the case of the PSI ring cyclotron complex, benefits from the DC proton beam structure which yields a lower instantaneous muon rate and is therefore less prone to pileup and accidental overlap events than at a pulsed machine.

The essential requirements for such a stopped muon beam are:

- a copious supply of low-energy muons, so-called “surface muons”, produced from stopped pion decay at the surface of the primary production target [9]. The maximum intensity peaks at around 28 MeV/c, close to the kinematic edge of the two-body momentum spectrum of pion decay at rest as can be seen in the measured muon momentum spectrum shown in Figure 4.1;
- compact high transmission optics at 28 MeV/c;
- a small beam emittance, allowing a minimal stopping target size;
- a moderate momentum-byte of around 10%, with an achromatic final focus, giving a balance between beam intensity and stopping density, allowing a thin muon stopping target to be used for the almost monochromatic beam. This in turn minimizes multiple scattering of both the incoming muons and the out-going secondary particles, as well as reducing photon production from

Bremsstrahlung or annihilation-in-flight and Bhabha scattering;

- Minimization and separation of beam-related backgrounds such as either Michel e^+ from μ^+ - decay or e^+ produced from π^0 -decays in the production target or from decay-in-flight particles produced along the beam line. Simulated production spectra of secondary particles produced from PSI’s Target E are shown in Figure 4.2. At 28 MeV/c the positive pion contamination is highly suppressed due to the short pion lifetime;
- Background minimization along the beam line due to material interactions such as windows or degraders. This requires the maximal extension of the vacuum system or the use of a helium environment to keep the multiple scattering to a minimum.

4.2 The Compact Muon Beam Line (CMBL) Phase I Design

As already outlined, a multi-staged approach to the experiment will be sought. For the initial phase, muon intensities between $(10^7 - 10^8)$ muons/s will be required, corresponding to the start-up scenario with just the central detectors and leading to maximum intensity when the recur stations are added. The demand for the highest muon intensities leaves only one choice of facility in the world, PSI’s $\pi E5$ channel. However, this channel must be shared with the upgrade version of the MEG experiment – MEG II [10],




Figure 4.1: Measured muon momentum spectrum in $\pi E5$, with fully open slits. Each point is obtained by optimizing the whole beam line for the corresponding central momentum and measuring the full beam-spot intensity. The red line is a fit to the data, based on a theoretical $p^{3.5}$ behaviour, folded with a Gaussian resolution function corresponding to the momentum-byte plus a constant cloud-muon background.

whose large detector and infrastructure are permanently located in the rear-part of the $\pi E5$ area. Since both experiments have similar beam requirements, a solution to house both experiments in the same area, with a minimal changeover effort, was sought. The solution found was a very compact Mu3e beam line which allows the full-sized 3.2 m long superconducting Mu3e solenoid to also be placed in the front-part of the $\pi E5$ area. This “Compact Muon Beam Line” (CMBL) solution shares the elements that both experiments require, while it requires only the replacement of the superconducting beam transport solenoid of MEG II by a dipole magnet for Mu3e.

The initial optical design of the CMBL was modelled using the beam optics matrix code programs GRAPHIC TRANSPORT [11] and GRAPHIC TURTLE [12], while the detailed modelling was undertaken using the newer GEANT4 based tool G4BEAMLINE (G4BL) [13]. An example of the CMBL optical model is shown in Figure 4.3.

The model is based on the elements of the $\pi E5$ channel from Target E up to the last dipole magnet in the shielding ASC41, just prior to the quadrupole triplet TI. In order to suppress the almost seven-times more intense inherent beam positron background associated with Michel decays in and around the production target, as well as π^0 -decay in the target producing electron positron showers from the converted photons of the decay, a 200 kV Wien-filter ($E \times B$ field separator) is used. This element together with the

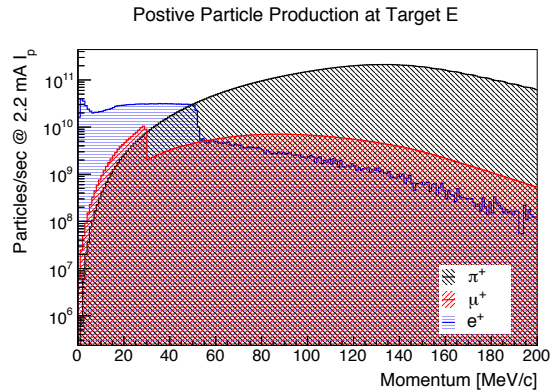


Figure 4.2: Simulated beam-related background spectra produced from PSI’s Target E. Based on 10^9 proton interactions using the GEANT4 Bertini cascade model.

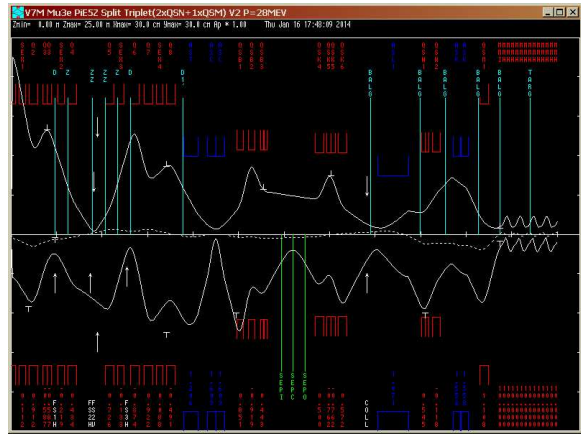


Figure 4.3: CMBL TRANSPORT solution, showing the beam envelope fits for a 28 MeV/c muon beam with an 7% FWHM momentum-byte. The upper-half shows the vertical and the lower-half the horizontal envelopes along the beam-axis (z -axis) from the origin, close to Target E, to the Mu3e solenoid. The red elements represent the multipoles, while the blue ones represent dipole magnets. The vertical/horizontal scales have different scale factors of 6 cm and 2m per division.

two sets of quadrupole triplet magnets TI, TII, which ensure an optimal transmission through the separator and downstream collimator system are shared with the MEG II beam line, since the requirements of the two experiments are the same. The remaining two dipole and three quadrupole elements are dedicated to Mu3e. The compact “split-triplet” design after the 90° dipole ASL41




Figure 4.4: (Top) - CAD model of the entire $\pi E5$ channel & CMBL used as a basis for the G4BL models. G4BL models - (Lower Left) Full CMBL beam line from Target E to the Mu3e solenoid. (Lower Right) Short version from Triplet II to the Mu3e detector solenoid.

magnet, involves the 65° bend by ASK41 and the quadrupole doublet QSO41/42 and quadrupole singlet magnet QSM41. This configuration not only halves the number of magnets that would be conventionally used (two triplets) but achieves a compact solution with similar focusing properties enhanced by the vertical edge focusing of the ASK dipole, so allowing the 3.2 m long Mu3e detector solenoid to fit into the front area.

Two G4BL models were constructed (cf. Figure 4.4), one with the full $\pi E5$ channel and Target E, simulating the whole pion production process by protons in the primary target, to surface muon production and transport to the Mu3e target, while a shorter version starts from Triplet II, just upstream of the intermediate focus at the collimator system, where measured phase space parameters determine the initial beam used for the simulation. The shorter version predictions were used as a direct comparison to the following measurement, since the long version is still being optimized.

Initial predictions from TURTLE and G4BL using previously measured MEG optimized phase space parameters at the collimator gave a transmission of $\sim 90\%$ between the intermediate focus and the final focus at the injection-point to the Mu3e solenoid with a beam size of $\sigma_x=27$ mm $\sigma_y=23$ mm (stage III).

4.3 CMBL Test Commissioning

Commissioning of the CMBL test beam layout was undertaken in a staged approach in two 4-week beam periods November and December 2014 and May 2015, using mostly existent elements. The commissioning was performed in three stages, as shown in figure 4.5. Stage I marks the intermediate focus at the collimator system, Stage II the exit from the first part of the split triplet while Stage III is the injection-point into the future Mu3e detector solenoid.

Optimization, phase space and intensity measurements were made at each of the three stages of the setup, using a 2-D automated scanner system with a pill scintillator (2 mm diameter, 2 mm thick) mounted on a Hamamatsu R9880U miniature photomultiplier. A typical setup is shown in Figure 4.6.

Stage I beam tuning yielded a surface muon intensity of $1.1\text{-}1.3 \times 10^8 \mu^+/\text{s}$ at 2.2 mA of proton current on the production Target E, depending on the proton beam centering. An example of a measured beam profile together with a Gaussian fit yielding an intensity of $1.14 \times 10^8 \mu^+/\text{s}$ at 2.2 mA proton current is shown in Figure 4.7.

The measured separation quality achieved by the Wien-filter is shown in figure 4.8. This rep-




Figure 4.5: Commissioning test setup of the CMBL showing the staged setup.

resents a separator magnet current scan taken at low pill-counter threshold so as to capture both positron and muon components. Based on the calibration of the system the separation between muon and positron peaks at the collimator represents $5.7 \sigma_\mu$ for the transmission optics to the solenoid. This corresponds to a physical separation between peaks of 85 mm.

Comparative results of both measured and simulated quantities for Stage I are shown in Figures 4.9, 4.10 and 4.11, these are based on measured multi-envelope versus quadrupole current fits in TRANSPORT. The so determined phase space upstream of the last quadrupole is propagated back through the magnet using the final settings to the position of the pill-counter at the intermediate focus and compared to the measured profiles. Figure 4.9 shows good agreement between the reconstructed and measured phase space at the intermediate focus.

Further very good agreement is demonstrated in figure 4.10 which shows a comparison between the measured horizontal and vertical beam-spot

sizes at the final focus (Stage III) and the corresponding G4BL simulated ones for varying QSM41 quadrupole currents.

Figure 4.11 shows the predicted beam size at the injection point to the Mu3e solenoid (Stage III), based on the above phase space parameters, and the measured profiles. The agreement is seen to be very good. The final focus measurements made here show a transmission factor of 68% giving $\mathcal{O}(10^8) \mu^+/s$ at 2.2 mA, so achieving proof-of-principle for this test setup using partly non-optimized elements. However, this is less than expected. The G4BL study shows that the main losses can be attributed to aperture constraints on the dipole vacuum chambers, in particular the 60° bending magnet ASK41.

The present yoke gaps for both dipoles are 18 cm, which can be increased to 24 cm, but necessitates new vacuum chambers, also for the quadrupole magnets “cross-slit” chambers would help. Simulations show that implementing these changes would increase the transmission efficiency from 63% to 88% thereby achieving the transmis-




Figure 4.6: Shows scanner system setup with pill scintillation counter mounted behind the beam window.

sion of approximately $1 \times 10^8 \mu^+$ /s at 2.2 mA proton current. The comparative results of this simulation are shown in figure 4.12, where the transmission efficiency versus beam-axis coordinate between the collimator system (Stage I) and the injection to the solenoid (Stage III) is shown.

The central detector region inside the solenoid magnet is dominated by a beam vacuum-pipe of inner diameter 60 mm containing a 600 μm thick MYLAR degrader followed by a 2 cm thick lead collimator and ending with a 35 μm MYLAR vacuum window. A double-cone MYLAR target of radius 19 mm, length 100 mm and total thickness of 160 μm (see chapter 6) is located close to the window at the centre of the solenoid. The warm-bore of the solenoid is filled with helium gas essentially at atmospheric pressure to reduce multiple scattering. The collimator system within the beam-pipe minimizes the number of muons hitting the inner silicon detector.

It is the inner silicon detector diameters and the associated beam-pipe size that determine the

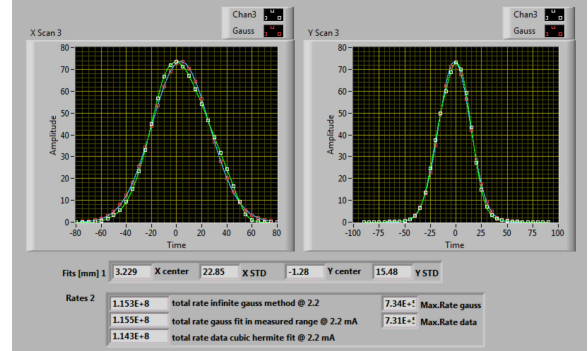


Figure 4.7: Example of Stage I collimator position beam envelope scan results with incorporated fitting $\sigma_x \sim 23$ mm, $\sigma_y \sim 15$ mm with an intensity of $1.14 \times 10^8 \mu^+$ /s at 2.2 mA of proton current.

inner bore layout of the experiment, in particular the stopping target diameter, which has been maximized to a radius of 19 mm. This however is not well matched to the minimum possible beam size but is a compromise between stopping rate, occupancy and the number of readout channels. Based on the G4BL simulation together with the measured beam emittance at the injection-point to the solenoid, a relative beam-loss of 41% is expected at the entrance to the beam-pipe compared to the intensity at the final focus (cf. Figure 4.13). This implies approximately $6 \times 10^7 \mu^+$ /s would reach the target. The expected beam size at the target is a round spot of rms width between 8-9 mm. The beam size together with a stopping efficiency of 90% would yield a stopping rate of about $5 \times 10^7 \mu^+$ /s at 2.2 mA.

4.4 Area Layout & Infrastructure

Due to the spatial restrictions in the $\pi\text{E}5$ front area and the substantial infrastructure needs of the experiment, an optimized area layout is necessary. Both electrical power requirements and cooling-water needs must be upgraded and require modifications to their distribution. Furthermore, due to safety requirements an additional new access route must also be envisaged to enter the front area. This is under study, with radiation safety dose-rate measurements already performed along the planned route. So from the safety point-of-view this is a feasible option. The current concept plans to incorporate the zone access with two infrastructure platforms above the area, placed on

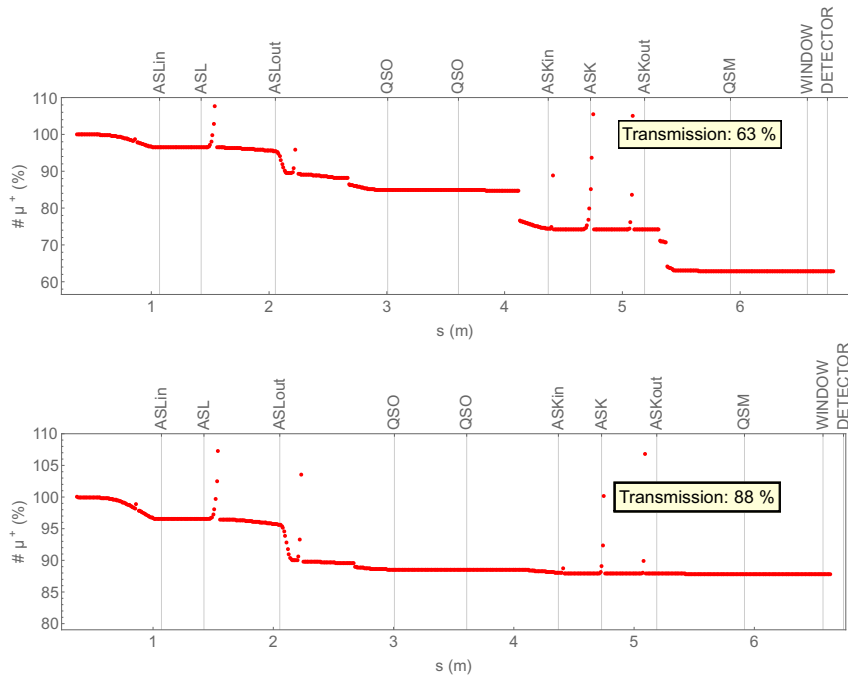


Figure 4.12: (Top) shows current simulated transmission efficiency from the collimator to the injection point of the solenoid. The main losses are at the second dipole magnet ASK41. (Bottom) shows the comparative transmission efficiency when wider aperture vacuum chambers are used, yielding lower losses.

the shielding walls. Removable stairways to an intermediate small platform placed above Triplet II, which will also be removable, would serve as an access route from the outside, rear-side of $\pi E5$ to the front area via a walkway placed on the rear shielding wall. A new controlled access door would be placed at floor-level outside of the rear area giving access to the walkway. This current planned area layout is shown in Figure 4.14.

The lower stairway closest to the magnet will be easily removable, allowing the magnet to be moved on an air-cushion system for ease-of-access.

A local 2.5-3 ton crane system is also envisaged for use with the magnet, primarily to remove the upstream and downstream shielding doors.

Of the two planned infrastructure platforms, the one above the beam-entrance wall must also be removable, in order to service the $\pi E5$ channel elements once per year during the accelerator shutdown period. This platform being closest to the magnet and detector will house the cooling elements such as compressors for the cryogenic cold-heads as well as the secondary water cooling circuits for the magnet and detector. The

connections to this platform will be made removable. The other larger platform will not be removable and will carry the magnet power supplies, quench detection system and electronics and power-control circuitry associated with both magnet and detectors.

Outside of the area a control room, either as a dedicated separate barrack or as part of a larger facility control room, sharable with other experiments is in planning. This together with a separate dedicated barrack housing the front-end computing farm will serve to control the experiment.

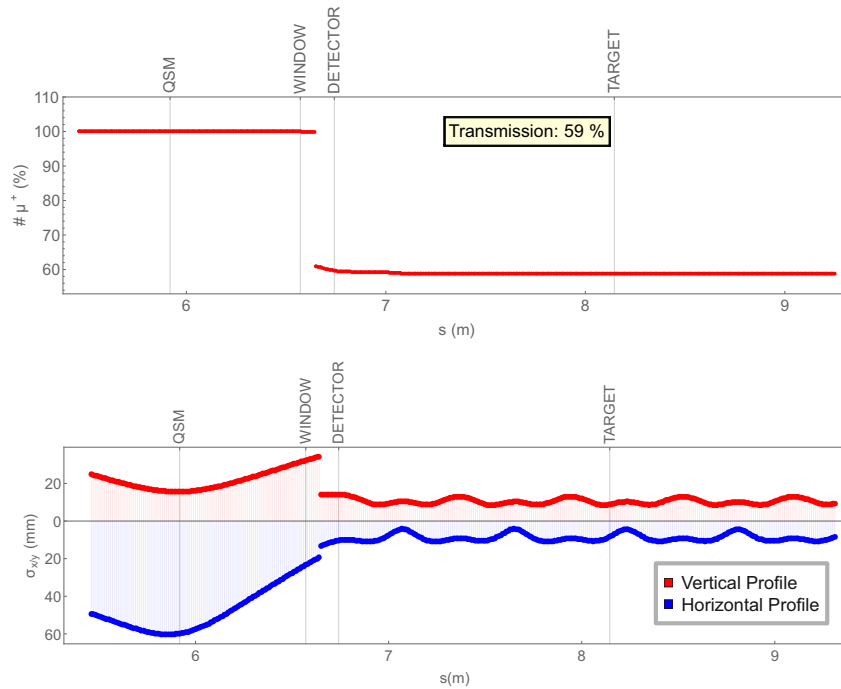


Figure 4.13: (Top) shows the transmission efficiency from the last quadrupole QSM41 to the end of the solenoid. (Bottom) shows the corresponding beam envelope sizes.

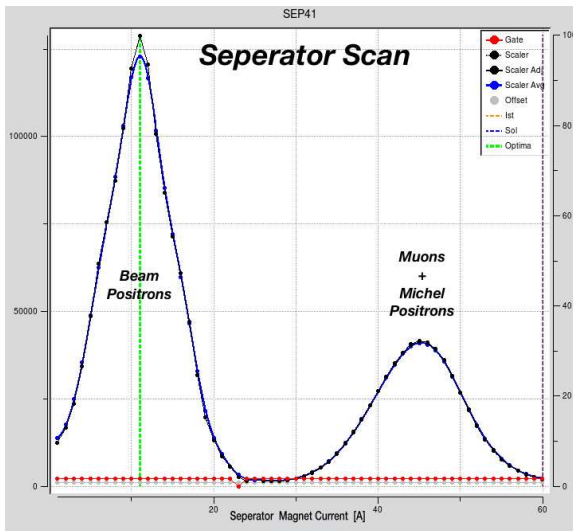


Figure 4.8: Separation quality at the intermediate focus by scanning the Wien-filter magnet current.




Figure 4.9: Comparison of reconstructed and measured beam envelope sizes versus quadrupole current obtained from phase space determined upstream and propagated to the measurement point. Good agreement is seen.




Figure 4.14: The planned $\pi E5$ Area layout for the Mu3e experiment. Two infrastructure platforms and a stairway access structure are shown in purple. The access walkway is planned along the rear shielding wall, ending parallel to the left-side wall at floor-level.




Figure 4.10: Comparison of simulated and measured beam-spot sizes at the intermediate focus versus quadrupole current

Figure 4.11: Simulated and measured 2-D beam-size at the Mu3e solenoid injection point, showing reasonable agreement.



5.1 Requirements

The magnet for the Mu3e experiment has to provide a homogeneous solenoidal magnetic field for the precise momentum determination of the muon decay products. In addition it will also serve as beam optical element guiding the muon beam to the target. The basic parameters of the superconducting solenoid magnet are given in Table 5.1. The outer dimensions also include an iron field shield.

The nominal magnetic field strength is 1.0 T in the central part, providing the optimum bending radius in terms of resolution for the proposed experimental design. A higher magnetic field would lead to a loss of acceptance as the low momentum particles would not reach the central outer pixel layers (see Figure 5.1). A lower magnetic field would lead to less magnetic deflection at constant multiple scattering, leading to worse momentum resolution (see Figure 5.2). For systematic studies and to allow for possible reuses of the magnet for other experimental measurements, the field can be varied between 0.5 and 2.0 T.

The dimensions of the cylindrical warm bore of the magnet are 1 m in diameter and 2.7 m in length. The minimum diameter is given by four times the bending radius of the highest momentum (53 MeV/c) decay products at the field of $B = 0.8 \text{ T}$ ¹ plus the target diameter. In addition the detector support and extraction rail system has to be taken into account when choosing the warm bore diameter.

¹This is the smallest field at which we still want to be able to measure re-curling tracks.

The total length is a compromise between geometric acceptance for recurling particles and the very tight space constraints for the phase I experimental area at $\pi E5$. In principle a longer solenoidal magnet would provide an intrinsically more homogeneous field. At both ends of the magnet it is foreseen to have full access by means of removable flanges.

5.2 Magnet Construction

Whilst the ideal magnet would have a constant field throughout the inner volume, real solenoid magnets show a drop in field to 50% at the end of the coil. The simplest solution would be a longer magnet, which however does not fit inside the phase I area $\pi E5$. Instead, correction coils at both ends of the magnet are appended, such that the high field region can be extended. The magnet design foresees five equal coils with a single power

MAGNET PARAMETER	VALUE
field for experiment	1.0 T
field range	0.5 – 2.0 T
warm bore diameter	1.0 m
warm bore length	2.7 m
field description $\Delta B/B$	$\leq 10^{-4}$
field stability $\Delta B/B$ (100 days)	$\leq 10^{-4}$
outer dimensions: length	< 3.2 m
width	< 2.0 m
height	< 3.5 m

Table 5.1: Properties of the Mu3e magnet




Figure 5.1: Reconstruction efficiency as a function of track transverse momentum for different magnetic fields, with recurl stations and without fiber detector.




Figure 5.2: Momentum resolution as a function of track momentum for different magnetic fields, with recurl stations and without fiber detector.

supply plus two compensating coils at the ends with separate power supplies. The field change along the z-axis has to be taken into account for the reconstruction of tracks in the recurl stations by using a look up table for the field map. Choosing the right granularity for the look-up table, a linear interpolation of the field will be sufficient to reach an approximation of $\Delta B/B \leq 10^{-4}$. For the fast online selection of events the assumption of a constant field in the active part of the experiment will be sufficient. Although the assumption of a constant (maximum) field leads to a systematic bias towards larger momenta and an increase

of online selected background events from internal conversion muon decays, no signal events will be lost.

The superconducting magnet is made from five coils of equal size, which has advantages over one long coil in terms of mechanical stability and ease of production. The small dips in the magnetic field can be treated numerically in the same way as the roll-off of the field towards the ends of the magnet. The choice for superconducting wires or conductors is driven by commercial availability, since standard superconducting components allow for the desired 2.0 T maximum field strength. A warm normal conducting magnet is no option because of size, cost (copper price), and worse operational stability. Superconducting magnets have an intrinsic immunity against absolute field changes, as they have to run at a constant (low) temperature. A dry cooling system has been chosen because this will make reliable long-term operation easier.

There will be a magnetic shield around the magnet. The shielding is required since the experimental hall is densely populated with other experiments and infrastructure. Also for the read-out of the proposed experiment it will be much easier to work in a low field environment. A beneficial side effect of the shielding is a gain of field homogeneity inside the magnet and less field dependence on variation of outside parameters.

The long term stability of the magnetic field should be $\Delta B/B \leq 10^{-4}$ over each 100 days data taking period. This can be achieved by using state of the art magnet power supplies and by permanently monitoring the absolute field with Hall probes inside the apparatus.

The cool-down time for a system of the projected size will be one week and the ramp up time will be in the order of a few hours.

The company Danfysik² was assigned to develop and produce the Mu3e solenoid magnet with an expected delivery to PSI in fall 2016. Danfysik has prepared the final technical design report of the complete magnet system.

5.2.1 MAGNETIC BOTTLE EFFECTS

A slightly lower field in the center of the magnet relative to higher z values could lead to a weak

²Danfysik A/S, Gregersensvej 8, DK-2630 Taastrup, Denmark




Figure 5.3: Axial field along the z axis in the configurations used for the magnetic bottle study. All maps are normalized to 1 T field at the magnet center. The lowest curve (black) represents a pure solenoid (configuration 0), whilst the top curve (dark green) has a strong over-current in the compensation coils, leading to a 0.12 T deep bottle (configuration 11).

magnetic bottle effect, reflecting particles back towards the central detector region emitted close to perpendicular to the field lines. This in turn would have adverse effects on the central detector occupancy and ultimately the reconstruction performance. In order to study the magnetic bottle effect, we have generated a series of field maps with various current strengths in the compensation coils (see figure 5.3), and then simulated the detector occupancy using the full simulation described in chapter 27. Whilst the bottle effect is clearly visible, its overall magnitude is small even for the rather extreme configurations studied, as can be seen in figures 5.4 to 5.7.




Figure 5.4: Rate of pixel hits in the complete phase Ib detector ($1 \cdot 10^8$ muon stops/s) for the field configurations shown in figure 5.3.




Figure 5.5: Rate of pixel hits in the innermost (vertex) layer of the phase Ib detector ($1 \cdot 10^8$ muon stops/s) for the field configurations shown in figure 5.3.




Figure 5.6: Rate of pixel hits in the outermost central layer of the phase Ib detector ($1 \cdot 10^8$ muon stops/s) for the field configurations shown in figure 5.3.




Figure 5.7: Rate of pixel hits in the outer layer of the backward phase Ib recurv detector ($1 \cdot 10^8$ muon stops/s) for the field configurations shown in figure 5.3.



STOPPING TARGET

The main challenge for the design of the stopping target is to optimize the stopping power on one hand and to minimize the total amount of material on the other hand in order to reduce both backgrounds and the impact on the track measurement. Therefore the stopping target should contain just enough material in the beam direction to stop most of the muons, which is facilitated by a degrader in the final part of the beam line but should be as thin as possible in the flight direction of decay electrons measured in the detector acceptance. Usage of a low Z material is advantageous as photon conversion and large angle Coulomb scattering are suppressed. In addition, the decay vertices should be spread out as widely as possible in order to reduce accidental coincidences of track vertices and to produce a more or less even occupancy in the innermost detector layer.

6.1 Baseline Design

These requirements can be met by a hollow double cone target à la SINDRUM [14,15]. In our baseline design (see Figure 6.1), the target is made from $75\ \mu\text{m}$ of Mylar in the front part and $85\ \mu\text{m}$ Mylar in the back part, with a total length of $100\ \text{mm}$ and a radius of $19\ \text{mm}$. This leads to an incline of 20.8° of the target surface with regards to the beam direction. The projected thickness is thus $211\ \mu\text{m}$ for the front and $239\ \mu\text{m}$ for the back part, giving a total of $450\ \mu\text{m}$ of Mylar corresponding to 0.16% of a radiation length. The mass of the

Mylar in the target is $0.671\ \text{g}$. The total area of the target is $6386\ \text{mm}^2$.

Stopping $1 \cdot 10^8\ \text{Hz}$ muons in the target corresponds to about $0.05\ \text{mW}$ of power. Compared with the power dissipation of the sensor chips, this is negligible and easily taken care of by the helium cooling.

6.2 Alternative Geometries

We have studied the stopping power and material budget for a variety of target shapes (see figure 6.2 and table 6.1) and found that for the given




Figure 6.1: Dimensions of the baseline design target. Note that the material thickness is not to scale.



Target shape	Thickness for > 90% stopping μm	Mass g
Inclined plane	160	0.72
Garland	240	1.35
Reverse garland	90% not reachable	
Two-turn garland	200	1.07
Reverse two-turn garland	200	1.07
Double cone	70/80	0.623

Table 6.1: Comparison of target shapes at 90% stopping power, assuming a 600 μm degrader.




Figure 6.2: Target shapes studied. a) Is the default hollow double cone, b) a simple plane, c) a single-turn garland and d) a double-turn garland. For the chiral shapes c) and d), both senses of rotation were tried.

beam parameters and geometrical constraints, the double cone offers the highest stopping fraction with the least material. The simulation was performed with Mylar as the target material, a previous study using aluminium however gave very similar results.

6.3 Production

At PSI, a manufacturing procedure was developed and a full scale target was produced which could be used for the phase I of Mu3e data taking. This first target has still a slightly different diameter as the final and optimized dimension was not yet fixed at that time. Nevertheless the manufacturing procedures will remain unchanged and only the tools need to be adapted to the final dimensions.




Figure 6.3: Procedure to machine the correct shape of the Mylar foils. A bunch of foils is clamped between two aluminum plates and machined according to the winding off surface of the cone.




Figure 6.4: Examples of Mylar and aluminum foils machined with the correct shape for the production of the cones.




Figure 6.5: The moulds are used to shape the Mylar foil during a heat treatment and the gluing procedure.




Figure 6.6: The glued cone is cut at the right length and perpendicular to the vertical axis with a turning steel.




Figure 6.7: Prototype of the double-cone muon stopping target made of Mylar foil. The target is mounted on a stick for handling and demonstration purposes

Most likely a new target with the final dimensions will be produced for the mechanical test of the target support scheme.

Basically, each single hollow cone of the double cone structure is manufactured by itself and is a sandwich structure consisting of 2 or 3 rolled up thin Mylar foils glued together with epoxy glue. The thickness of the individual Mylar foils and the combination of several foils are chosen to match best with the desired final thickness. Finally the two individual cones are glued together to build up the hollow double cone structure.

In the first step of the manufacturing procedure, the shape of the Mylar foils is prepared. For this purpose a bunch of foils is clamped between two aluminium plates and machined according to the winding off surface of the cone (see figure 6.3 and figure 6.4).

In order to release the restoring forces of the later bent Mylar foil, each foil undergoes a heat treatment at around 130°C before the gluing procedure. During this treatment the foil is wound around a dedicated steel cone and inserted in a Teflon mould, both with the shape of the final cone (see figure 6.5).

After that the individual foils for a single cone are glued together. To achieve the correct shape the gluing takes place in a special mould which is made out of Teflon and from which the cone can be removed rather easily due to the repellent property of the Teflon surface (see also figure 6.5).



In case the cone consists of two layers of Mylar the two foils are rotated by 180° with respect to each other so that the seams of the abutting edges of the two foils are vis-à-vis, in case that three foils are building up one cone the foils are rotated by 120° .

After the glue is cured and the cone is removed from the Teflon mould the cone is mounted between a steel spike and cap on a turning machine. With a turning steel the sandwiched foil is cut at the right length of the cone. It is important to note, that with this set-up the cut is perpendicular to the central axis of the cone and not perpendicular to the outer surface. This allows that the upstream and downstream cone can be glued together facing each other with parallel surfaces where a tiny amount of glue can be applied.

Triggered by the experience from the MEG experiment there is a discussion ongoing whether to

implement dedicated holes in the target for alignment and track reconstruction purposes. To prove the feasibility a kind of stamping technique was used to apply several holes with a diameter of 1.5 mm in the cone.

Figure 6.7 shows a prototype for the muon stopping target made of Mylar foil. Please note, that the radius is only 10 mm instead of 19 mm. The target is mounted on a stick for handling and demonstration purposes.

6.4 Support

The target support will be realized by three thin nylon lines at each side of the target. The nylon lines are attached to the metal end rings of the vertex layer support structure for the pixel detector. These six strings in total define the target position in all three space coordinates.

PART III



PIXEL DETECTOR



MILESTONES FOR THE PIXEL DETECTOR

1. Produce mechanical prototype for inner layers.
2. Produce a large sensor.
3. Develop wafer test hard- and software.
4. Produce flexprint prototype.
5. Bond onto flexprint.
6. Fabricate inner layer module.
7. Develop module test hard- and software.
8. Produce full inner pixel detector.
9. Electrical, mechanical and cooling integration of the inner pixel layers.
10. Commissioning of inner pixel layers in the muon beam.
11. Produce flexprint for outer layers.
12. Fabricate outer layer module.
13. Produce full central outer pixel detector station.
14. Electrical, mechanical and cooling integration of the central outer pixel station.
15. Commissioning of the central pixel station in the muon beam.
16. Produce recurl pixel detector stations.
17. Electrical, mechanical and cooling integration of the recurl pixel stations.
18. Commissioning of the full pixel detector in the muon beam.



THE MU3E PIXEL SENSOR

The Mu3e pixel tracker is based on a dedicated High-Voltage Monolithic Active Pixel Sensor (HV-MAPS) thinned to $50\ \mu\text{m}$, called MUPIX. Signal and power lines are aluminum traces on a Kapton flex-print, which, together with a folded Kapton foil frame, also serves as a support structure. The detector is cooled with gaseous helium.

In section 7.1 the chosen HV-MAPS technology is described and motivated. Detailed sensor specifications are given in section 7.2. It is followed by the description of the required steps to achieve this goals in section 7.3. Then the different available and planned prototype versions are described in section 7.4. This is followed by section 7.5, describing a stand alone telescope setup which demonstrates the successful integration of multiple chips in one tracking system. A detailed description of performed measurements and obtained results with the MUPIX prototypes in the laboratory as well as numerous test-beams can be found in section 7.6.

7.1 HV-Maps Sensor

We use High-Voltage Monolithic Active Pixel Sensors (HV-MAPS) as tracking detectors as they integrate sensor and readout functionalities in the same device and thus greatly reduce the material budget. Classical concepts like hybrid designs have usually a higher material budget due to additional interconnects (bonds) and extra readout chips, which compromise the track reconstruction performance especially at low track momentum.

In the first MAPS designs ionization charge was collected mainly by diffusion with a time constant of several hundreds of nanoseconds. MAPS designs with high bias voltages exceeding $50\ \text{V}$, however overcome this problem by collecting charge via drift and provide time resolutions of better than $15\ \text{ns}$.

We will use the High-Voltage MAPS design with the amplifier electronics completely implemented inside the deep pixel N-well, which was first proposed in [16] and since successfully tested [17, 18], see also section 7.6. Figure 7.1 shows a sketch of a High-Voltage Monolithic Pixel Sensor. The readout circuitry at the chip edge allows an efficient zero suppression of pixel information and the implementation of timestamps to facilitate the assignment of hits between different pixel layers.

For readout designs providing $10 - 20\ \text{ns}$ time resolution power consumption of about $250\ \text{mW}/\text{cm}^2$ are expected [19, 20].

Because of the small size of the active depletion zone, the detectors can be thinned down to $50\ \mu\text{m}$ or less. By thinning the material budget can be




Figure 7.1: Sketch of the HV-MAPS detector design from [16].



Pixel Size [μm^2]	80×80
Sensor Size [cm^2]	2×2
Sub-module assembly	1×3
Sub-module assembly size [cm^2]	2×6
Thickness [μm]	50
Bond Pad Size [μm^2]	150×150
Max. LVDS links	3
Bandwidth per link [Gbit/s]	1.25
Spatial resolution μm	≤ 50
Time resolution [ns]	≤ 20
Efficiency [%]	≥ 99
Signal to noise	≥ 20
Power consumption [mW/cm^2]	≤ 400
Operation temperature range [$^\circ\text{C}$]	0 to 70

Table 7.1: Sensor specifications MUPIX

significantly reduced to 0.1% of X_0 per tracking layer and becomes comparable to gaseous detectors.

A further advantage is that HV-MAPS can be implemented in a “cheap” commercial process. We use the AMS/IBM 180nm HV-CMOS process [21], which was developed mainly for the automotive industry, and thus offers long-term availability as well as specifications covering a wide range of operating conditions. The process offers a maximum reticle size of about $2 \times 2 \text{ cm}^2$. A new process (AH18) produced by AMS alone allows changing the substrate resistivity from currently $20 \Omega \text{ cm}$ to $80 \Omega \text{ cm}$.

7.2 Sensor specification

The Mu3e HV-MAPS chip (MUPIX) is tailored to build a silicon tracking detector with very low radiation length to limit multiple coulomb scattering. 10^8 muon decays on target per second together with a trigger-less readout require on-chip zero-suppression and high data output bandwidth. In addition the high event rate makes a good time resolution necessary in order to suppress combinatorial background.

7.2.1 MECHANICAL DIMENSIONS

We plan to use one single type of sensor in the Mu3e experiment for the inner layers and the outer layers, see table 7.1. This sensor chip has a size of $2 \times 2 \text{ cm}^2$. If the yield permits, we will cut strips of three subsequent sensors from the wafers and mount them in one piece, ensuring

good relative alignment. The chips are thinned to $50 \mu\text{m}$ corresponding to 0.054% of one radiation length X_0 in order to reduce the multiple coulomb scattering which dominates the spatial resolution. The pixel size of $80 \times 80 \mu\text{m}^2$ contributes only slightly to the spatial resolution of $\leq 50 \mu\text{m}$. Half a millimeter of the pixel sensor at the edge is assumed to be inactive, the rest is filled by the pixel matrix. The inactive region is shared between the bond pads and the on-chip digital readout. It has to be scrutinized whether or not the inactive region can be kept this small, since extrapolation from current MUPIX prototypes suggest that 2mm might be inactive.

7.2.2 FEATURES AND PERFORMANCE

A large fraction of the front end readout electronics is integrated into the MUPIX chip itself, since an additional readout chip in the active tracking volume would lead to extra multiple coulomb scattering. This includes the pre-amplification of the charge signal, discriminators, time to digital converters (TDC), the readout logic with zero-suppression and multiple high speed serial data outputs. The circuits for the high speed serial data output and the per pixel TDCs require very clean clock signals, which can best be generated with the help of an integrated voltage controlled oscillator (VCO.) The MUPIX chip has a specified efficiency of $\geq 99\%$ combined with a signal to noise ratio of ≥ 20 for a ^{55}Fe -signal, which is comparable to the signal from an electron track in Mu3e. A time resolution of $\leq 20 \text{ ns}$ for the time stamp produced when a MUPIX pixel is hit must be achieved in order to reduce combinatorial background. To operate efficiently in the Mu3e readout system, the data transmission has to have a bit error rate of less than 10^{-14} per link and must use 8b/10b encoding for DC balancing. At the same time the power consumption has to be kept at $\leq 400 \text{ mW}/\text{cm}^2$ so the MUPIX chips can be cooled by a flow of helium gas as described in section 19. The temperature range in which the MUPIX sensor has to be fully operational is between 0 and 70°C . A dedicated temperature sensor must be implemented in the MUPIX chip, which will be read out before and during the physics data taking. It would be convenient if the temperature could be measured with the help of an ADC inside the MUPIX chip and read out via the slow control bus. Switching off individual pixels via



slow control must be possible, in order to mask noisy pixels which would otherwise clog the read-out system. On top of this and the current slow control functionality it must be possible to read back the register values from the MUPIX.

7.2.3 READOUT AND FLEX-PRINT INTEGRATION

For the inner layers three MUPIX chips and for the outer layers nine MUPIX chips are mounted together on one flex-print, see section 8.2. These flex-prints provide the operating voltages 1.5 V, 1.8 V the high voltage of -90 V, ground potential and all signal traces for the readout, slow control and monitoring of the MUPIX chip. All other voltages like threshold and baseline have to be generated on the MUPIX chip. The MUPIX output data consists of time-stamps and pixel addresses of hits plus header words, serialized on fast low voltage differential signaling (LVDS) links running at 1.250 Gbit/s. For the two innermost detector layers three high speed links per chip are needed, the sensors used in the outer layers need only one serial output. The sensor is configured via a serial programming interface, which must be implemented with LVDS in- and outputs in order to mitigate cross-talk. All other signals like clock, synchronous reset, injection and hit monitor are implemented in LVDS standard as well. The $100\ \Omega$ termination for all LVDS input signals must be integrated on the MUPIX chip, so the LVDS termination can be done on the last chip on a flex-print, since the extra material of external resistors would be unacceptable. Stringent space constraints on the flex-print limit the number of pads (and thus bond wires) including those for supply voltages to about 30 per MUPIX chip. These pads have to be on one side of the chip below the digital part, limiting the inactive region to one side of the chip. With the help of an overlap between this inactive region and the active pixel matrix of the neighboring MUPIX good geometrical acceptance of the detector can be restored, see section 8.1. Only ten extra traces per individual chip can be accommodated on the flex-print, all the remaining signals and voltages have to be shared with the other MUPIX chips on the same flex-print. Bond pads must be $\geq 150 \times 150\ \mu\text{m}^2$ to ensure successful module production with high yield. In particular the application of a new bonding technique where the aluminum traces of the flex-print are directly

bonded without extra wires is only possible with bond pads of at least $150 \times 150\ \mu\text{m}^2$ size.

7.3 Path towards the Full Sensor

In the following section a road-map from the current MUPIX prototypes to the MUPIX chips used for the phase I Mu3e detector will be outlined.

At the current stage a fully functional small HV-MAPS prototype, the MUPIX 7, is available. While the MUPIX 7 chip already has a large fraction of the features required for Mu3e operation, some features still have to be added.

In order to go from the small $0.3 \times 0.3\ \text{cm}^2$ MUPIX 7 chip to a full size sensor, it is planned to submit a larger chip (MUPIX 8) of $1.3 \times 2.3\ \text{cm}^2$ in 2016. This MUPIX 8 will have an analog and digital part similar to the one of MUPIX 7. All pads of the MUPIX 8 will be on the same side below the digital part and of $\geq 150 \times 150\ \mu\text{m}^2$ size which is necessary for chip integration on a module. The MUPIX 8 will be used to study the effects of much longer pixel columns and to engineer the mechanical and electrical integration on sub-module assemblies. The changes between MUPIX 7 and MUPIX 8 should include the generation of baseline and threshold voltage on chip, slow control over one single bus, read-back of the chip DAC settings and masking of hot pixels. Ideally comparator ringing observed for MUPIX 7, more stable baseline tune settings (TDACs), a better HV distribution grid and $80 \times 80\ \mu\text{m}^2$ pixel size can be addressed in the design of MUPIX 8. Another important step will be the increase of the substrate resistivity from $20\ \Omega\ \text{cm}$ to $80\ \Omega\ \text{cm}$. The increased substrate resistivity will lead to a larger depletion zone and to larger signal to noise ratio. If the MUPIX 8 is operational, it is planned to build full vertex modules based on MUPIX 8.

The next step, MUPIX 9, would be a bug-fix of MUPIX 8. MUPIX 9 is also planned to have all features detailed in 7.2. If all is going well and MUPIX 9 is fully operational, MUPIX 10 (MUPIX X) will be a full sized $2 \times 2\ \text{cm}^2$ chip with all features and the performance to build the Mu3e phase I detector.

7.3.1 FUTURE MUPIX IMPROVEMENTS

Since the Mu3e experiment can profit from an improved MUPIX performance, several features beyond the baseline design are discussed.



On-chip DC-DC converters

To keep the material budget of the pixel tracker layers as low as possible, it should be envisaged to integrate DC-DC converters in the MUPIX chips, such that one single input voltage of 20 V can be used to generate 1.5 V, 1.8 V and possibly even the high voltage for the substrate on-chip. Using such a supply voltage at low currents will ensure a low voltage drop over the long flex-print traces of the outer pixel layers. Since the DC-DC conversion on-chip can lead to x-talk, wrong voltages or worse results, this feature should be developed in parallel and should only be implemented in a production version of the MUPIX after successful silicon proof with small scale prototypes.

Time-walk Correction

The time resolution of the current MUPIX 7 chip of around 10 ns is largely dominated by time-walk. Ongoing studies indicate that the intrinsic charge collection time via drift is in the order 400 ps and time resolution for single hits of ≤ 1 ns are achievable, greatly facilitating the reduction of combinatorial background. Time-walk correction can be introduced with a fully analogue circuit or by combining dual thresholds at the rising edge of the charge signal with a look up table in the digital part. Studies with small scale prototypes will be carried out to compare the different techniques and qualify them for application in a large sensor.

7.4 The MUPIX Prototypes

First purpose-built sensor prototypes (the MUPIX series of chips) became available in 2011. A brief summary of the chip development is given below.

7.4.1 MUPIX 1 AND 2

The MUPIX 1 and 2 are small demonstration prototypes with a matrix of 42×36 pixels of $30 \times 39 \mu\text{m}^2$ size for an active area of approximately 1.8 mm^2 , see Figures 7.2 and 7.4. Each pixel consists of the sensor diode, a charge-sensitive amplifier and a source follower to drive the signal to the chip periphery. In addition there is a capacity allowing to inject test charges. See Figure 7.3 for an overview of the pixel electronics. On the periphery, a comparator turns the analogue signal into a digital time-over-threshold (ToT) signal. The




Figure 7.2: Design view of the MUPIX 2 chip (actual size about 1.8 by 2.5 mm).

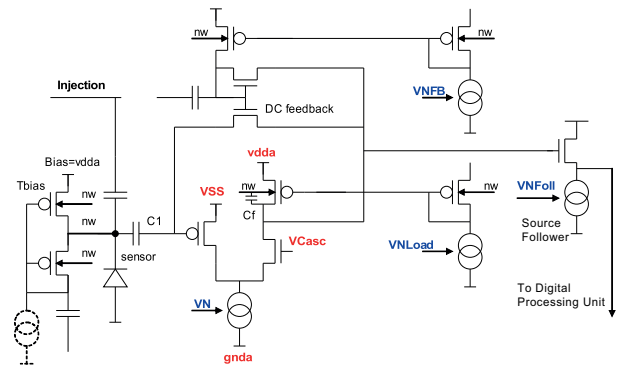


Figure 7.3: Schematic of the pixel cell analog electronics in the MUPIX chips. *nw* stands for n-well.

threshold of the comparators is set globally for the chip and adjusted pixel-per-pixel with a 4-bit tune digital to analogue converter (TDAC). In the test chips, the comparator output of an individual pixel can be observed via a dedicated output line. Alternatively, the whole chip can be read out via a shift register, where all the available information is whether a particular pixel saw a signal during an active gate.

The MUPIX 1 chip had an issue with feedback in the comparator that occasionally led to double pulses. This issue has been fixed in MUPIX 2, which in addition contains temperature sensors.

7.4.2 MUPIX 3 TO MUPIX 7

The MUPIX prototypes number 3 to 7 were deployed to develop HV-MAPS chips with fully integrated readout electronics.




Figure 7.6: Overall layout of the MUPIX 7 chip. The pixel matrix, the digital part and the periphery with the bond pads can be clearly distinguished. The red rectangle in the bottom left contains the readout state machine and serializer designed using synthesis tools, the rest of the layout was done manually.




Figure 7.4: MUPIX 2 sensor on ceramic carrier with one Euro-cent for scale comparison.

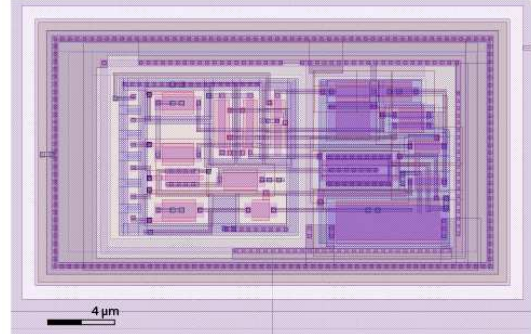


Figure 7.7: Layout of the in-pixel circuitry (first amplification stage) in the MUPIX 7.




Figure 7.5: Design view of the MUPIX 3 chip (actual size about 4 by 5 mm).




Figure 7.8: Layout of the pixel unit cell in the MUPIX 7. The nine charge collecting diodes and the in-pixel circuitry shown in figure 7.7




Figure 7.9: Layout of the digital pixel unit cell in the MUPIX 7.



In August 2012, we submitted the MUPIX 3 chip. The chip has 40×32 pixels of $92 \times 80 \mu\text{m}^2$ size for an active area of approximately 9.4 mm^2 , see Figure 7.5. It implements the full digital column logic, allowing for address generation and serial readout of zero-suppressed data. In addition, MUPIX 3 has faster signal shaping than MUPIX 2.

The main conceptual differences compared to the final sensor are the lack of a high-speed LVDS outputs and the chip-wide hit collection logic. For this prototype, the corresponding logic is emulated off-chip in a FPGA. Also several control voltages which should be generated on-chip in the final version are currently produced on a test printed circuit board (PCB) in order to allow for easier debugging.

MUPIX 3 could not be reliably used because of a bug in the masking of noisy pixel cells which led to an unpredictable number of pixels which were switched off.

MUPIX 4 was a very successful prototype which allowed full zero-suppressed readout and was the first HV-CMOS sensor with a per pixel counter TDC. The MUPIX 4 has a time resolution of better than 10 ns RMS, see below. Because of a timing violation in the readout part, the timestamps and row addresses of half of the pixels could not be read out.

This readout timing violation was fixed for MUPIX 6 which is fully functional. In addition to these bug fixes, the MUPIX 6 has for its first four columns single stage pre amplifiers and for the remaining 28 columns dual stage pre amplification. Single stage pre-amplification with low output swing was implemented already in MUPIX 4, requiring extremely stable voltages for threshold, baseline and supply voltage. The dual stage amplification ensures stable operation even for non-ideal conditions.

MUPIX 7 is the first ASIC with a fully integrated readout state machine, high speed clock generation circuitry (PLL and VCO) and a fast serial output running at 1.25 Gbit/s, see figure 7.6 for the chip layout. The analogue part of MUPIX 7 is almost identical to the one of MUPIX 6 with a total active area of 10.55 mm^2 . The in-pixel amplification circuitry and a full pixel including the charge collection electrodes can be seen in figures 7.7 and 7.8. The digital part for a single pixel is shown in figure 7.9 and the end-of-column logic is shown in figure 7.10. This full system on chip




Figure 7.10: Layout of the double end of column in the MUPIX 7.




Figure 7.12: MUPIX 7 telescope setup in the PSI testbeam in October 2015. The single plane setup and reference scintillators can also be seen on the right. In the back, the low and high-voltage power supplies as well as the two readout PCs are visible.

was successfully tested at multiple test beam campaigns 7.6.2 and in the laboratory.

An overview of the current prototypes is given in tables 7.2 and 7.3.

7.5 The MUPIX Telescope

Using the fourth MUPIX generation, a tracking telescope was developed [22]. The telescope consists of four sensor planes, see figure 7.11. To realize a flexible mechanical setup, which can be easily renewed in the case of activation during a test-beam or damage, we based the mechanics on commercially available components. The required precision is delivered by optomechanical components (Thorlabs). Two parallel optical rails, spaced by 10 cm mounted on a $15 \times 60 \text{ cm}^2$ optical breadboard pointing in beam direction, build a solid base. A pair of optical post holders connected to the rails with clamps provide a stable movable stage for each layer. The PCBs with the chips are mounted on a custom PCB holder made by the mechanical workshop of the physics institute in Heidelberg. This PCB holder is attached to two optical posts. A system of micrometer screws and springs integrated in the PCB holder give the flexibility to move the individual planes with a precision of $10 \mu\text{m}$.

The development started during autumn 2013 and first results were obtained at the DESY II

testbeam facility during the early spring 2014. The first MUPIX Telescope version used two Stratix IV FPGAs, each one controlling and reading out two MUPIX sensors. The data streams were sorted online and frames with identical 125 MHz time stamp were merged and analyzed offline. The maximal track rate for this telescope version was about 500 kHz, limited by the readout wait cycles in the FPGA state machine.

For the MUPIX 7 prototype, which operates in a streaming mode, continuously sending its own data, a single FPGA is sufficient to control and read out the four sensors, due to the reduced number of required signals. On the FPGA, the data sorter (see section 24.4.2) is running, providing sorted data frames with a range of 32 time stamp bins. This simplifies the software and opens the door for a telescope consisting of up to eight layers, which will improve pointing resolution and helps to reject the noise contribution in the tracking. The MUPIX 7 telescope can handle up to 1 MHz tracks, limited by the polling readout between FPGA and computer. Using direct memory access (DMA), which is developed for the Mu3e readout chain, the maximal track rate can further be increased to the data logging speed on a normal HDD in the order of 500 Mbyte/s. This corresponds to a track rate in the order of 4 MHz. The MUPIX 7 telescope is shown in operation in figure 7.12.

Besides the usage as tools for characterizing one of the MUPIX planes at test-beams, the telescope is used to test the integration of MUPIX prototypes in a more complex system. Due to the successful synchronous operation of a four layer telescope, which is one sensor more than half the sensor strip of the vertex detector, a milestone towards the Mu3e pixel detector has been reached.

7.6 Characterization of the Prototypes

The properties of the MUPIX prototypes have been tested using injection pulses, LEDs, laser diodes, X-rays, radioactive sources and test beam measurements. In the following, we will outline the core results, details of the findings can be found in the referenced master and bachelor theses.

The characterization setup is based on a printed circuit board which houses the chip and serves as



	MuPix 3	MuPix 4	MuPix 6	MuPix 7
Functionality	Sensor analog and digital electronics	Sensor analog and digital electronics	Sensor analog and digital electronics	Sensor analog and digital electronics fast serial readout
Bugs	Random pixel on/off	Time stamp and row address for only 50% pixels	None	Minor features in state machine control and wrong byte order in words
Improvements	First digital readout	Working digital readout and timestamp	Timestamp and address readout bug fixed	Fast serial readout

Table 7.2: Sensor prototype feature overview

	Target value	MuPix 3	MuPix 4	MuPix 6	MuPix 7
Pixel Size [μm^2]	80×80	92×80	92×80	103×80	103×80
Sensor Size [cm^2]	2×2	0.3×0.3	0.3×0.3	0.3×0.3	0.3×0.3
Active area [mm^2]	390	9.42	9.42	10.55	10.55
Assembly	1×3	1	1	1	1
Assembly size [cm^2]	2×6	0.3×0.3	0.3×0.3	0.3×0.3	0.3×0.3
Thickness [μm]	50	90, 600	50	250	50, 63, 75, 250
Max. LVDS links	3	-	-	-	1
Bandwidth [Gbit/s]	3.75	low	low	low	1.25
Spatial resolution $\leq 100 \mu\text{m}$	≤ 50	-	≤ 30	?	≤ 30
Power consumption [mW/cm^2]	≤ 400	?	?	≥ 223	≥ 225
Time resolution [ns]	≤ 20	-	≤ 16.6	≤ 15.5	?
Efficiency [%]	≥ 99	-	≥ 99	≥ 94	≥ 99
Signal to noise	≥ 20	?	≥ 31.5	≥ 10	?
Noise rate at 99% efficiency [Hz/pixel]	≤ 10	-	-	-	≤ 2

Table 7.3: Sensor prototype overview

an interface to the control FPGA. The chip test board provides voltage regulators for the supply voltages and digital to analogue converters for the threshold and the injection pulse height. To communicate with the FPGA, signals can be transmitted and received via a flat ribbon cable or shielded twisted pair SCSI cable. The FPGA is connected to a PC which allows to steer it via a graphical user interface written in C++.

The setup has been developed and improved continuously over the prototype generations.

All current prototype chips need to be provided with external baseline and threshold voltages. The chips can be configured with the help of 6-bit DACs and configuration bits on the chip, which can be set via a shift register. Further all

prototypes provide the possibility of an individual pixel tuning by adjusting the relative offset of threshold and baseline on the comparator (see chapter 30).

The measurements are performed either using the digital readout of the chip, which allows to measure all pixels at the same time or by using the so called hitbus which is the digital output signal of a single pixel's comparator. Both oscilloscope measurements and data taking via the FPGA board are used.

7.6.1 LABORATORY RESULTS

The first step of chip characterization is done in the lab, which allows for a general check of functionality and optimization of the chip's perfor-




Figure 7.11: Schematic view of the four-plane MUPix telescope.

mance. The core results and development over the prototype generations are summarized in the following.

Pixel response


We refer to pixel response as the output pulse of the pixels comparator. The pulse length corresponds to the time the analogue signal from the amplification stage crossed the threshold, the so called time-over-threshold (ToT). For pulsed or triggered signals, the first edge of the signal can be used for latency measurements, e.g. section 7.6.2. The comparator output can be observed for a single pixel with the help of the hitbus signal.

The analogue pulse can be reconstructed by using a pulsed signal source and measuring latency and ToT for different thresholds. For the earliest prototypes (Fig.7.13) this revealed pulse widths larger than $1\ \mu\text{s}$, which corresponds to a large dead time. In the later prototypes this was improved by a faster shaping which creates a pulse length of 100 ns.

The measured ToT is proportional to the initial input signal which is created by the collected charges. Therefore, the ToT is proportional to the number of collected charges and by this also directly mirrors the amount of deposited energy.

Heiko: Add energy calibration plots by Jan

By using an infrared laser with 855 nm wavelength, which has a large penetration depth in silicon compared to the depletion zone thickness, an approximately uniform charge carrier production can be achieved along the depletion zone. The fast charge pulse is dominated by charges collected via drift from the thin depletion zone. The pulse height is thus approximately proportional to the


 Figure 7.13: Measured pulse shape of a MUPix pixel in response to an LED pulse, fitted with the expectation from $RC - CR$ shaping, from [23].

depth of the depletion zone, which varies with the applied high voltage as expected, see figure 7.14.

- HV-dependence
- Pulse shape
- ToT
- ref to time resolution

Signal-to-Noise

To measure the signal-to-noise ratio (SNR) a reliable/known reference signal is needed. The signal of choice is the 5.9 keV photon of the ^{55}Fe decay, which is converting via the photo effect producing a $\approx 1\ \mu\text{m}$ diameter cloud of electron-hole-pairs. For a single pixel the iron source produces a distinct peak in the ToT spectrum (Fig.7.15), which allows to match the injection signal to reproduce




Figure 7.14: Position of the ^{55}Fe peak in dependence of the applied high voltage, from [23].




Figure 7.16: Signal to noise rate for the MuPIX 2 chip. The signal size is taken from measurements with a ^{55}Fe source, whilst the noise is measured in a injection pulse scan, see [23] for details.




Figure 7.15: Time-over-threshold (corresponding to energy) spectrum of a ^{55}Fe radioactive source, from [23].

this peak. By using the tuned injection signal the SNR can be determined with a threshold scan, which measures the s-curve of the pixel response as shown in Fig.7.16. With the help of an error function fit the mean signal height and the noise can be determined. In all cases a SNR larger than 10 has been measured. The results are summarized in Tab. 7.3.

- Fe55 signal
- threshold scan
- fit
- ENC

Bias current optimization

With the help of the bias currents, the chip performance is steered. The bias currents are controlled by 6 bit digital to analogue converters (DACs). The DAC settings thus control the amplifiers and comparators and steer the amplification and shaping times. At the same time they also dictate the power consumption of the pixel electronics. In [20] first results of a bias current optimization are presented which aim for a reduced power consumption. By measuring the currents drawn by the chip, the power consumption of this setting was measured to be $\approx 225 \text{ mW/cm}^2$. In test beam measurements several power settings have been tried. The most promising settings up to now offers a very good efficiency at $\approx 350 \text{ mW/cm}^2$

- Performance
- Power Consumption

Temperature dependence

In the detector environment the pixel detector faces temperatures from 0 to 70 °C. Possible performance differences due to varying temperatures need to be understood and compensated for. The main focus of these investigations is the pulse shaping and for the latest prototype the stability of the VCO and PLL circuits.

- MuPIX 2
- PLL & VCO




Figure 7.17: Temperature dependence of the position of the ^{55}Fe peak, from [23].

- Signal

Radiation tolerance

As the Mu3e experiment is performed at a muon beam-line the requirements on radiation hardness are not comparable to those at hadron colliders like the LHC. The radiation tests of HV-MAPS sensors done so far all indicated that there will be no radiation damage even at highest muon rates in Mu3e phase II.

Radiation-tolerance studies of HV-MAPS sensors in 180 nm and 350 nm technology are ongoing for other projects (e.g. ATLAS pixel R&D). Several test chips with similar pixel electronics as MUPIX have been irradiated at CERN up to doses up to 600 MRad (x-ray plus protons) and up to $10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ [24, 25]. The results are promising. A HV-CMOS sensor with the comparators capacitive coupled to the readout chip (CCPD) was irradiated at the Jožef Stefan Institute¹ to $10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ at a temperature of 5° and at -55 V bias. This CCPDv2 had a signal to noise ratio after neutron irradiation of 20 and a shaping time of 100 ns. The CCPDv2 irradiated to $10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ had an efficiency of about 95% instead of 99% for the unirradiated sample. Another CCPD detector works after irradiation to 600 MRad with x-rays plus protons showing increased noise to about 150e.

¹Institut "Jožef Stefan", Jamova cesta 39, 1000 Ljubljana, Slovenija

7.6.2 TESTBEAM RESULTS

In the years 2012 to 2015 fourteen testbeam campaigns were carried out in order to characterize the MUPIX prototypes, an overview is given in table 7.4. Below the test-beam campaigns are described and the core results summarized.

CERN SPS testbeam 2012

In August 2012 the MUPIX 2 chip was tested at the CERN SPS. The beam from the SPS was a 170 GeV pion beam, chosen for little multiple scattering at the telescope and the device under test (DUT). In two separate data taking periods the MUPIX 2 was mounted inside the TimePix telescope [26]. This silicon pixel telescope has four layers before and four layers after the DUT, providing very precise pointing resolution of $5 \mu\text{m}$. The MUPIX 2 chip was tested facing the beam and under an angle of 45° . For both beam periods a threshold scan was performed in order to derive the efficiency as a function of threshold, but no threshold calibration to equalize the gain of the pixels was performed. The measured resolution of $11.2 \mu\text{m}$ in x and $15.4 \mu\text{m}$ in y corresponds to the expected resolution given the telescope resolution and the pixel size.

DESY testbeam March 2013

The DESY testbeam in March 2013 was aiming at a precise study of the MUPIX 2. Unfortunately the integration of the MUPIX 2 readout together with the EUTelescope (ACONITE) suffered from a bad choice of the readout window. Since the intrinsic readout window of the MIMOSA26 chips in the EUTelescope is $2 \times 115 \mu\text{s}$ and the MUPIX 2 readout window was set to $1 \mu\text{s}$, the results of the efficiency study reflect the low time overlap of these two systems, see [27].

DESY testbeam June 2013

During the June 2013 testbeam at DESY both the MUPIX 2 and the MUPIX 3 prototype were tested with a 3 GeV electron beam. The hit integration time for the MUPIX 3 in the device under test position was set to cover the entire period between two triggers, but the random pixel enabling bug made a determination of the chip efficiency impossible. The readout of the time over threshold (ToT) for the MUPIX 3 was successfully used but



also here a random subset of pixels was contributing to the output signal. The high voltage and threshold dependence of the MUPIX 3 ToT was studied. The ToT decreased insignificantly with increasing high voltage. With increasing threshold the ToT decreased significantly from around $7\ \mu\text{s}$ at $0.83\ \text{V}$ to $\leq 0.2\ \mu\text{s}$ at $0.90\ \text{V}$ [27]. One MUPIX 3 thinned down to $90\ \mu\text{m}$ thickness was compared to another MUPIX 3 which was $600\ \mu\text{m}$ thick. HV and threshold scans were performed and no significant differences in ToT behavior could be found [27].

PSI testbeam October 2013

In this campaign the πM1 beamline at PSI was used providing a mixed proton, pion, muon and electron beam at a momentum range of $100\ \text{MeV}/c$ to $500\ \text{MeV}/c$, in the following results for particle momenta of $193.2\ \text{MeV}/c$ are discussed. Since no reference telescope system was available, all measurements were taken in the ToT operation mode with a scintillating tile time reference system. Thick ($600\ \mu\text{m}$) and thin ($80\ \mu\text{m}$) MUPIX 2 and thick ($600\ \mu\text{m}$) and thin ($90\ \mu\text{m}$) MUPIX 3 were tested. For the following measurements the beam electrons were selected with the help of the scintillating tile pulse height. For MUPIX 2 HV scans at a fixed threshold of $840\ \text{mV}$ were taken. The ToT for central pixels was shorter than the ToT at the edge pixels for both thick and thinned chip. The ToT for pixels of the thin chip is significantly longer than for the thick MUPIX 2. At $-90\ \text{V}$ the thin MUPIX 2 chip does behave differently, which could be attributed to either the thinning or the individual chip since no measurement could be taken before thinning [22,27]. For the thin and thick MUPIX 3 chip HV scans for the ToT did not show any significant trends. The ToT dependance on the location of the pixel could not be studied for the MUPIX 3, since an undefined subset of pixels was read out with the same ToT event. The thick MUPIX 3 chip had a significantly lower latency (scintillator time to start of ToT signal) of $160\ \text{ns}$ as compared to the thin MUPIX 3 chip $180\ \text{ns}$ [27].

DESY testbeam October 2013

In the October 2013 testbeam the MUPIX 4 chip prototype was successfully tested. Because of cross talk issues in the readout system developed in Heidelberg at the Physikalische Institut, the




Figure 7.18: MUPIX 4 sensor prototype inside the EUTElescope. The MUPIX 4 readout system developed by Ivan Perić was used here.

test system of Ivan Perić was used for these measurements. The MUPIX 4 was installed as device under test (DUT) in the EUTElescope ACONITE. To match the data sampling time of the MIMOSA26 chip of the EUTElescope, the MUPIX 4 hit collection time was set to twice the integration time of the MIMOSA chips before and after the trigger signal. The MUPIX integration time for the dedicated time resolution measurements had to be shorter.

In order to measure the time at which a particle hit was detected, the MUPIX 4 samples the bit pattern of an external Gray counter. This time stamp is then read out together with the pixel column and row address. In this test beam the MUPIX 4 timestamp was compared to the time of the trigger signal which the trigger logic unit (TLU) generates when pairs of finger scintillators before and after the telescope have been hit. Choosing a $100\ \text{MHz}$ Gray counter frequency for the MUPIX 4, the distribution of the time difference between MUPIX 4 and TLU has a sigma of $17\ \text{ns}$, see Figure 7.19. This performance made the MUPIX 4 chip the MAPS chip with the best time resolution.

The efficiency of the MUPIX 4 was measured with the EUTElescope as a reference with an electron beam of $3\ \text{GeV}$. Figure 7.20 shows the efficiency of the MUPIX 4 in the DUT position of the EUTElescope, on average it is around 99% . Further improvement in efficiency can be achieved by tuning the pixels individually, see below.




Figure 7.19: Time difference between MUPIX 4 pixel timestamp and TLU trigger signal time in units of 10 ns, from [27].




Figure 7.20: Efficiency for the MUPIX 4 chip. The efficiency for each pixel is shown in color scale. The average efficiency for each column and row are given at the axis, from [28].

DESY testbeam February 2014

The testbeam at DESY in February was used to further investigate the performance of the MUPIX 4 prototype. As in the previous testbeam the test-setup developed by Ivan Perić was used. The hit integration time for the MUPIX 4 readout setup to start from the last trigger and stop 200 μs after the current trigger signal.

The measurement of the efficiency of the MUPIX 4 in the device under test position of the EUTelescope was studied in detail. The high voltage (substrate voltage) was set either to -50 V or to -70 V , threshold scans in steps of 5 mV were taken and the MUPIX 4 was rotated to angles of




Figure 7.21: Efficiency for the MUPIX 4 chip for a high voltage of -50 V and a 3 GeV electron beam. Threshold scans for the MUPIX 4 chip were carried out unrotated as well as for angles of 22.5° and 45° , [28].

22.5° and 45° , see Figures 7.21 and 7.22. The efficiencies measured for the MUPIX 4 under an angle of 45° are up to 99.9%. The measurements under no angle and the one under 22.5° still exceed 99.5% for a high voltage of -70 V . The tuning of individual pixels was successfully applied to the MUPIX 4 for these measurements, as can be seen in the measurement of the per pixel efficiency shown in Figure 7.23. This efficiency map of the MUPIX 4 chip indicates a very high and homogeneous efficiency, pixels shown in white do not have an entry because of missing statistics. The spatial resolution of the MUPIX 4 is dominated by the pixel size of $92 \times 80\ \mu\text{m}^2$. Figure 7.24 shows the residual distribution of tracks determined with the EUTelescope as a reference system.

At the February 2014 testbeam at DESY the MUPIX Telescope (see section 7.5) was tested for the first time in a particle beam [22]. The MUPIX Telescope was equipped with four planes of the MUPIX 4 chips together with the readout system developed in Heidelberg, which still suffered from shortcomings concerning digital crosstalk. Nevertheless correlations between the column and row addresses of sensor hits in different planes gave a clear indication that the MUPIX telescope detects particle tracks.




Figure 7.22: Efficiency for the MUPIX 4 chip for a high voltage of -70 V and a 3 GeV electron beam. Threshold scans for the MUPIX 4 chip were carried out unrotated as well as for angles of 22.5° and 45° , [28].




Figure 7.24: Residual distribution of tracks through the MUPIX 4 prototype using the EU-Telescope as track reference, from [28].




Figure 7.23: Efficiency for the MUPIX 4 chip after tuning the pixels individually. The efficiency for each pixel is shown in color scale. The average efficiency for each column and row are given at the axis, from [28].

PSI testbeam June 2014

It was planned to test the new MUPIX 6 chip prototype at the PSI beam line πM1 with the mixed π^+ , μ^+ and positron beam at a momentum range of $100\text{ MeV}/c$ to $500\text{ MeV}/c$. The MUPIX 6 is a bug-fix of the MUPIX 4 chip, with functional readout of all row addresses and timestamps and a second amplification stage for the right 28 columns of the pixel matrix. Since the accelerator had to be repaired during the scheduled testbeam time, no beam was delivered to πM1 . Therefore the time

was used to test the new readout system hardware including LVDS data links mitigating the cross talk problem of the previously used flat ribbon cable [22]. The telescope setup was assembled from MUPIX 6 chips thinned to $250\text{ }\mu\text{m}$ on thinned MUPIX 4v1 boards with LVDS adapter cards. Calibration runs were taken using two different ^{90}Sr -sources of 3.7 MBq and 250 kBq . While there was still a cross talk issue in the setup, especially for the time stamps which were input without Gray encoding, some successful runs could be taken at low clock speed. Runs with a time stamp bin size of 640 ns and a readout cycle length of $12.7\text{ }\mu\text{s}$ were successfully taken and clear time correlations between the hits of different sensor planes could be shown. The scattering of the low energy 0.55 MeV electrons from the source prevented a successful analysis of spatial correlation in the four telescope planes.

PSI testbeam July 2014

A successful testbeam campaign was carried out at the πM1 beamline in July 2014 running at a mixed π^+ , μ^+ and positron beam with a momentum of $250\text{ MeV}/c$ [20]. On one optical breadboard a single MUPIX sensor plane and the MUPIX telescope with four sensor planes were mounted, both setups were complemented with two scintillating tiles for trigger and time information. The scintillating tiles have an area of $4 \times 4\text{ mm}^2$ facing the beam.

For all five sensor planes MUPIX 6 prototypes of $250\text{ }\mu\text{m}$ thickness on thinned MUPIX 4v1 boards



were used. The horizontal alignment of the MUPIX sensor plane could be adjusted by micrometer screws, but the vertical alignment both of the MUPIX sensor planes and of the scintillating tiles had to be done roughly by adjusting the post height. In this testbeam campaign the bias current settings for the single setup were adapted to the need for lower power consumption and heat production, cutting the three main bias currents to half (VN, VN2, VPComp=30). HV scans from -5 -80 V and threshold scans from 0.66 V to 0 V were performed for the single MUPIX chip setup. The online monitoring allowed amongst others to check the relative time of the scintillating tiles and the MUPIX 6 in the single setup, ensuring meaningful data sets for offline analysis. The time resolution of the MUPIX 6 chip was tested in the single setup with a time stamp bin size of 10 ns. For offline data analysis one hot pixel was removed, for the remaining pixels single hit events with trigger signals from both scintillating tiles were used. Figure 7.25 shows the correlation between the MUPIX 6 timestamp and the time from the scintillating tiles. The difference between MUPIX 6 time stamp and scintillator time was closely inspected, see Figure 7.26. A fit to the central part of the time difference distribution gives a σ of 15.48 ns, which can be considered as an upper limit of the MUPIX 6 time resolution. For one single pixel the comparator output (hitbus signal) was sampled together with the trigger scintillator signal with a 400 MHz clock equivalent to 2.5 ns bin size, the distribution of the time difference between hitbus time and trigger time can be seen in Figure 7.27. The data for this plot was accumulated over 10 runs at a threshold of 0.64 V and a high voltage of -60 V. The Gaussian fit to the peak of this distribution gives a σ of 15.1 ns. Since the timestamp to trigger distribution in Figure 7.26 and the hitbus to trigger distribution from a single pixel, Figure 7.27 both show a tail, it cannot be attributed to the pixel to pixel variation. The hitbus signal is sampled with the FPGA, so the latency from the particle passage (and detection by the trigger scintillators) and the time over threshold (ToT) are known for each event with a hitbus signal. A plot of the ToT against the latency, see Figure 7.28, shows the time walk of the MUPIX 6 and the potential improvement of the time resolution utilizing an on-chip time-walk correction.




Figure 7.25: Timestamp correlation between MUPIX 6 with bin size of 10 ns and scintillating tiles, from [20].




Figure 7.26: Difference between timestamp of the MUPIX 6 with bin size of 10 ns and scintillating tile time, from [20].




Figure 7.27: Difference between hitbus time of the MUPIX 6 with bin size of 2.5 ns and scintillating tile time, from [20].

PSI testbeam October 2014

The October testbeam at the π M1 beamline at PSI was used to continue testing the MUPIX 4 and MUPIX 6 prototypes with a π^+ , μ^+ and positron




Figure 7.28: Time over threshold (ToT) versus latency of the hitbus signal of MUPIX 6, from [20].

beam set to a momentum of 250 MeV/c [20]. Using a similar hardware as for the July PSI test-beam, both a single MUPIX setup and the MUPIX telescope were mounted onto an optical base plate together with a pair of scintillating tiles for trigger and timing purposes. The single setup was based on a MUPIX 6 chip glued on a 25 μm thick Kapton foil, the so called floating MUPIX. The MUPIX telescope had three reference planes of MUPIX 6 chips and the device under test (DUT) was either a floating thin MUPIX 4 or a floating MUPIX 6. The efficiency of the 50 μm thin MUPIX 4 chip was investigated with the help of the telescope in order to detect potential deviations from the thick sensors. For both the single setup and the telescope setup various bias current setting were studied in order to find the best compromise between detection efficiency and power consumption of the MUPIX. Furthermore the newly implemented accelerated readout was tested, operation was possible up to rates of 1 MHz of hits per layer, the rate used for most measurements was 333 kHz of hits per layer allowing for some safety margin. The maximum rate of this readout implementation is limited by the communication between CPU and FPGA. The efficiency of the MUPIX 6 in the device under test position of the MUPIX telescope was precisely determined. Figure 7.29 shows the efficiency of the MUPIX 6 as function of the high voltage applied to the pixel substrate. Clearly the efficiency goes up even between -75 V and -80 V , pointing at a more efficient charge collection and some charge amplification for the highest voltages applied. While the high voltage scan was taken at a moderate threshold of 0.65 V, Figure 7.30 shows the efficiency of the MUPIX 6 going up even around 0.7 V reaching 94 %. It can be concluded that the best working point for the MUPIX




Figure 7.29: MUPIX 6 efficiency versus high voltage, from [29].




Figure 7.30: MUPIX 6 efficiency versus threshold, from [29].

6 is at a voltage beyond -80 V and a threshold above 0.7 V, which was confirmed in later test-beam campaigns. In order to run at thresholds close to the baseline voltage of the MUPIX chip, it is important to tune all pixel cells to a similar level of noise, otherwise the noise of a fraction of the pixels would make the readout of the chip impossible. The programming of the tune digital to analog converters (TDACs) was successfully implemented and tested during this testbeam campaign and helped reaching high efficiencies at low noise in further test-beams.

DESY testbeam March 2015

The March testbeam campaign at DESY was used to continue the measurements with MUPIX 6 and to perform the first measurements with the MUPIX 7 prototype. The beam energy chosen for these tests was 5 GeV. The MUPIX 7 chip read out over a parallel bus with an external state machine as for the MUPIX 6 chip with similar firmware and software, which provides direct comparison between the two chip generations. The new fea-



tures of MUPIX 7 like the internal state machine for the readout, the internal clock filtering, fast PLL and the 1.25 Gbit/s serial readout were not used during this testbeam. The MUPIX telescope was equipped with three floating MUPIX 6 in the reference planes and one floating MUPIX 7 as device under test. The device under test in the EU-Telescope DATURA was either a single MUPIX 6 or MUPIX 7 plane.

The MUPIX telescope was run at a high voltage of -85 V and a time bin size of 80 ns , the DAC values can be found in [30]. The mechanical alignment was done with the help of precision screws resulting in a mechanical misalignment of around $100\text{ }\mu\text{m}$, Figure 7.31. The software alignment done offline resulted in mean residuals of less than $6\text{ }\mu\text{m}$. The efficiency as a function of threshold for the MUPIX 7 in parallel readout mode is shown in Figure 7.32. The maximum efficiency measured for this scan was 99.31% at a relatively high noise rate of 668 Hz/pixel . A good working point for this chip would be a threshold between 0.74 V and 0.75 V , for which the efficiency is above 98% and the noise is below 200 Hz/pixel . Better performance can be reached by carefully tuning the individual pixels with the help of the TDACs, see below.

In the EU-Telescope a single MUPIX 7 was used as device under test. It was read-out with the readout scheme developed for MUPIX 4 and MUPIX 6 using the external state machine on the FPGA, external timestamp counters and parallel data transmission via low voltage differential signaling (LVDS). For some studies of the MUPIX 7 each pixel baseline was tuned individually, aiming at an equalized noise rate for one specific threshold. Threshold scans were performed for detailed studies of the efficiency for high voltages of -85 V , -90 V . The bias current settings were chosen either to allow for maximum chip performance or to save power which is important for a very light weight overall design of the on detector services. For 0.772 V threshold and a high voltage of -85 V an efficiency of 97% was determined, for 0.775 V threshold and a high voltage of -90 V the efficiency was 98.6% . In the power saving settings (225 mW/cm^2), 0.762 V threshold and a high voltage of -90 V the efficiency of 93.5% was measured, see Figure 7.33. The high pointing resolution of the MIMOSA26 reference planes in the EU-Telescope of around $5\text{ }\mu\text{m}$ was used to study




Figure 7.31: Row row correlation after mechanical alignment, from [30].




Figure 7.32: Efficiency for threshold scan with MUPIX 7 in parallel readout, from [30].




Figure 7.33: Efficiency for a MUPIX 7 prototype at a high voltage of -90 V and power saving settings, the measurement was performed with the EU-Telescope DATURA as reference and after tuning all pixel baselines individually.

sub-pixel effects, see Figure 7.34. No sub-pixel structure was observed for the MUPIX 7.




Figure 7.34: Sub-pixel efficiency for a MUPIX 7 prototype at a high voltage of -90 V and 0.777 V threshold, the measurement was performed with the EUTElescope DATURA as reference and after tuning all pixel baselines individually.

Mainz testbeam spring 2015

In spring 2015 the MUPIX 6/7 telescope was set up in Mainz in order to perform measurements behind the tagger magnet in the A2 photon scattering hall. Here ≈ 1 GeV electrons from MAMI are available for parasitic running after they radiate the photons used in the crystal ball experiment. The MUPIX 6 and MUPIX 7 sensors were the ones used at DESY before and had a thickness of $250 \mu\text{m}$, the setup and especially the bias DAC values can be found in [30]. The baseline for each pixel was tuned individually in order to achieve a homogenous noise distribution at the target threshold. A high voltage of -85 V was applied to the sensors and the time bin size was set to 80 ns. As for the DESY testbeam data analysis a χ^2 -cut of 10, a time window of 240 ns but a search radius of $320 \mu\text{m}$ was chosen for the track reconstruction [30]. For a threshold of 0.76 V an efficiency of 98.87% and a noise rate of 4 Hz/pixel was found. The relatively low noise rate was a result of performing the baseline tune for the individual pixels at the same threshold of 0.76 V. Figure 7.35 shows the MUPIX 7 efficiency for a threshold scan, revealing high efficiencies of up to 99.87% at the cost of a very high noise rate of 1960 Hz/pixel.




Figure 7.35: Threshold scan for a MUPIX 7 prototype at a high voltage of -85 V, the measurement was performed after tuning all pixel baselines individually in order to achieve a low and homogenous noise per pixel at 0.76 V, from [30].

CERN SPS testbeam July 2015

At the CERN SPS testbeam in July 2015 the MUPIX 7 chip was tested in a 180 GeV pion beam. The MUPIX 7 chip was for the first time in a particle beam operated with internal readout state machine, built-in clock circuitry and high speed serial data output at 1.25 Gbit/s. The MUPIX telescope for this campaign was assembled from four planes of MUPIX 7 chips on MUPIX 7 PCBs each connected to the same FPGA via a SCSI cable. Fast clock and data signals are running over the SCSI cable as well as the slow control communication all of which are implemented in LVDS standard. Threshold scans were performed both with bias current settings at high power consumption (935 mW/cm²) and for power saving bias settings (225 mW/cm²), see Figure 7.36 and 7.37. These results prove that the MUPIX 7 HV-MAPS full-system-on-chip runs in continuous readout mode with an efficiency around 99% , marking one of the major milestones of the Mu3e R&D work.

PSI testbeam October 2015

As for previous measurements at the π M1 beamline at PSI a π^+ , μ^+ and positron beam set to a momentum of 250 MeV/c was used. The hardware setup comprised a MUPIX 7 telescope with four planes and a single MUPIX 7 setup. A pair of plastic scintillators with PMT readout were used as time and trigger reference. During this campaign the efficiency of the MUPIX




Figure 7.36: Threshold scan for a MUPIX 7 prototype in *high power* (935 mW/cm^2) bias current configuration at a high voltage of -85 V , the measurement was performed after tuning.




Figure 7.37: Threshold scan for a MUPIX 7 prototype in *low power* (225 mW/cm^2) bias current configuration at a high voltage of -85 V , the measurement was performed after tuning.

7 for four different bias current configurations was examined, ranging from 225 to 935 mW/cm^2 . Beside the so called *high power* 935 mW/cm^2 and *low power* 225 mW/cm^2 settings, a *medium power* 400 mW/cm^2 and a *low-medium power* 300 mW/cm^2 setting were chosen in the attempt to find a good compromise between overall chip performance and power consumption, see 7.38. The performance of the MUPIX 7 chip was best in the *low-medium power* settings, showing efficiencies of above 99% while maintaining a noise rate of around 10 Hz/pixel .

DESY testbeam October 2015

An electron beam of 2 GeV to 6 GeV was used for the measurements. Most of the threshold scans were taken at a beam energy of 4 GeV . At the DESY testbeam area TB22 one EUTelescope with a single MUPIX 7 plane was used, downstream a MUPIX 7 telescope with four working MUPIX 7 planes and another MUPIX 7 telescope with three




Figure 7.38: Threshold scan for a MUPIX 7 prototype in *low power* 225 mW/cm^2 , *low-medium power* 300 mW/cm^2 , *medium power* 400 mW/cm^2 and *high power* 935 mW/cm^2 bias current configuration at a high voltage of -85 V , the measurement was performed after tuning.

working planes were installed. The last telescope was run all the time in order to commission a second MUPIX 7 telescope. It was tried unsuccessfully to integrate the single MUPIX 7 plane in the Eudaq system. The operation of the middle MUPIX telescope was very successful, threshold scans for the second MUPIX 7 plane were performed with the dual stage amplification, the single amplification stage (MUPIX 4 mode). In addition scan were taken for the sensor rotated by 0° , 15° , 30° , 45° . The high voltage was set to -85 V , the bias currents were set to intermediate power settings 300 mW/cm^2 and the pixels were tuned for homogeneous noise distribution for all sensor planes.

For the threshold scans under angles of 0° to 45° first software alignment for each angle was performed, yielding residuals in x and y -direction of less than $10 \mu\text{m}$, for many runs even below $2 \mu\text{m}$. After the selection of successful runs, the dependency of the MUPIX 7 efficiency on χ^2 and the angle at a threshold of 0.74 V was plotted, see Figure 7.39. Efficiencies of above 99% at a noise rate of around 2 Hz/pixel have been measured for MUPIX 7.

The test-beam results obtained over the last years are summarized in table 7.4. These results reflect the performance not only of the MUPIX prototype chips but also the performance of the over-all system including electronic hardware, FPGA firmware and software. Even though many detailed measurements with the MUPIX 7 are either ongoing or scheduled for the first half




Figure 7.39: Angular scan for a MUPIX 7 prototype from 0° to 45° at a high voltage of -85 V, the measurement was performed after equilibrating the pixels.

of 2016, the current results encourage Mu3e to aim for the fabrication of a large scale prototype in spring 2016. This large prototype will allow studying the scalability of the MUPIX chip and to start the HV-MAPS module integration tests.



Beam	Period	Particles	Energy [GeV]	MuPIX type	Eff. [%]	Time res. [ns]	Noise @ 99 % eff. [Hz/pixel]
SPS	Aug 2012	pions	170	MuPIX 2	≥ 86	-	-
DESY	Mar 2013	electrons	2 to 5	MuPIX 2		-	-
DESY	Jun 2013	electrons	3 to 5	MuPIX 2, MuPIX 3		-	-
PSI	Sep 2013	pions	0.16 to 0.45	MuPIX 2, MuPIX 3		-	-
DESY	Oct 2013	electrons	3 to 5	MuPIX 4	≥ 99	≤ 17	-
DESY	Feb 2014	electrons	5	MuPIX 4		-	-
PSI	Jun 2014	^{90}Sr	0.5510^{-3}	MuPIX 6	-	-	-
PSI	Jul 2014	pions	0.25	MuPIX 4, MuPIX 6	≥ 86	15.5	-
PSI	Oct 2014	pions	0.25	MuPIX 4, MuPIX 6	≥ 94	-	-
DESY	Mar 2015	electrons	5	MuPIX 6, MuPIX 7	99.31		668 (≤ 5 at 98.5% eff.)
Mainz	Q2 2015	electrons	1	MuPIX 6, MuPIX 7	98.9		4
SPS	Jul 2015	pions, p	180	MuPIX 7	≥ 99	?	11
PSI	Oct 2015	pions	0.25?	MuPIX 7	≥ 99.5	?	3
DESY	Oct 2015	electrons	2-6	MuPIX 7	≥ 99	?	≤ 2

Table 7.4: Test beam overview



PIXEL MECHANICS AND SUPPLY




Figure 8.2: Tool for layer 3 segment assembly.

The pixel detector mechanics has been optimized for very low material budget in the active detector region. Additional requirements are mechanical stability, resistance to temperatures over a wide range and a modular design for ease of assembly and repair. The gaseous helium cooling distribution to the HV-MAPS chips is part of the mechanical design of the pixel detector, see chapter 19.

It is proposed to build the frame for the silicon pixel detector from thin Kapton foil. The HV-MAPS chips are glued and bonded on a flex-print and then mounted onto the Kapton frame. As the digital readout circuits of the pixel chips create an approximately 0.5 mm wide dead area, there is a 1 mm overlap to the adjacent HV-MAPS chip on the next flex-print. The Kapton foil used for the mechanical frame construction is 25 μm thin. It gains in mechanical stability as it is folded around a prism-shaped template and glued to plastic end-pieces. This structure, composed of HV-MAPS chips, flex-print, Kapton frame and end-pieces is regarded as a pixel module. The outer pixel modules gain extra mechanical stability from V-




Figure 8.3: Tool for layer 3 segment assembly.

shaped folds in the support structure, which are positioned under the middle of each row of HV-MAPS chips.

The pixel detector layers are composed of multiple pixel modules, with each module. These layers have four different sizes and prism shapes, see Figure 8.4.

Layers 1 and 2 are each build from two modules, layer 3 and 4 are build from six and seven modules respectively. The inner double layers have 12 cm active length. Layer 1 has 8 and layer 2 10 sides of 19 mm width. Each inner layer is assembled from two half-modules. As a consequence the plastic end-pieces are half moon shaped. All four half modules of the inner detector layers are mounted to two thin rim wheels. A mechanical prototype for the inner double layer has been constructed from 25 μm Kapton foil both for the frame and the flex-print layer, while the pixel chips have been simulated with 100 μm thick glass plates. Glass




Figure 8.1: Screenshot of the Mu3e detector mechanics 3-D model.

of 100 μm thickness is of comparable flexibility as thinned silicon, which is shown in Figure 8.5. The resulting mechanical unit is surprisingly sturdy and fully self supporting, as can be seen in Figure 8.6.

Layer 3 and layer 4 have a three times larger active length of 36 cm. The sides of these outer layers are 19 mm wide. The layers 3 and 4 have 24 and 28 sides. An outer double layer module combines four sides, so layer 3 consists of 6 and layer 4 consists of 7 modules, see Figure 8.7.

The modules of the pixel detector layers are mounted onto pairs of end-wheels, forming detector stations. In Mu3e phase I there are three detector stations, one in the center around the target, one recurl stations upstream and one recurl station downstream of the central station. The central station consists of the target, the pixel layers 1 and 2 forming the vertex detector, the fiber tracker and the pixel layers 3 and 4, forming the central outer pixel detector. Assembly of each the central station is done in a special mounting frame, combining the inner two layers with the modules of the outer two layers on large rim wheels (Figure 8.8). Each of the two recurl stations has a metal beam pipe in the center onto which the tile-detector and the pixel layers 3 and 4 are mounted.




Figure 8.4: Mechanics of the central pixel detector

Figure 8.1 shows the status of the mechanical design of the Mu3e detector.

8.1 Mechanical design of the four pixel layers

8.1.1 LAYER 1

The innermost layer (layer 1) of the Mu3e pixel detector is built around the target and, together with layer 2, forms the vertex tracker of the Mu3e experiment. It has prism shape with an active length of 12 cm and eight sides of 19 mm width each, see Figure 8.9. It is composed of two modules with four sides each, see page 246 in the appendix. The



Figure 8.5: A 50 μm thin silicon wafer.




Figure 8.8: Mounting tool for the central pixel detector.



Figure 8.6: Mechanical prototype of the inner pixel layers. Thin glass plates replace the silicon chips.




Figure 8.7: Segmentation of the central outer layers.




Figure 8.9: Mechanical design of the inner pixel layer 1.




Figure 8.10: Mechanical design of the inner pixel layer 2.




Figure 8.11: Mechanical design of the outer pixel layer 3.

layer 1 module has one end-piece on each end and a Kapton foil bridging the active area. Each end-piece is milled from PEI (polyetherimide), which can be used at temperatures of up to 170 °C. The Kapton foil in the active area is 25 μm thin and has cut-outs for the bond-wires between the flex-print and the MuPix chips. It also has holes in the four corners for the relative alignment between the foil and the end-pieces during the fabrication process. In addition it is foreseen to cutout part of the foil under the MuPix sensors in order to reduce the amount of material. Each side of the layer 1 module is equipped with six MuPix chips of $2 \times 2 \text{ cm}^2$. Ideally three MuPix chips can be cut in one strip from the production wafer, making the relative alignment much easier. This depends strongly on the MuPix yield on the wafers, because three good chips in a row are required. The upstream and downstream ends of the layer 1 pixel modules are read out separately, using flex-prints with control signals, clock, power and high-voltage shared by the three MuPix chips on each end, see section 8.2.

8.1.2 LAYER 2

The second layer (layer 2) of the pixel detector is constructed very similarly to layer 1, see section 8.1.1. It also has an active length of 12 cm but its prism shape has ten sides, see Figure 8.10 and page 247 in the appendix. Layer 2 is also made from two half-modules, each having five sides with six MuPix chips per side. The modules are composed from two PEI end-pieces and a Kapton foil acting as a mechanical frame in the active area. The flex-prints with the glued and bonded MuPix chips in one piece is exactly the same as for layer 1.

In contrast to the module for layer 1, the layer 2 modules have inlets for the gaseous helium used to cool the MuPix chips. The helium gas is injected into the gap between the layer 1 and the layer 2 by two gas inlets per side, so ten inlets per module end-piece. The hot helium gas is extracted through identical outlets on the opposite side, see chapter 19.

8.1.3 LAYER 3

The third pixel detector layer, which is a prism shape with a length of 36 cm and twenty-four sides of 19 mm width each, forms together with layer 4 the outer pixel tracker. It is used in the central part around the target, vertex layers (layer 1 and 2) and Fiber Detector and in the recurl stations up- and downstream around the Tile Detector. There are six layer 3 pixel detector modules in each detector station. A layer 3 module combines four sides of 19 mm width and 36 cm length, see Figure 8.11 and page 248 in the appendix. The modules are built from PEI end-pieces on both sides and a 25 μm thin Kapton foil acting as a mechanical frame in the central part. The Kapton foil has a v-shaped fold along each side underneath the flex-prints and MuPix sensors, ensuring mechanical stability and offering a channel for the helium coolant. The Kapton frame also has laser cut holes to make space for the bond-wires between flex-print and chips. The end-pieces have a v-shaped groove in the middle of each side of the detector matching the fold in the Kapton foil frame. Additionally the end-pieces have inlets for the cold helium gas, which flows inside the v-shaped fold and in opposite direction also on both




Figure 8.12: Mechanical design of the outer pixel layer 4.

sides of the fold. Each layer 3 module is equipped with eighteen MuPix chips of $2 \times 2 \text{ cm}^2$ per side, so 72 chips in total. The chips are ideally cut in groups of three from the wafer. On this flex-print power and read-out of the nine chips in the upstream half and the ones in the downstream half are separate. During the fabrication process the four flex-prints are glued onto the Kapton frame and to the end-pieces one by one, among others closing up the v-shaped cooling channels.

8.1.4 LAYER 4

Layer 4 of the pixel detector is very similar to layer 3 and has a prism shape with 36 cm and twenty-eight sides of 19 mm width each. Each layer 4 pixel detector station has seven modules with four sides each, see Figure 8.12 and page 247 in the appendix. The layer 4 modules are quite similar to those of layer 3 and the flex-prints for the module 4 construction are identical.

8.2 Flex-print assembly

The MuPix chips must be powered, properly grounded and require all the connections for communication to and from the chip. The flex-print assembly takes care of this. It consists of multiple layers, each for different functions:

- **Power layer:** Distributes the low-voltages to the chip. Made of aluminum film of $25 \mu\text{m}$ thickness in order to keep the voltage drop across lines within an acceptable level of less than 20 mV. Two power grids are needed

(V_{DD} and V_{ss}) and carried out as point-to-point connections from power connector to chip, avoiding extra material from vias.

- **Signal layer:** Provides the control signals to the MuPix chips and hosts the readout lines. All signal connections will be implemented as impedance-controlled differential pairs ($Z_0 \approx 50 \Omega$ and $Z_{\text{diff}} \approx 100 \Omega$). This layer will also distribute the HV. Currently it is considered to use either commercial technology (HiCoFlexTM) which allows to produce a copper layer on a polyimide substrate with a thickness of a few μm , or another technology (LED Technologies of Ukraine) based on etched aluminum on polyimide substrate.
- **Ground layer:** One layer of aluminum film, similar to the power layer but without a trace structure and separate for the upstream and the downstream end of the detector module.

The three layers will be laminated together using epoxy glue (Stycast 1266), which can be applied in thin layers grace to its low viscosity. To apply the glue, two options will be evaluated. The first option is using a robotic gantry with a dispenser head attached. The second option is a stamping technique where glue from a reservoir wets a rubber stamp and the pattern will be imprinted on the target substrate. Both options have been used successfully in multiple particle physics experiments before. The method with the best results in terms of minimized glue amount and uniformity of the results will be chosen, see also section 9. Using this approach, the material load per single pixel layer can be achieved to match the breakdown as given in Tab. 8.1.

8.2.1 DESIGN CONSIDERATIONS

A multi-layer design offers the following advantages over a single layer design: It is possible to use a monolithic plane with a low resistance for ground. Moreover, the ground plane acts as shielding between the sensors and the signals on the other layers of the flex-print. Furthermore, it is possible to divide this plane into ground and power distribution.

A dedicated signal layer can be optimized regarding trace width and separation to match the impedance ($Z_0 \approx 50 \Omega$ for single traces and $Z_{\text{diff}} \approx 100 \Omega$ for differential pairs) and also to distribute



the high voltage. The impedance is controlled by the ground and power layers such that there is always a power/ground plane above and/or below a signal trace. Here, no large currents are expected to flow, so the layer can be thin. Therefore, this layer can be made of copper as it is not the most critical component in the material budget (see Table 8.1). To accommodate bus connections, two signal layers are required, which however does not have a large impact on the material budget.

A two-layered signal HiCoFlexTM solution with additional in-house produced aluminium layer for power and ground distribution has been studied.

Similarly a design based on the technology offered by LED Technologies of Ukraine has been studied. This technology makes up to two layers with etched aluminum traces on polyimide substrate available. It provides also inter-connectivity between the two layers. Using these possibilities a layout with one layer for signal traces and one layer for the signal fan-out of the chips, power and ground has been developed.

Current design specifications for the differential signal lines can be found in Table 8.2. The specifications have been found using the layer stacks for the HiCoFlex and LTU solutions as shown in Figures 8.13 and 8.14.

The power layer will be optimized to keep the power dissipation along the flex-print as low as possible, which is yet a challenge for the outer layers where the power for 9 sensors has to be supplied with point-to-point connections. As an advantage, point-to-point connections allow to regu-

COMPONENT	THICKNESS [μm]	X/X ₀ [%]
Kapton frame	25	0.009
Kapton flex-print	20	0.007
Aluminum power lines (99% coverage)	2 × 25	0.056
Copper traces (33% coverage)	4	0.009
HV-MAPS	50	0.054
Adhesive	2 × 2	0.001
Full detector layer	199	0.136

Table 8.1: The pixel detector layer radiation length is dominated by the HV-MAPS chips and the aluminum in the power and ground layers.

TRACE PARAMETER	DIMENSION	DIMENSION
	[μm] HiCoFlex TM	[μm] LTU
width (top layer)	40	65
separation (top layer)	80	130
width (bottom layer)	20	65
separation (bottom layer)	40	130

Table 8.2: Differential pair parameters in the signal layer to match $Z_0 \approx 50 \Omega$ and $Z_{\text{diff}} \approx 100 \Omega$ for both vendors (HighTec and LTU).

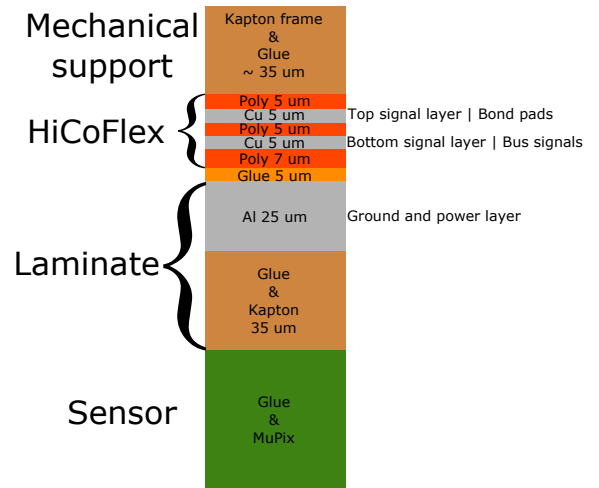


Figure 8.13: Stack chosen for the HiCoFlex solution with additional aluminium layer for ground and power distribution.

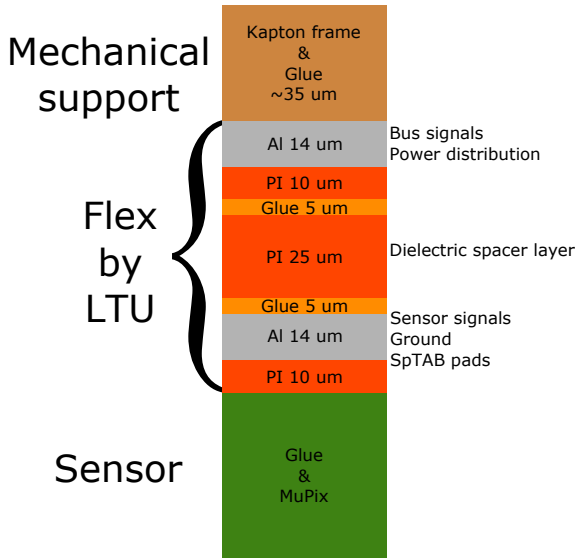


Figure 8.14: Stack chosen for the LED Technologies of Ukraine (LTU) 2-layer solution.

late the power for each sensor to compensate the voltage drop along the line. A power distribution by a bus line would have the advantage of a smaller resistance, but the voltage drop along the line would exceed our targeted 20 mV. This solution would require on-chip power regulators.

For the HiCoFlex solution, all signals, power connections and ground will be connected to the sensors by wire bonding. The bond pads for the signals on the copper layers will be on the top signal layer, while there will be holes in the HiCoFlex to connect ground and power on the aluminium layer below. This solution requires a considerable amount of vias to minimize the material. Currently under investigation is the addition of a rigid-flex PCB, with one end glued to the flex-print at the module end-piece and the flexible part acting as the cable going to the front-end PCB. Clearly this would ease the mechanical and electrical integration, but must be included in the overall mechanical design. The connections to the flex-rigid PCB will be done using solder connections through holes. Similar to the pads for bonding to the MuPix the solder holes for the signals will be on the top signal layer, while power and ground will be soldered directly on the aluminium layer.

The current design using the HiCoFlexTM technology, see Figures 8.15 and 8.16, foresees a daisy-chaining of common signals for all sensors, which include the reference clock and the synchronous




Figure 8.15: Composite view of the current flex-print design for a vertex module showing the power traces on the top layer (blue) and the solder connections on the left. The ground layer is shown in purple. It is foreseen that 3 MuPix sensors can be connected.




Figure 8.16: The inner layer of the flex-print design showing the signal traces and the high voltage distribution for 3 sensors.

reset. This is required to keep the number of signals per module as small as possible in order to facilitate the operation of the module.

For the 2-layer LTU solution, connection of the flex and the sensors will be done using the SpTAB technique (cite!). Here, the aluminium traces of the flex are connected to the pads by applying pressure. The reliability of this process has to be determined to estimate the yield of fully functional modules.

8.2.2 FLEX-PRINT STUDIES

The feasibility of high speed data transmission with aluminum-Kapton flex-prints has been studied in [31] and [32]. Starting from a foil laminate with 12 μm or 25 μm aluminum and 25 μm Kapton, with 25 μm glue in-between, differential traces with a width and separation of 120 μm have been produced by means of a laser platform (see Figure 8.20). The 12 μm aluminum foil would be preferable regarding radiation length, however long term tests have shown that bare traces with 12 μm failed after being exposed to the lab environment for a longer time. It has also been found



DATA RATE [Mbps]	RUN TIME [h]	CHANNELS	BIT ERRORS OBSERVED	BIT ERROR RATE
1000	29	7	0	$< 7.1 \cdot 10^{-15}$
1600	512	7	0	$\leq 1.8 \cdot 10^{-16}$
2500	19	5	0	$\leq 4.3 \cdot 10^{-15}$
3200	398	5	36	$1.6 \cdot 10^{-16}$

Table 8.3: Bit error rates measured for a 20 cm aluminum-Kapton flex-print prototype. When 0 errors are observed, the upper limit at a 95 % confidence level is given.




Figure 8.19: Composite view of the 2-layer aluminum Flex (LTU solution).



Figure 8.20: Aluminum-Kapton flex-print prototype produced in-house. The flex-print is 10 cm long and 1.8 cm wide. It comprises 17 differential signal pairs.




Figure 8.17: Top layer of the 2-layer aluminum Flex (LTU solution).




Figure 8.18: Bottom layer of the 2-layer aluminum Flex (LTU solution).

that the traces do not survive several connection cycles using a ZIF¹-connector.

To control impedance, another layer of aluminum-Kapton laminate has been glued to the aluminum side of the flex-print. The impedance of a differential pair can be tested with a time domain reflectometer. With structures as mentioned before the single line impedance has been measured to be in the range of (45-55) Ω [32] which matches the goal impedance of 50 Ω within a 10 % tolerance.

To test the quality of data transmission, bit error rate tests have been performed using a Stratix V FPGA development board and a custom designed PCB to connect the flex prints. For flex-prints with a length of 20 cm different data rates have been tested (see Table 8.3). The flex-prints transmitted data up to rates of 2500 Mbps without any errors in our tests, exceeding even the M_U3e design data rate of 1250 Mbps by a factor of two.

The flex-prints produced in-house so far show good performance, however the trace width is limited by the current laser setup to be 120 μm or

¹Zero insertion force



larger. In a layout where the flex-print is sandwiched between the MuPix and the Kapton frame, smaller sizes are required to match the impedance. Therefore, a full in-house design solution is not

feasible at the moment and the hybrid solution with a signal layer produced in a commercial technology is the current choice.



PIXEL MODULE PRODUCTION AND TESTING

Pixel modules are an assembly made of MuPix chips attached to a circuit providing the electrical connections to the data acquisition further downstream. The latter consists of multiple layers, where the signal layer is manufactured by an external company and the powering/ground-layers will be made in house. The lamination of those layers together, the attachment of the MuPix chips, the soldering and the wire-bonding are shown in Fig. 9.1, which in addition illustrates the dependencies of all steps and the testing (“quality control points”) performed at each stage of the process.

9.1 Manufacturing steps

In manufacturing, the parts get assembled in order to make a working module, ready for mounting on the mechanical structure. The general workflow is outlined in Fig. 9.2, crucial aspects which are taken care during manufacturing are:

- **Flex-print assembly:** The flex-prints for the signal layer get delivered from the manufacturer in singulated units. After an inspection they will be ready to use. The layers for powering and ground are made in-house. For this, aluminium films will be laser-cut to size on a temporary handling tape and then laminated to the signal flex-print on both sides using a thin film of epoxy glue.

If we find a provider for a three layer flex-print made entirely with aluminium we will use this source.

- **PCB assembly:** The PCB makes the connection between the flex cable and the flex-print assembly. The cable is an integral part of the PCB, which is made of a flex-rigid-technology, i.e. the flexible layers will be sandwiched between the rigid layers of the PCB. First, some electronic components like resistors and repeaters need to be mounted, all in SMD packaging. This will be done using industry-standard SMD reflow soldering. A mandatory electrical test will ensure that the PCBs are in working condition and that the flexible cable attached to the PCB shows the required electrical properties. In the second step, the electrical connection between the flex-print and the PCB will be made via solder-holes in the flex-print. The aluminium layers will be connected using ultrasonic soldering¹.

- **Sensor chips:** The sensors will be delivered as wafers from the chip foundry. A thorough inspection will take place to single out the working items and bad chips will be marked (see section 9.3.2). The chips will then get diced out at an external company. After a visual inspection (surface sanity) and an electrical test (similar to the wafer test) they are

¹Soldering is restricted to peripheral parts outside the active sensor volume in order to avoid high- Z materials in the tracking volume.

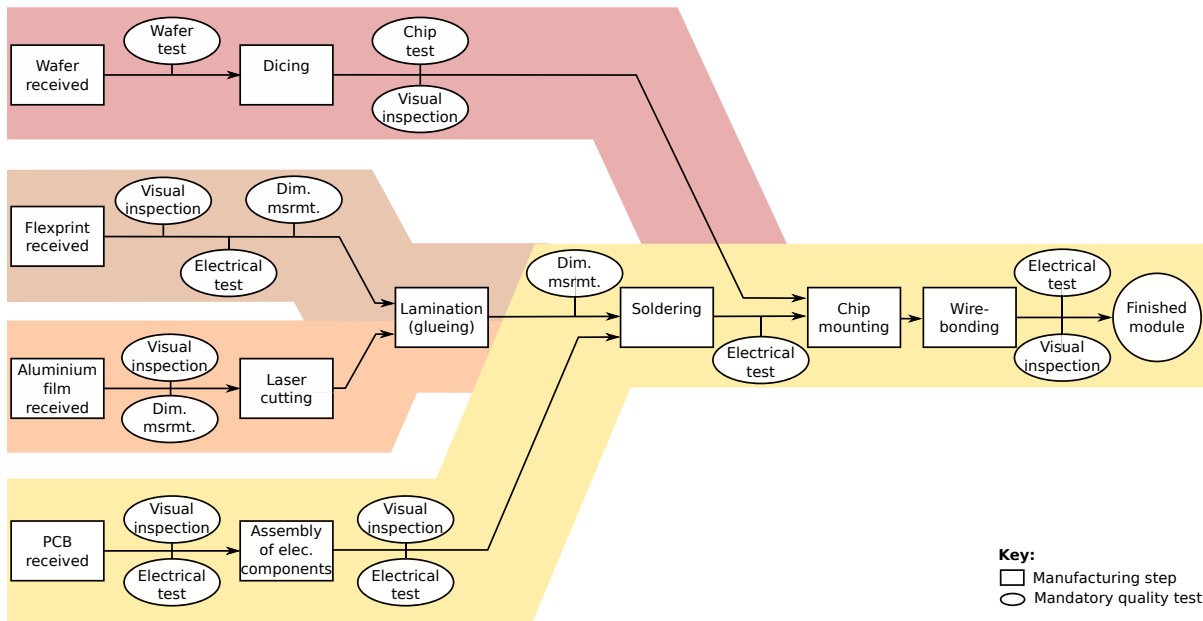


Figure 9.1: Production workflow and quality control points. Note: Dim. msrmt. stands for dimension measurement.

ready for mounting. The chips will be glued to the flex-print assembly using a thin layer of epoxy glue.

- Sensor mounting and finalisation of modules:** After curing, the electrical connections between the chip and the flex-print will be made using industry-standard aluminium wedge bonding (wire thickness: $25\ \mu\text{m}$, wedge with long bond-flat of about $75\ \mu\text{m}$). Finally, the modules are fully functional and can be tested thoroughly before getting mounted on the mechanical support structure. A close-up picture of the full module will be taken for the records. This will allow to measure the relative alignment of the sensors to its neighbours from visible features on the chips.

Manufacturing takes place in appropriate environments, i.e. wafers will be handled in clean-room environments, any semiconductor part will be handled with proper ESD protection. Should transportation be required (i.e. shipping wafers to the dicing company), proper packaging will be used to safeguard the materials.

The tools needed for manufacturing will be developed over the course of the next few months. They will be designed in-house to facilitate the needed production steps. For handling flat module components and putting them together, tools inspired by proven designs of the CMS pixel group at Paul Scherrer Institut will be evaluated. As an example, Fig. 9.3 shows a versatile tool for lamination and glue application purposes, shown in a configuration for precision placing of a part on top of another part. A precision linear guide controls the vertical movement. It is balanced with a counter-weight (not visible) to ease operation and to control the force applied when left free. Operation is entirely manual and parts are held by means of vacuum. Alignment is controlled via adjustable micrometer screw stops. We envision to fabricate such tools in the configuration needed for making MuPix modules in the sizes needed. For pick-and-place tasks we envision the use of an existing four-axis robot used for the precise assembly of PCBs. To control glue application, various techniques are being evaluated and the best will be chosen. Among those are precision dosing of glue using syringes in a glue-robot and stamp techniques like CMS uses.

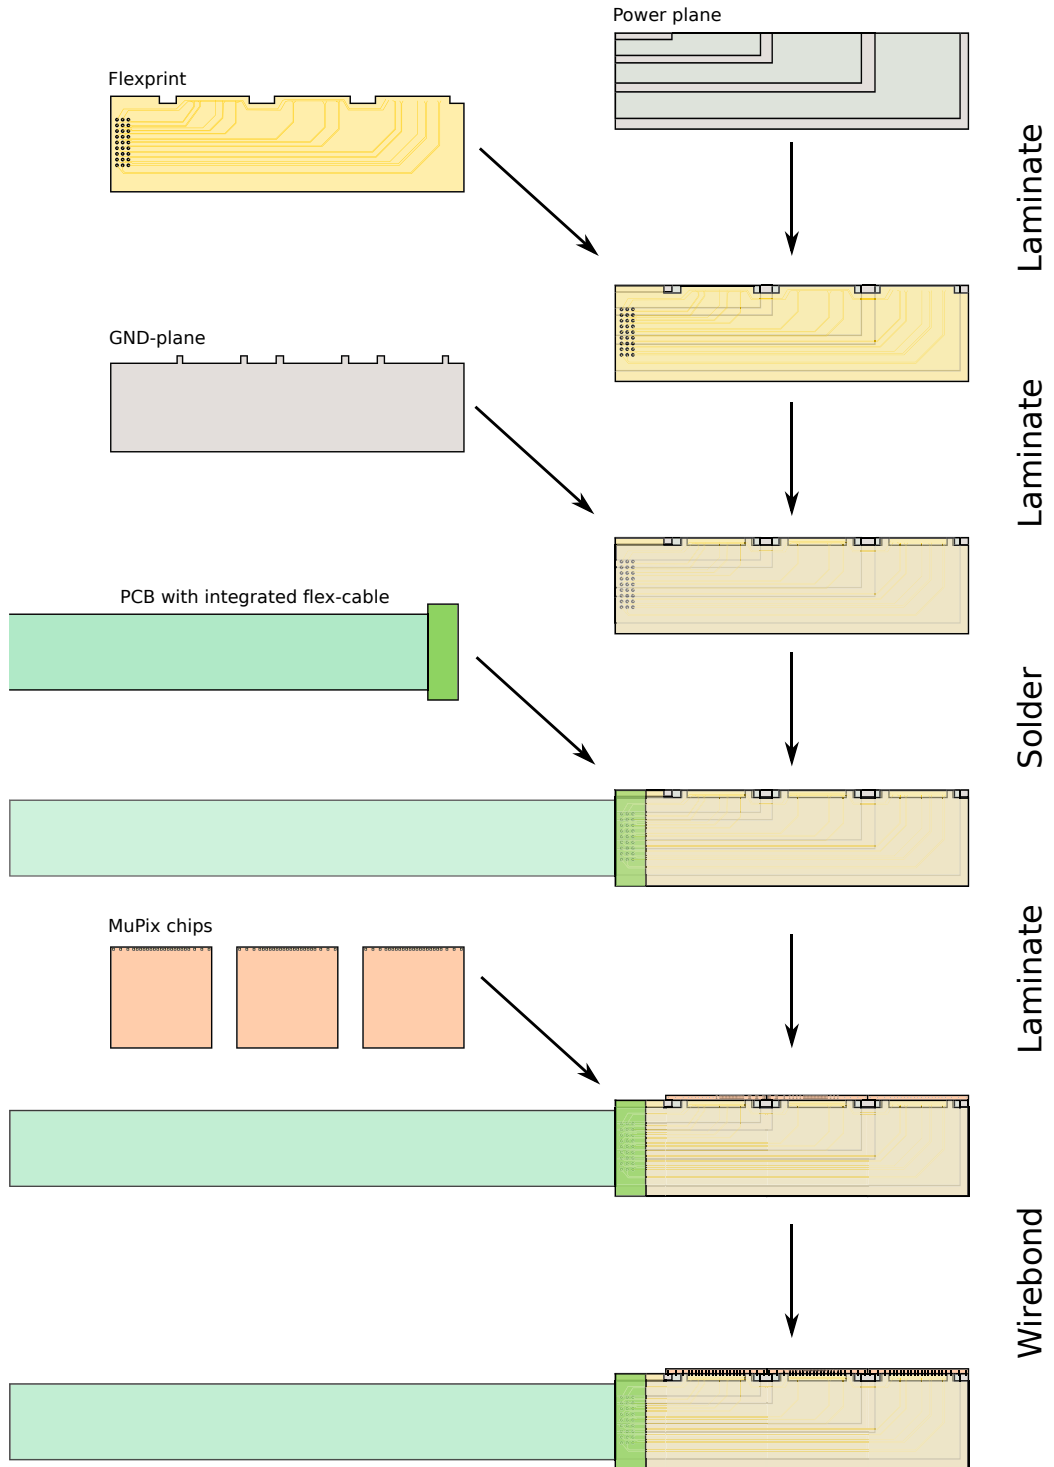


Figure 9.2: Main production steps for manufacturing a pixel module.




Figure 9.3: Versatile tool for lamination, gluing and placing of parts as needed during manufacturing steps. In this configuration, the bottom part holds a module via vacuum, aligned by micrometer screw controlled stops. The top part is movable vertically by precision linear guides. Here it is used for holding a part before placing it precisely (image shows state after that). Size indication: module dimensions approx. 22 mm \times 66 mm. Image used by kind permission of CMS pixel group at Paul Scherrer Institut.



9.2 Part tracking

Every part used for manufacturing modules is identified to allow tracking.

Identification. Unique numbers are assigned to parts at different levels.

Part no. Every part receives a number, be it a commodity product (screws, tape, etc.) or a custom manufactured part (flex-print, chip dice, etc.). For each part, specifications and test procedures are required.

Lot no. Every delivery of a part gets assigned a lot number. A manufacturer's own lot or batch number will be kept on record, but not used for identification within the manufacturing process.

Serial no. All items of a certain complexity will receive this number, e.g. chip wafers, individual chips, flex-prints.

If items exist in more than one version, the same part no. shall only be used if the versions are interchangeable. If a version is incompatible to another, a new item no. needs to be assigned.

Part tracking will be done using a database, accessible via the web. Access control will be in place to protect the data. Where appropriate, labels will be attached to storage boxes or individual items. Labels will carry basic information (part no., lot no., serial no., item name) and a machine readable identifier (bar code or 2D-code). To minimise material for finished modules, laser engraving will be used to identify those parts. The database will provide backtracking functionalities to enable drill-down investigations. For this, any assembly step will allow to map involved parts on a lot or serial number level. E.g. it will be possible to tell from the id of a finished module which lot of flex-prints was used, which wafer(s) the chips came from, which location the chip had on the wafer, and so on.

Tracking scope. Any activity with a part will be reflected by an update in the database, e.g. any operation as shown in Fig. 9.1. Data kept for this purpose depends on the event and may cover (but is not limited to) the following information:

- **Manufacturing step.** Such steps alter the properties of the parts, hence this information needs to be recorded.

Minimal information: Unique material identifier(s), new identifier (if the result is a new part), step taken, success information, date and time, responsible person.

- **Quality test.** Results of tests need to be recorded. Raw data shall be kept either in the database or at a safe location. For the latter, the database needs information on where to find the raw data.

Minimal information: Unique material identifier, raw data (or pointer), summarised result (pass/fail or grade), equipment used (i.e. S/N of a measuring device), date and time, responsible person.

- **Transfer.** Any transfer of a material from one location to another requires an entry in the database.

Minimal information: Unique material identifier, initial and final destination, date and time, responsible person.

- **State changes.** Parts may alter their status outside of normal production steps, e.g. a part gets destroyed, lost in a shipment etc. Such changes need to be kept on records as well.

- **Additional information.** The database will allow to store additional information like notes, pictures, miscellaneous observations etc.

The database should allow users to determine the status of each item, handle orders of stock items, allow to analyse the quality of the overall production and to monitor progress over time.

9.3 Quality control

Test results of quality control points will be stored, including acceptance decisions. For each product and quality step, acceptance criteria will be defined. This can be either a simple accept/reject scheme (simple parts) or a more elaborate grading scheme (chips, modules). Only accepted parts will be used for subsequent manufacturing steps. For items where a grading scheme will be used, the best parts shall be used for making the modules for the final detector.



9.3.1 COMMON QUALITY TESTS

Visual inspection and dimension testing is done wherever appropriate. Visual inspection is carried out to spot obvious mistakes and quality issues like dirt, cuts, damaged surfaces etc. Images of visually inspected items will be taken and stored in the database. Dimension testing is especially required for steps where the material thickness needs to be under tight control, e.g. in gluing steps.

9.3.2 WAFER TEST

A needle card, connected to an FPGA via a custom PCB is used to test the wafers. The custom PCB will be derived from the PCBs developed for the single MuPix test. The software tool will be based on the existing software and firmware, which is developed for the prototype testing. To monitor the status of every chip and the testing procedure, the test results are processed online and the results displayed in a GUI. In addition, all results will be stored in a data base. Due to the high number of sensors which have to be tested, the procedure should be fast and the required time to test one sensor should be in the order of minutes. The test can be divided into three stages:

1. Powering Test
2. General Functionality
3. Performance Test

Failure in one of the steps will lead to a reduced grading and eventually rejection of the chip.

Powering Test. At the beginning, the resistance between HV, power and ground is measured. If there is any short, the chip is rejected. The 1.8 V required to operate the chip are applied and the current is measured. Subsequently, the HV is first ramped up to -15 V, then to -85 V and the current is also measured. For the HV, a current limit of $1 \mu\text{A}$ is set in order to protect the chip. If the HV cannot be applied, the chip is rejected and the testing procedure stops.

General Functionality. To test the general functionality, a set of default DAC values is applied and the PLL is switched on. The DAC values are read back in order to confirm the functionality of the slow control. If the PLL can be locked, different sets of DAC value are set and the measured current is compared to the expected one. Now, the

baseline of the chip is read back and compared to the applied value. Then, different thresholds are set and read back. If all tests are passed, a LED array is switched on, the hit map read out, and the LED spots identified.

Performance Test. In order to keep the chip efficient while controlling the noise, an individual pixel tuning is required. Therefore, the tuning routine is applied. Currently, this is the time consuming part of the testing and has to be accelerated. After the tuning, a quick threshold scan to measure a noise threshold curve is performed. Afterwards, the LEDs are switched on again and the hit maps are stored as well as a plot correlating the time stamps of the hits with the LED trigger times. The phase between the LED pulses and the reference clock of the system is scanned in eight steps and the corresponding time stamps are stored in separate histograms. From these histograms the time resolution and the analogue performance of the chip can be extracted.

9.3.3 PCB TESTING

The PCB will be visually inspected upon receiving and a basic electrical acceptance test will be carried out prior to assembly. This will include a check for connectivity and shorts, functionality of the repeaters and a check for the correct values of the passive components (termination resistors etc.).

9.3.4 FLEX-PRINT TESTING

The flex-prints undergo several steps. At every step, the items will be visually inspected and electrically tested. For this, pogo-pin cards will be designed to quickly test basic electrical properties to deduce conformity to requirements. This includes testing for connectivity, opens and shorts within and among the layers. Samples will be tested for wire-bonding capabilities to ensure process suitability.

9.3.5 FINAL MODULE TESTING

Modules will be fully tested to show that they meet the quality criteria. Testing procedures can be classified in type testing (which includes destructible tests) and individual module testing (non-destructive by design).



TEST	STEP	TARGET VALUES TO PASS
Power	1	1.8 V at 120-200 mA
High Voltage	1	-85 V at 1 μ A
PLL	2	lock the PLL in a reasonable time
DAC	2	Measure the expected currents
Baseline	2	Measure a Baseline of 0.800 V
Threshold	2	Measured Threshold = Applied Threshold
LED 1	2	See a LED spot on different positions on the sensor
Tuning	3	Perform tuning
Threshold scan	3	Noise-Threshold curve should be step function like
LED 2	3	LED spots and time stamp LED trigger correlations

Table 9.1: Overview of the quality criteria for the wafer test. Currents are based on the MuPix7.

Type testing. A sub-sample of modules will be used to thoroughly test their properties. Some of these tests may be destructive. Tests will include:

- **Temperature cycling.** The modules contain laminates, which may get deformed under thermal stress. While the choice of materials have been matched to minimise such effects, robustness of the final module design has to be demonstrated. A stress test with multiple cycles over a temperature range exceeding nominal operation under powered condition will be carried out. If possible, the test may be extended to extreme conditions to provoke failure.
- **Corrosion testing.** Although normal operation should not create corrosive conditions, a test operating a module under humid conditions will be carried out. (Normal operation of the pixel detector happens at temperatures around 10 to 70 °C and the dew point of the experimental hall is assumed to be well below 20 °C.) After success, the test will be extended to condensating conditions, where a module eventually may fail.
- **Resonance test of wire-bonds.** Other experiments observed resonant wire-bond failures while operating them inside a magnetic field, e.g. the CDF-detector acts as a prominent example [33]. The usual mitigation us-

ing encapsulation substances is not suitable for a ultra-low mass design as in Mu3e, hence the need to show that no frequencies occur that may trigger resonance failure. In a magnet with controlled field, modules should be operated and the wire-bonds should be observed using a microscope. A technique using video frames may be envisioned, as described e.g. in [34].

- **Irradiation test.** Expected radiation doses for pixel modules within Mu3e are moderate. To test the radiation hardness of a module under realistic conditions, a campaign at an electron accelerator facility (e.g. MAMI) shall be envisioned. At such a place, an early production grade module will be operated for an extended period of weeks or months, e.g. parasitically behind another experiment.

Individual module testing. Every module will be tested thoroughly to demonstrate its quality. This will comprise the sanity of the pixels, power consumption under nominal operating conditions, check that the module can be properly calibrated and the calibration values are within specifications, response to visible laser light and radiation (using a radioactive source). The procedure will be similar to the chip tests carried out during wafer testing.



PIXEL READOUT

This chapter describes the digital part of the MUPIX sensor as planned for the production device. The digital part of the MUPIX 7 prototype is very similar and was tested in the lab as well as in test beams and found to perform as expected. The scheme described here however incorporates some improvements and additional features.

10.1 Clocking

The MUPIX is synchronized to a single external clock and generates a variety of frequencies for timestamps and to control the readout internally. The following sections describe the clock generation and the usage of the various frequencies in the system.

10.1.1 PLL AND EXTERNAL CLOCKS

The MuPix has an internal voltage controlled oscillator (VCO) which can be tuned with two bias currents (VPVCO and VNVCO). It generates the fast base clock `c1k800p` used for data serialization and from which all other clocks are derived. In order to synchronize the sensors, a phase locked loop (PLL) can be enabled which adjusts the VCO with respect to a reference clock. The PLL locks over a wide range of VCO settings and an optimal working point with <100 ps jitter can be found (Fig. 10.1).

A 125 MHz reference clock has to be provided to run the LVDS link at 1250 Mbit/s. The readout base clock `c1k8n` can be picked up from the sensor to analyse the VCO frequency and the locking to the reference clock (see Figure 10.3).

A detailed study of the VCO and PLL properties is topic of a current bachelors thesis [?]. The temperature dependence was investigated and yielded very promising results for out chosen working point. In an observed temperature range from 0°C - 60°C the PLL locked successfully with a negligible change in jitter as shown in figure 10.2.

For redundancy, the fast base clock `c1k800p` can also be provided by a fast external clock which bypasses the VCO for testing purposes or in case of VCO malfunction.

10.1.2 CLOCKING IN THE DIGITAL PART

The digital part receives a single fast input clock `c1k800p`, for which we assume 625 MHz frequency in the following (the name is due to the fact that the design was targeted at 800 MHz and runs well at even higher frequencies); for other input clock speeds, the multiplier stays the same, see table




Figure 10.1: The measured jitter between reference clock (125 MHz) and chip clock for different VCO settings. The optimal working point for $\text{VPVCO} = 0$ is clearly visible.



10.1. The LVDS data link runs at twice the frequency of clk800p .

The various stages of the serializer, the readout state machine and the timestamp Gray counter (see below) all run at frequencies which are integer divisions of the fast base clock.

Note that `TimestampClk` and `ReadoutClk` are not named clocks in the Verilog used for synthesis of the digital part and are adjustable (towards slower frequencies) via the `ckdivend` and `timerend` configuration registers (6 bits and 4 bits respectively). The signals to the analog part (e.g. `RdCol`) are driven by `ReadoutClk`, and have thus at maximum half that frequency, which is also the maximum hit frequency during a readout cycle.




Figure 10.2: A VPVCO scan around the working point for different temperatures. It shows a small shift of the jitter minimum, which can be explained by an increased current for higher temperatures.



Figure 10.3: MuPix phase locked to reference clock. Orange: 62.5 MHz MuPix clock output clk8n , Purple: 125 MHz reference clock from FPGA.

The frequencies of the timestamps and the readout can be adjusted individually and do not need to be linked, the only concern being that the readout is fast enough to resolve timestamp overflows.

`TimestampClk` drives a 24 bit binary counter, which is transmitted at the beginning of each readout sequence and allows to keep different sensors synchronized even if there are no hits. The lowest 8 bit of the binary counter are converted to a Gray counter which is used for the hit timestamps. The counters can be reset using the `SyncReset` signal, allowing to align all the sensors at the start of a run.

For the production sensor, the phase of `TimestampClk` should be adjustable with regard to clk800p in steps of 1.6 ns (one period of clk800p) in order to be able to synchronize the detector (compensating for cable delays etc.).

The readout state machine generates 32 bit words with a frequency of `ReadoutClk/2`, which are then passed to the serializer part which multiplies the frequency up to two bit at clk800p and then clocks out on both edges of clk800p .

10.2 Readout Link

10.2.1 LINK ENCODING

The MUX sensors have up to three 1250 Mbit/s LVDS output links for data. All data on the link are 8bit/10bit encoded, using the standard IBM encoding [35, 36]. This leaves 1000 MBit/s user data bandwidth.

10.2.2 SYNCHRONIZATION AND ALIGNMENT

After power-up reset, the MUX sends the comma-word K28.5 (0xBC in unencoded form, encoded alternatively as 0x0FA and 0x305). The length of this synchronization sequence can be adjusted with the four bit `resetckdivend` register, which serves as a clock divider controlling an 8 bit counter. The comma word is sent $N_{res} = 2 \times 256 \times \text{resetckdivend}$ times. This long sequence of comma words is necessary for the receiver on the FPGA to lock onto the phase of the incoming signal and then determine the boundaries between 10 bit words. After this, the readout sequence begins.



Clock name	Frequency at 625 MHz base	Frequency at f base	Duty cycle	use
c1k800p	625 MHz	f	50%	base clock, 2bit serializer
c1k4n	125 MHz	$f/5$	40%	byte serializer, 8b/10b
c1k3n2	156.25 MHz	$f/4$	50%	encoded byte serializer
c1k1n6	312.5 MHz	$f/2$	50%	nibble serializer
c1k8n	62.5 MHz	$f/10$	50%	readout base clock
TimestampClk	$625 \text{ MHz}/(\text{ckdivend}+1)$	$\frac{f}{(\text{ckdivend}+1)}$	50%	timestamps
ReadoutClk	$62.5 \text{ MHz}/(\text{timerend}+1)$	$\frac{f}{10(\text{timerend}+1)}$	50%	RO state machine

Table 10.1: Clocks in the digital part of the MuPIX.

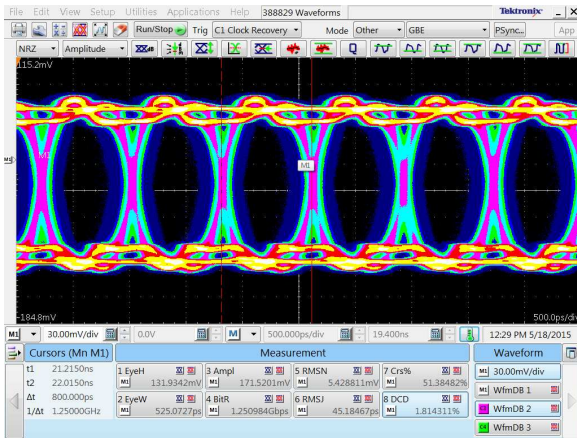


Figure 10.4: MuPIX serial data output at 1250 Mbit/s. Eye height is above 130 mV and the eye width is larger than 0.65 UI (unit interval = time per bit).

10.2.3 LINK PERFORMANCE

The LVDS link can be adjusted using several bias currents. The differential output voltage V_{OD} can be primarily adjusted via the bias current V_{NLVDS} . At maximum, V_{OD} reaches about 200 mV. Pre-emphasis of the signal can be controlled via bias currents to reduce rise and fall times, but decreases the overall amplitude. Figure 10.4 shows the eye diagram of the MuPIX at 1250 Mbit/s. For this data rate, no pre-emphasis is required.

The quality of the serial data stream was checked in a bit error rate measurement searching for 8bit/10bit code and disparity errors. The signal was transferred over a 2 m long SCSI-2 cable from the MuPix test PCB to a Stratix IV FPGA. In a 10 h long test no errors were found, giving an upper limit on the bit error rate $BER \leq 5 \cdot 10^{-14}$ at 90% confidence level.

10.3 Pixel and Column Logic

Figure 10.5 shows the digital logic in each pixel cell and in the column periphery. Note that one physical column is matched to two logical columns, one for the even and one for the odd pixels. In the following we describe the working principle of the digital logic using only active high signals; in the hardware, signals typically change from active high to active low and back with every logic gate in order to use only NAND and NOR gates and thus reduce transistor count and delay times.

If a signal from the pixel surpasses the threshold set in the comparator, the comparator output will go high and the edge detector will see a rising edge, which is then latched in the first SR-latch. The latched signal opens the switch connecting the timestamp storage capacitance with the Gray counter - the current Gray counter value is now stored on the floating capacitance, which will slowly discharge (ms timescale) if not read out. The pixel is now dead until it is read out and will not register new hits. The LdPix signal from the state machine transfers the hit signal into a second SR-latch; the pixel is now ready for read-out.

The pixel priority logic receives a **PriorityIn** signal if no pixel in the priority chain before it has seen a hit. If the pixel under consideration also has not seen a hit, the **PriorityIn** signal is passed down the chain to **PriorityOut**. Otherwise, **PriorityOut** goes low, the D input of the read flip-flop goes high and the pixel waits for a **Read** signal from the periphery in order to write the hit data to the bus.

The column bus is pulled low by the **PullDown** signal. In the column periphery, the LdCol signal registers the inverted output of the pixel priority chain (if the priority chain output is low, at least one pixel has seen a hit). The output of this register in coincidence with LdCol generates the **Read** signal, which generates the **ReadEn** signal in

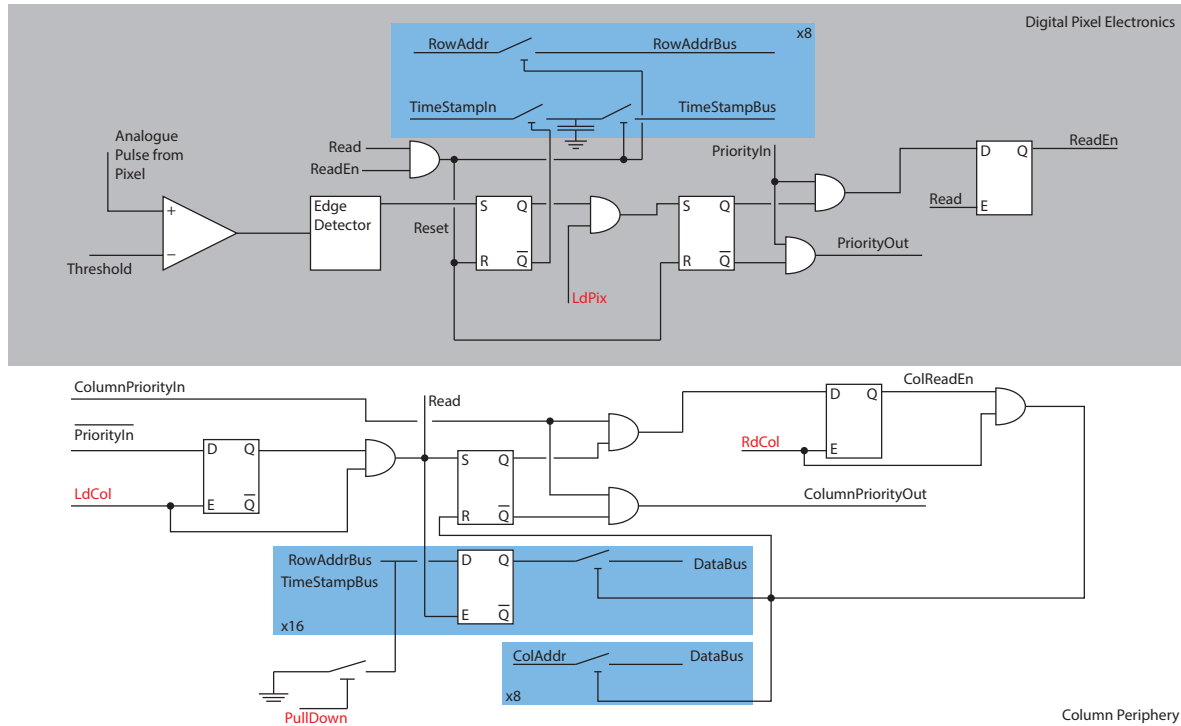


Figure 10.5: Schematic of the logic in the pixel digital part (top) and the column periphery (bottom). Note that the priority signals are chained through the pixels or the columns; the inverse of the last `ColumnPriorityOut` is the `PriFromDet` used in the state machine. For this schematic, the logic has been simplified/changed to use only active high signals, in hardware, signal levels are frequently inverted in order to produce appropriate delays and/or minimize the transistor count. See text for a full description of the functionality.

the pixel currently interrupting the priority chain. This pixel will now open the switches connecting pixel row address and time stamp capacitor to the column bus and reset both SR-latches, which has the pixel ready to accept the next hit. In the column periphery, the values on the column bus are registered in a set of 16 flip-flops. The fact that there is a hit in the column is registered in an SR-latch, which serves as an input to the column priority logic, a copy of the pixel priority logic: If a column has a hit stored, the chain is interrupted and the column ready to write to the chip-wide data bus, otherwise the priority signal is passed on. The inverted output of the priority signal of the last column is seen by the state machine and indicates the presence of at least one hit.

The `RdCol` signal then triggers the write of the column address as well as the pixel data to the chip bus and a reset of the column periphery.

10.4 Readout State Machine

10.4.1 READOUT SEQUENCE

The readout sequence is controlled by the state machine shown in figure 10.6. On exit from the `SyncState` described in section 10.2.2, the state machine checks the `sendcounter` signal. If it is on, the state machine will change to a debug mode sending only counters as described in section 10.4.2. Otherwise the readout sequence starts and continuously cycles through the readout states:

- `StatePD1` The `Pulldown` signal is on and clears the column buses.
- `StatePD2` The `Pulldown` signal is off again.
- `StateLdCol1` The `LdCol` signal moves one hit per column to the column periphery.

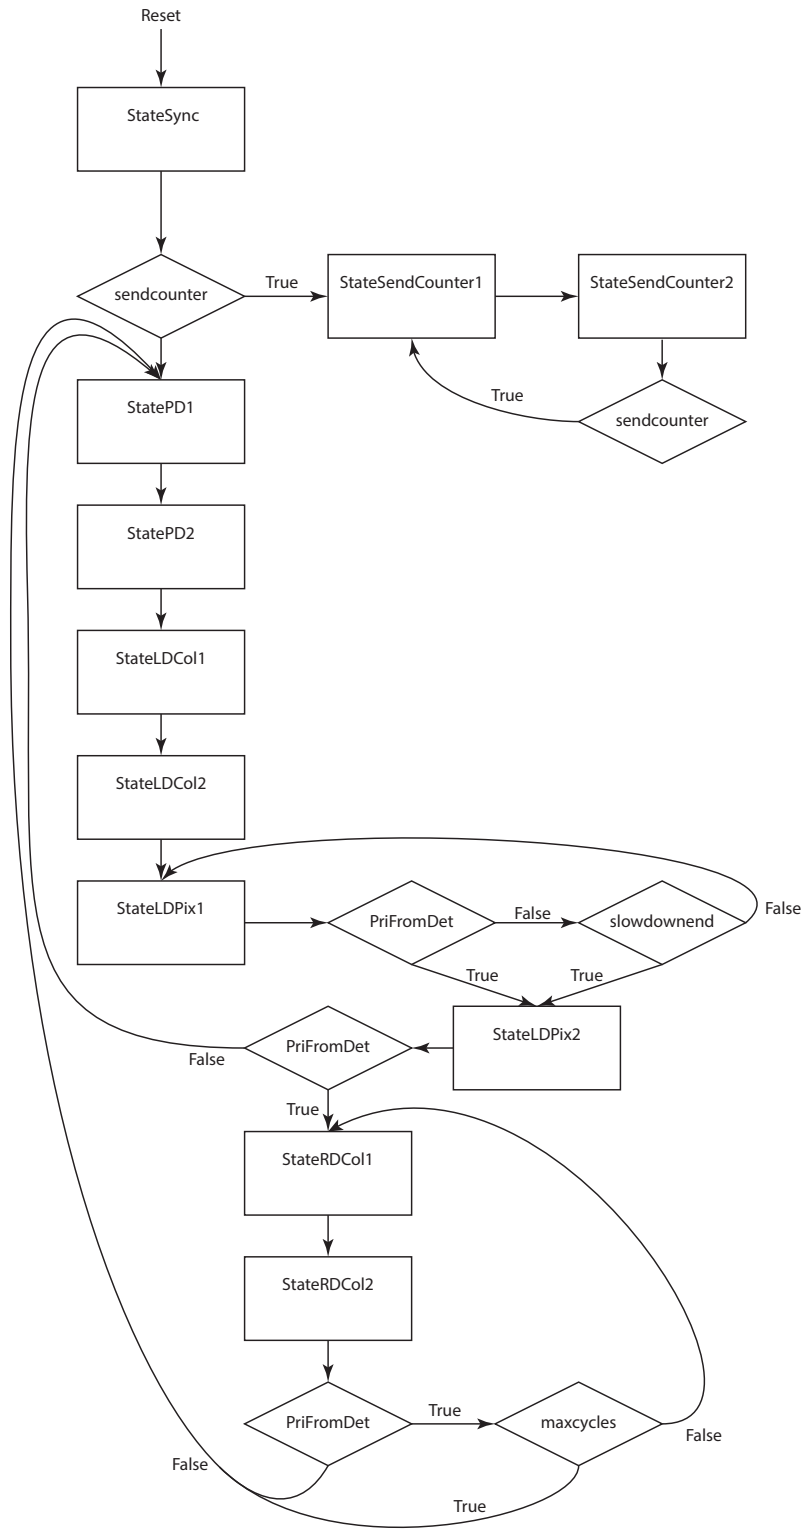


Figure 10.6: MuPIX Readout State Machine



Cycle	State	Data
0	SyncState	K28.5
1	SyncState	K28.5
...	SyncState	K28.5
$N_{res} - 1$	SyncState	K28.5
N_{res}	StatePD1	K28.5
$N_{res} + 1$	StatePD1	K28.5
$N_{res} + 2$	StatePD2	K28.5
$N_{res} + 3$	StatePD2	K28.5
$N_{res} + 4$	StateLdCol1	K28.5
$N_{res} + 5$	StateLdCol1	K28.5
$N_{res} + 6$	StateLdCol2	K28.5
$N_{res} + 7$	StateLdCol2	K28.5
$N_{res} + 8$	StateLdPix1	BinaryCounter[23 down to 16]
$N_{res} + 9$	StateLdPix1	BinaryCounter[15 down to 8]
$N_{res} + 10$	StateLdPix2	BinaryCounter[7 down to 0]
$N_{res} + 11$	StateLdPix2	GrayCounter
$N_{res} + 12$	StateRdCol1	K28.1
$N_{res} + 13$	StateRdCol1	Timestamp
$N_{res} + 14$	StateRdCol2	Column
$N_{res} + 15$	StateRdCol2	Row
...	StateRdCol1	K28.1
...	StateRdCol1	Timestamp
...	StateRdCol2	Column
...	StateRdCol2	Row

Repeat maxcycend times, then back to cycle N_{res}

Table 10.2: Data format for the MUPIX sensors. Note that the information in `BinaryCounter[7 down to 0]` and `GrayCounter` is redundant and allows to test the proper functioning of the Gray encoding.

Cycle	State	Data
0	SyncState	K28.5
1	SyncState	K28.5
...	SyncState	K28.5
$N_{res} - 1$	SyncState	K28.5
N_{res}	StateSendCounter1	BinaryCounter[23 down to 16]
$N_{res} + 1$	StateSendCounter1	BinaryCounter[15 down to 8]
$N_{res} + 2$	StateSendCounter2	BinaryCounter[7 down to 0]
$N_{res} + 3$	StateSendCounter2	GrayCounter

Back to cycle N_{res}

Table 10.3: Data format in the counter mode

- **StateLdCol2** The LdCol signal goes off again.
- **StateLdPix1** The LdPix signal registers hits; if there are hits left from the last cycle, directly go to the next state, else wait either for hits or for `slowdownend` cycles. The binary counter and timestamp are copied to the data output.
- **StateLdPix2** The LdPix signal goes off again. If there are hits (`Priout` is on), move to the next state, else start again at `StatePD1`.
- **StateRdCol1** The RdCol signal is on. Data are sampled and written to output.
- **StateRdCol2** The RdCol signal is off again. If there are still hits and we have read less



than `maxcycend` hits in this cycle, go back to `StateRdCol1` else go to `StatePD1`.

One cycle of the readout state machine corresponds to two eight bit words to the encoder or 20 bits on the serial output. The state machine can be run at a slower speed by setting the `timerend` register to a value larger than 0, which will then keep every state active for `timerend + 1` cycles.

The data format at the encoder/serializer is thus as listed in table 10.2.

10.4.2 COUNTER DEBUG MODE

In the counter debug mode, no readout takes place and the state machine cycles between the `StateSendCounter1` and `StateSendCounter2` af-

ter reset, sending data as shown in table 10.3. The counter mode is ideally suited for extended link tests.

10.4.3 RUN START

The individual pixels in the sensor serve as a very large buffer without the possibility of a clear; the pixel readout thus has to be kept running during run stops. The procedure foreseen for the run start is to assert `SyncReset` for an extended time, which will lead to hits with timestamp zero. At the run start, `SyncReset` is de-asserted at the same time for all sensors, and the front-end FPGA will discard received hits until the received timestamps are different from zero.

PART IV



FIBRE DETECTOR



MILESTONES FOR THE FIBRE DETECTOR

1. Decision on fibre type (round or square) (*Q2/16*)
2. SiPM selection (*Q2/16*)
3. Full simulation and reconstruction of Fibre Detector; complete integration of Fibre Detector into overall reconstruction. (*Q2/16*)
4. SiPM radiation hardness (*Q3/16*)
5. Construction of a technical prototype for the fibre mechanics (attachment, cooling, services, possibly fibre fan-out). (*Q4/16*)
6. Construction of a readout prototype including SiPM arrays, PCB, power distribution and slow control (*Q4/16*)
7. MuSTiC chip prototype (*Q4/16*)
8. Fibre readout integration into experiment DAQ and slow control (Midas) (*Q2/17*)
9. Manufacturing and quality management strategy for fibre ribbon/module production (*Q1/17*)
10. Fibre detector alignment and calibration scheme



SCINTILLATING FIBRE DETECTOR

The cylindrical Fibre Detector is located in the central barrel of the detector at a radius of ~ 6 cm right inside the outer pixel double layers. As shown in figures 11.1 and 11.2 it consists of 2 or 4 staggered layers of round or square multi-clad scintillating fibres with a diameter/size of $250 \mu\text{m}$ and a length of at least 28 cm. The fibres are read-out on both sides with state of the art Silicon Photomultipliers (SiPM). The Fibre Detector aims at providing precise timing information with an efficiency close to 100% of the particle tracks in the order of a few hundred ps by keeping the material budget at the same time as low as possible. This complementary timing information is crucial to efficiently identify coincident signals of electron/positron triplets and suppress accidental background, as well as for determining the sense of rotation and thus charge of recurling tracks.

The baseline design, to be seen as a starting point for many ongoing optimizations, consists of the following parts:

- cylindrical at ~ 6 cm; length (28-30) cm
 - $\leq 900 \mu\text{m}$ total thickness of round or square fibres; 2-4 layers
 - SiPM array readout; both ends (LHCb like)
 - MuSTiC (no preamp) same as Tile Detector
- and has to meet the following constraints:
- very tight space for cables, ASICs and cooling
 - rate (occupancy) up to 250 kHz/fibre
 - time resolution better than 500 ps
 - thickness below $0.4\% X_0$

The present chapter describes the motivation for the inclusion of a Fibre Detector in the Mu3e



Figure 11.2: Round fibre ribbon example.


experiment, the performance requirements and the design concept. The following chapter 12 summarizes the fibre detector R&D status. The detector readout and interfaces to the experiment DAQ and slow control are described in chapter 13. A possible solution for the mechanical integration of the Fibre Detector into the whole experiment is outlined in chapter 14.

11.1 Motivation

As described in chapter 24.3 the Pixel Tracker provides reconstructed tracks with 16 ns or 8 ns timestamps arranged in 40-50 ns reconstruction frames. This corresponds to $\mathcal{O}(10)$ tracks per frame in phase I and $\mathcal{O}(100)$ in a later phase II. Complementary to the vertex-fit (see 28.5), more precise timing is needed for an efficient signal identification and to suppress accidental background further. The additional timing is given by the Fibre Detector in the central detector region and complemented by the Tile Detector (see chapter 15) in the forward and backward regions, which provides even more precise timing. Figure 11.3 shows the fraction of tracks in the geometrical acceptance of the detector which do not get any additional tim-



a) The Fibre Detector inside the Experiment.



b) One out of 6 Modules consisting of 4 ribbons on support wheel.

Figure 11.1: Renderings of the Fibre Detector. The fibre ribbons in blue, the readout boards in green and the pixel support structure in orange.




Figure 11.3: Fraction of tracks in phase Ib which obtain the dominant timing information by the respective detector.

ing information from the timing detectors, those whose timing is given only by the Fibre Detector and the tracks which also end up in the tile detector. Since the time resolution of the Tile Detector is much better, the timing of tracks reaching this sub-detector is dominated by it. Roughly 1/3 of the tracks do not end up in the Tile Detector, for these tracks the best timing is given by the Fibre Detector.

In addition, the timing from the Fibre Detector can be used to improve the tracking algorithms, in particular for determining the charge of tracks approximately perpendicular to the beam which re-curl through the target.

11.2 Requirements

Figure 11.4 shows the timing suppressions depending on the detectors' time resolution. The timing suppression is expressed in a survival fraction for the dominating accidental background from Bhabha electron-positron pairs combining with a Michel positron. The suppression depends




Figure 11.4: Timing suppression (2σ cut) of Bhabha pair accidental background in different timing detector configurations.

on the used reconstruction frame length, where 50 ns (blue), 40 ns (green) and 16 ns (red) were considered. The plots on top show the suppression if only the respective (fibre or tile) detector is present, those at the bottom the combination of the two timing detectors. Time resolutions of 70 ps for the Tile Detector and 400 ps for the Fibre Detector were used in 50 ns reconstruction frames. Using only the Fibre Detector, provides a survival fraction of Bhabha pair accidental backgrounds of $\mathcal{O}(2 \cdot 10^{-2})$, the two timing detectors in combination, $\mathcal{O}(7 \cdot 10^{-3})$. More elaborated timing suppression studies are presented in chapter 34.2.2.

In contrast to the Tile Detector, the Fibre Detector is compelled to minimize multiple scattering, minimizing the material. Multiple scattering in the fibres reduces on one hand the momentum resolution of the detector and effects on the other hand the efficiency of the reconstruction algorithm. (Mainly because the additional




Figure 11.5: Multiple Scattering θ_0 depending on electron/positron momentum and fibre ribbon thickness. r indicates the thickness of round and s of square fibres.

layer of scattering make the matching of segments from the inner to the outer part more demanding.) Figure 11.5 shows the angle straggling depending on particle momentum and ribbon thickness. The scattering scales with $1/\beta cp$ in momentum and $\sqrt{x/x_0 [1 + 0.0038 \log(x/X_0)]}$ in the thickness x , where the radiation length of fibres is $X_0 \sim 45$ cm [37, chapter 27.3].

Since the Fibre Detector, equivalently to the Tile Detector, does not provide any track information, the granularity of these detectors only has to allow a proper matching of hits to the reconstructed tracks within the time resolution of the pixel detector (8 or 16 ns timestamp). The needed granularity is mainly given by the single fibre rate (see table 13.1).

The Fibre Detector dimensions are set by the surrounding Pixel Detector. The minimal length of 28 cm is dictated by the acceptance of track reconstruction in the pixel detector for tracks originating from the target.

As described in chapter 3, re-curling tracks provide the needed precision in momentum measurement. Hence, it is not sensible to place the Fibre Detector outside the outer double layer of pixels, where it would destroy the excellent momentum resolution. On the other hand it is desired to be as far away from the target as possible to minimize the hit rate in the fibres. As described above, this is beneficial to properly match the detector hits to the reconstructed tracks. These two consider-




Figure 11.6: Schematic view of different fibre types.

ations lead to the position right inside the outer pixel double layer at 6 cm.

Due to the very compact design of the whole detector the space at the fibre ends for photon detection and digitization is very limited. Detailed overview over the available space can be found in section 14.2.

11.3 Detector Design

The Fibre Detector consists of 6 modules mounted on the support rings. One module consists of 4 fibre ribbons with a width of 16 mm each. A ribbon consists of 2 or 4 layers of staggered round or square fibres with a diameter/size of 250 μm , approximately 64 per ribbon per layer. Hence the detector consists of ~ 1500 fibres per layer which sums up to between 3000 (2 layers) and 6100 (4 layers) fibres in total. Two different fibre types, round or squares are discussed in more detail in section 12. In both cases Polystyrene (PS) is used as scintillating material surrounded by a double cladding consisting of Acrylic (PMMA) and Fluor-acrylic (FP) compounds. The properties of the two different fibres used can be found in table 11.1. Since the SiPM photosensitivity peaks in blue wavelengths it is beneficial to use fibres with an emission peak in the same area. In order to optimise the time response and minimise pile-up, fibres with short decay times are chosen.

Currently two approaches are investigated to reduce the optical cross-talk between fibres. Either Titanium-Dioxide is added to the glue used in the ribbon production (12.1) or the fibres get an additional ~ 100 nm aluminum coating (12.2.2).

The active core of multi-cladded fibres only covers a fraction of the total width of 92% in the round case and 88% in the squared case. As shown in section 12.1, the glue used to bundle the fibres into ribbons increases the dead area




	round	square
company	Kuraray	Saint-Gobain
core	Polystyrene (PS)	
type	SCSF-81M	BCF12
inner cladding	Acrylic (PMMA)	
outer cladding	Fluor-acrylic (FP)	
cladding [%]	3/1	4/2
refractive index	1.59/1.49/1.42	1.60/1.49/1.42
density [g/cm ³]	1.05/1.19/1.43	1.05
light yield [ph/MeV]	~8000	
trapping efficiency [%]	5.4	7.3
capture angle [deg]	26.7	27.4
attenuation length [m]	>3.5	>2.7
decay time [ns]	2.4	3.2
emission peak [nm]	437	435

Table 11.1: Properties of different fibre types from [38, 39]

between active areas by another few $10\ \mu\text{m}$. A further contribution may be added by a possible $\sim 100\ \text{nm}$ aluminium coating (11.6). Therefore, at least two layers of staggered layers are needed for very high detection efficiency. The measured detection efficiencies of fibre layers are summarized in table 12.2 in section 12.2.3, more detector layers are needed to keep the overall detector efficiency very high. At the same time the number of produced photons increased with more active thickness (It scales linearly with the interaction length). Since the timing resolution scales with $1/\sqrt{\text{number of detected photons}}$, a potential increase in timing performance is possible. The number of layers is limited by the requirement to use as little material as possible to reduce the multiple scattering.

The photons produced in the scintillating fibres are detected on both fibre ends in Silicon Photomultipliers (SiPM). Acquiring the signals on both sides increases the time resolution, helps to distinguish between noise and signal and increases the detection efficiency of the whole system. The latter can be seen in figure 12.6 where a combined (AND) fibre readout is compared with a side independent (OR). Based on standalone Geant4 photon propagation simulations, described in more detail in section 12.3, order of a few 10 photons are expected to arrive at the fibre ends. The shown numbers already include the photon detection efficiency (PDE) of SiPMs of around 40%.


 Figure 11.7: Simulated distribution of photons from round $250\ \mu\text{m}$ fibres on the SiPM active surface with different epoxy layer thicknesses in between.

As shown in figure 12.19 the photons tend to propagate predominantly in the cladding of the fibres and leave the fibres at angles up to 47° [38].

The fibre ends are coupled directly to the surface of SiPMs on both sides. Many different coupling methods such as gluing, optical grease/jelly or the use of constantly applied pressure are possible. The methods are currently under investigation in terms of light yield and detector maintainability.

A typical SiPM packaging comes with a (100-300) μm epoxy layer on top of the active area to protect the substrate. Taking the above mentioned exit angles of the photon into account, the needed active area to detect all photons of one fibre increases by the thickness of this epoxy layer.



To efficiently detect all photons of a fibre with a diameter of $250\ \mu\text{m}$ the active SiPM area has to be at least $750 \times 750\ \mu\text{m}^2$. Figures 11.7 and 12.20 (central) illustrate the photon distribution in the SiPM active layer.

11.3.1 READ-OUT SCHEMES

By grouping the fibres into vertical columns the fan-out can either be completely omitted if the ribbons are directly attached to the SiPMs or minimized if fibres are grouped to the centre of the columns to maximize the light yield. These two concepts are illustrated in figure 11.8.

The first scheme is very space-saving and simplifies the mechanical construction. Due to the dead area between the SiPM columns the photon detection efficiency is reduced.

The second scheme needs a more elaborated mechanical structure and more space for the fan-out zone and leads to the use of broader SiPM cells. As in the discarded single readout scheme, the cells need a minimal width of at least $750\ \mu\text{m}$ (see 11.7).

A column-wise readout scheme reduces the required space for a possible fan-out and reduces the number of channels below 1536 per side (< 50 MuSTiC per side). The occupancy, as listed in table 13.1, increases in this design.

To match each fibre to a SiPM with a rather large active area ($\sim 0.6\ \mu\text{m}$) the ribbons need to be fanned-out in a transition zone between active volume and detection plane. The geometrical fan-out of the ~ 6100 fibres violates the space constraints given by the detector design. Furthermore, the large number of channels translates into lots (e.g. more than 190 MuSTiC ASICs per side) of readout electronics, which again does not fulfil the space requirements.

11.3.2 SILICON PHOTOMULTIPLIERS

SiPMs are well suited to read out the Fibre Detector. Their compact form factors and moderate high voltage ($\sim 70\ \text{V}$) account for the limited space and they provide the required channel density. They are insensitive to magnetic fields. The photon detection efficiency (PDE¹) of the order of 40%, single photon detection and very fast intrinsic time response (below 1 ns) are the key features

¹With contributions from quantum efficiency and geometrical fill factors.




Figure 11.8: The two different readout schemes grouping fibres to columns.

for the use in the Fibre Detector. The high gain ($\sim 10^6$) allows the use of mixed-mode ASICs (like MuSTiC) without any pre-amplifications. Typical dark-rates of SiPMs, which scale linearly in active area and exponentially in temperature, are around a few kHz up to a few MHz at room temperature; depending on the SiPM type.

To read out fibres grouped into vertical columns SiPM arrays with different column width are currently investigated. The baseline concept foresees a SiPM array solution (similar to LHCb). Since the two experiments have different requirements a custom made solution in the same style is proposed.

As a reference, the properties of Hamamatsu S10943 SiPM arrays developed for LHCb are summarized below.

Hamamatsu S10943 SiPM arrays

The LHCb experiment plans to install a scintillating fibre tracker in the next upgrade [40]. Two SiPM manufacturers, Hamamatsu and KETEK, have developed dedicated devices for this tracker. The baseline concept foresees the use of the same Hamamatsu SiPM arrays shown in figure 11.9, which was developed for the LHCb experiment. Different than in LHCb, where the SiPMs are operated around $-40\ ^\circ\text{C}$, Mu3e sensors operate at room temperature.

The latest version of this array has a pixel size of $57.7 \times 62.5\ \mu\text{m}^2$ with trenches between the pixels to reduce the pixel to pixel cross-talk. This new technology allows to operate the sensors at higher over-voltage ($\sim 4\ \text{V}$) to increase the PDE. Table 11.2 summaries the most important features of the Hamamatsu S10943 SiPM arrays.

The arrays are powered in common cathode scheme to minimise the number of needed chan-



property	value
breakdown voltages	slightly below 55 V
temperature coefficient	53.7 mV/K at 25°
crosstalk at $\Delta V = 2$ V	~ 4 %
crosstalk at $\Delta V = 4$ V	~ 16 %
PDE peak	470 nm
PDE at $\Delta V = 2.5$ V	32 %
PDE at $\Delta V = 4.5$ V	45 %
geometrical fill factor	61 %
epoxy layer	(100-120) μ m

Table 11.2: Properties of Hamamatsu S10943, from [40, 41].




Figure 11.9: Picture of a 128 (2x64) channel LHCb SiPM array.

nels. The overall current consumption of one array is expected to be well below 1 mA.

Radiation Damage

The effect of radiation (mostly positrons) is under investigations. In the baseline design of phase Ib a positron/electron flux per mm^2 active sensor area inside the thin active SiPM layer of ~ 0.9 kHz, respectively ~ 1.7 kHz depending on the side, is expected. Figure 11.10 shows the expected spectrum of deposited energy, on average 42 keV. This sums to a integrated positron flux of $0.8 \cdot 10^{10}$ ($1.4 \cdot 10^{10}$) e^+ / mm^2 per year; 55 mJ (97 mJ) corresponding to 24 Gy (42 Gy) per year.

First results, as shown in 11.11 show a significant increase in the dark rate. It is to be noted that the different energy spectra prevent a direct translation from irradiation test with ^{90}Sr to radiation damage in the experiment. Recovery due to annealing is currently under investigation.

11.3.3 RO HARDWARE

In the following the Fibre Detector support hardware is discussed independently of the chosen SiPM arrays (baseline are Hamamatsu S10943 arrays) and readout electronics (baseline MuSTiC).




Figure 11.10: Energy deposition of positrons and electron in the thin active layer of the SiPM arrays of the Fibre Detector.




Figure 11.11: Dark rate increase of SiPM (Hamamatsu S12571-050P, $1 \times 1 \text{ mm}^2$ active area, $50 \times 50 \mu\text{m}^2$ pixels) irradiated by 20 MBq ^{90}Sr .



The SiPM arrays are either connected with flex-prints or directly soldered to a flex board linked to the same PCB where the readout electronics and HV distribution are located. The electronic components of this board are described in more detail in chapter 13. This fibre readout board sits directly on top of a cooled support plate. There is one cooling plate per module (6 modules) and one PCB per cooling plate/module. This PCB provides clock distribution from the reference clock, slow control, SPI, LVDS links to the main DAQ, power and ground distribution.

11.4 Calibration & Monitoring

Calibration and monitoring studies are currently ongoing. The idea to mount a commercial communication fibre alongside all SiPM arrays to have the possibility to inject light on one hand directly into all attached SiPM cells and on the other hand into the fibres needs further investigation.



FIBRE DETECTOR PROTOTYPES

This chapter describes the measurements performed with different detector prototypes and discusses the Monte Carlo simulation results. All prototypes are based on the thinnest available scintillating fibres. Section 12.1 refers to prototypes built with the $250\ \mu\text{m}$ diameter round fibres. Section 12.2 describes the prototypes made of $250 \times 250\ \mu\text{m}^2$ square fibres. And section 12.3 summarizes the current Monte Carlo simulation studies.

12.1 Round Fibres

Four layer ribbons with a length of 36 cm and a width of 16 mm were produced at University of Geneva.

Antoaneta, Sandro: Round Fibre Production, glue studies, picture

12.1.1 ROUND FIBRE SETUP

The setup shown in figure 12.1 was used to investigate the round fibre ribbons.

16 mm wide four layer ribbons with a length of 36 cm have been developed. Two layers with 8 adjacent fibers per layer have been individually read out to investigate the ribbon performance. The fibres have been fanned out and polished by hand. They have been connected to a Hamamatsu S12571-050P (active area $1 \times 1\ \text{mm}^2$, $50 \times 50\ \mu\text{m}^2$ pixels) SiPM. 16 sensors per side have been biased with the same voltage. The SiPM signal has been amplified by a custom three transistor amplifier that generates a signal amplitude of about




Figure 12.1: Round fibre ribbon setup with SiPM board and attached pre-amplifier board.

50 mV for the first photo-electron. For a better understanding of the fibre response the signal waveform has been digitized with several DRS4 evaluation boards [42] at 5 GSamples/s. The gain of each SiPM and of each analogue amplifier can vary slightly. Therefore an automatic calibration procedure has been implemented. To reduce the electronic noise and the electrical crosstalk the integral of the signals have been used to calculate the measured number of photons (Fig 12.2).

As an external trigger, the setup uses two crossed 1 mm thick round fibres connected to a photomultiplier tube. This trigger has been placed behind the fibre ribbon and ensures the study of minimum ionizing particles. Also the trigger signals have been digitized with the DRS board. This setup has been used in the labora-




Figure 12.2: Signal integral versus amplitude distribution. The smearing of the amplitudes due to electronic noise and signal crosstalk is clearly visible.

tory by irradiating it with a ^{90}Sr source and in the piM1 test facility at PSI.

The extraction of the time needs a post processing of the digitized data. The time is calculated by the center of gravity of the signal derivative in the interval between the zero-crossings (Fig 12.3).




Figure 12.3: Sample of a digitized waveform. The raw data is shown in blue, the processed data in black. The red dotted line shows the definition of the time-stamp.

12.1.2 ROUND FIBRE RESULTS

Antoaneta, Sandro: Efficiency, cross-talk (Ti, no Ti), time resolution

Figure 12.4 shows the time difference ($\Delta T = T_1 - T_2$) of two channels interlinked by round fibre without any additional titanium dioxide in the glue with a sigma ($\sigma_{\Delta T}$) of around 2.2 ns. This results in a intrinsic time resolution of the system of $\sigma = \sigma_{\Delta T}/2 = 1.1$ ns.




Figure 12.4: Time distribution of signals from the two ends of one round fibre inside a ribbon produced without additional titanium dioxide in the glue. ($\Delta T = T_1 - T_2$, with an arbitrary time offset)

12.2 Square Fibres

Several versions of prototypes have been built in order to become acquainted with this technology. This included the selection of the fibers (single clad versus multi-clad), the fiber quality (size, geometrical defects etc.), the array/ribbon production, the array optimization (Al coating over the total length of the fiber), the minimal fiber-MPPC relative alignment and the MPPC characteristics (photosensor size, pixel size, photosensor entrance window etc.).

12.2.1 SQUARE FIBRE SETUP

The prototypes are based on $250 \times 250 \mu\text{m}^2$ Saint-Gobain BCF12 multi-clad square fibres. The coupling between the fibre and the photosensor has been carefully studied, taking into account also the maximum allowed misalignment between fibre and SiPM without light collection losses, for a given SiPM. In all cases, the ends of the fibre were polished with a diamond cutting head. The prototypes were coupled to a standalone electronic and data acquisition system. A single fibre readout scheme was used to explore the ribbon performance in more detail. All SiPMs are biased at the same voltage (bias voltage: ≈ 66.5 V for the Hamamatsu S12825-050C, bias voltage: ≈ 55.0 V for the Hamamatsu S13360-1350CS), around the operating voltage value provided by the data sheet. A fine-tuning calibration of each




Figure 12.5: The double layer scintillating fibre prototype V4.1 and the lateral view of the fibre assembly.

SiPM is performed offline, using the reconstructed and well separated photoelectron peaks. The preamplifiers, based on the wideband mini-circuit MAR-6 amplifier, are optimized for signals of a few photoelectrons, the main characteristics of which being the high gain (a factor 100; 1 phe has an amplitude of ≈ 60 mV) and the low peak-to-peak noise (< 10 mV) [43]. The data acquisition system (DAQ) is based on several DRS evaluation boards operated in daisy chain mode [42]. All SiPM waveforms are digitized at a frequency of 5 GSamples/s. A custom waveform analysis is performed to extract the amplitude, the charge, the timing of each event together with other operational characteristics such as pile-up rejection.

We will describe first the prototype V4.1 (sec. 12.2.2) then the so called large prototype (sec. 12.2.3), based on the same technology as the prototype V4.1.

12.2.2 PROTOTYPE V4.1 PERFORMANCES

The prototype V4.1 is a segmented detector, where each fibre is an independent detection element. It is made of two staggered layers of Saint-Gobain BCF12 $250 \mu\text{m}$ multi-clad square scintillating fibres plus a single fibre used for trigger purposes. The mean fibre length is about 25 cm. Each fibre is coupled to two SiPMs (Hamamatsu S12825-050C, active area: $1.3 \times 1.3 \text{ mm}^2$, pixel size: $50 \mu\text{m}$), one per fibre end. Figure 12.5 shows the prototype.

The prototype detector was first characterized in the laboratory using a ^{90}Sr source and was then exposed to positron and muon beams, with an initial momentum of $28 \text{ MeV}/c$ impinging on the de-

tector after leaving the vacuum beam line through a $190 \mu\text{m}$ mylar window (πE5 channel at PSI).

The detector showed a homogeneous response: all channels worked properly, with a relative light collection variation less than 10%.

Figure 12.6 shows the typical charge spectrum of a single fibre for minimum ionizing particles (m.i.p.), when the charge collected by both SiPMs is combined in an OR logic (at least one SiPM with one or more photoelectrons, phe: $\text{SiPM}_1 \geq 0.5$ phe OR $\text{SiPM}_2 \geq 0.5$ phe) and in AND logic (both SiPM collected at least one phe: $\text{SiPM}_1 \geq 0.5$ phe AND $\text{SiPM}_2 \geq 0.5$ phe). For a single fibre a detection efficiency of $\epsilon(\text{OR}) \approx 92 \pm 2\%$ and $\epsilon(\text{AND}) \approx 74 \pm 2\%$ has been measured. For the single layer and the double layer we get respectively $\epsilon(\text{OR}) \approx 95 \pm 2\%$ – $\epsilon(\text{AND}) \approx 76 \pm 2\%$ and $\epsilon(\text{OR}) \approx 99 \pm 2\%$ – $\epsilon(\text{AND}) \approx 88 \pm 2\%$.

The optical cross-talk between adjacent fibres was negligible ($< 1\%$) when coating the fibres with 100 nm of aluminum, see figure 12.7. Under such a condition, spatial resolutions $< 100 \mu\text{m}$ are achievable.


Figure 12.8 (a) shows the timing resolution of the mean time for m.i.p. particles hitting a single fibre, giving a $\sigma = 541 \pm 2$ ps with a single gauss function fit (charge threshold ≥ 0.5 phe). It is evaluated by calculating the difference between the time of the two SiPMs divided by a factor 2: $(t_1 - t_2)/2$. An offline constant fraction discrimination is used to extract the time of each waveform, with a 20% threshold on the maximum signal amplitude. If two fibres fire (double hit events) we measured a timing resolution of $\sigma = 397 \pm 2$ ps (Figure 12.8 (b)). Results using a $500 \mu\text{m}$ multi-clad scintillating fibre can be found in [44].

12.2.3 LARGE PROTOTYPE PERFORMANCES


The Large Prototype is also a segmented detector, where each fibre is an independent detection element. It is the large version of the prototype version 4.1 with the following main differences:

- more channels
- four fibre layers
- longer fibre
- new SiPM series

It is made of four layers of Saint-Gobain BCF12 $250 \times 250 \mu\text{m}^2$ multi-clad square scintillating fibres, where one layer is used for trigger purposes.




a) OR logic




b) AND logic

Figure 12.6: The charge spectrum for m.i.p. in a single fibre coupled to two SiPMs, combining their charge. OR logic: at least one SiPM with one or more phe (a). AND logic: both SiPM collected at least one phe (b).



a) without Al



b) with Al

Figure 12.7: The correlated charge of two adjacent fibres are plotted one with respect to the other (N_{phe2} , N_{phe3}). The optical cross-talk without (a) and with (b) Al deposit around each fibre of the array.




Figure 12.8: The timing resolution of the mean time for m.i.p. going through a single fibre layer. Different than in figure 12.4 here $\Delta t = (t_1 - t_2)/2$ is shown.

The mean fibre length is about 50 cm. Each fiber has been coated with a 100 nm Al deposit. Each fibre is coupled to two SiPMs (Hamamatsu S13360-1350CS, active area: $1.3 \times 1.3 \text{ mm}^2$, pixel size: $50 \mu\text{m}$), one per fibre end, for a total of 64 channels. The main differences between the previous and the current SiPM series are listed in the Table 12.1. The S13360 series is characterized by a strongly reduced cross-talk probability P_{ct} ($P_{ct} \leq 2\%$ at the operating voltage and room temperature), a lower breakdown voltage V_{bd} ($V_{bd} \approx 52 \text{ V}$ at room temperature) and thanks to the low P_{ct} and the reduced darkcount rate R_{dc} ($R_{dc} \approx 100 \text{ KHz}$ at the operating voltage and room temperature) the overvoltage $\Delta V = V_{op} - V_{bd}$ can be at level of 3 Volt, with an increasing of the photon detection efficiency PDE of the photosensor (PDE $\approx 40\%$). More details can be found in the data sheets [?].

Figure 12.9 shows the prototype and the fiber arrangement (lateral view).

The four-layer prototype allows the performance study of several layer combinations: the four single layers, treating each layer as an independent detector, the four double-layer configurations, where two staggered/parallel layers are considered as a unique detector, the four three-layer configurations, where three staggered/parallel lay-

	S13360	S12825
V_{bd} [V]	53 ± 5	65 ± 10
$(V_{op} - V_{bd})$ [V]	3.0	2.6
PDE (at V_{op}) [%]	40	35
gain (at V_{op}) [10^6]	1.7	1.25
geometrical fill factor [%]	74	62
dark count (at V_{op}) [kcps]	90-270	170-350
crosstalk (at V_{op}) [%]	1	40
temperature coef. (mV/ $^\circ \text{C}$)	54	60

Table 12.1: Electrical and optical characteristics of the Hamamatsu S13360 SiPM and S12825 SiPM.




Figure 12.9: The four layer scintillating fibre prototype and the lateral view of the fibre assembly.

ers are considered as a unique detector. For each arrangement we have in principle at least four combinations due to the fact that each layer can be used as a trigger.

The prototype detector was first characterized in the laboratory using a ^{90}Sr source and was then exposed to a positron beam (with muon and pion contamination at a level of 10%), with an initial momentum of 115 MeV/c impinging on the detector after leaving the vacuum beam line through a 250 μm mylar window (πM1 channel at PSI).

The detector showed a homogeneous response and the detector performances achieved with prototype V4.1 have been confirmed. We describe here only the configurations that at the moment are considered realistic candidates for the experiment during phase I. These are the double layer and triple layer configurations, each one considered as an independent detector with an optimized “fiber grouping and photosensor coupling”, which minimizes the amount of readout channels, as shown in Figure 12.10. We report the single layer performances for completeness.

For a single layer a detection efficiency of $\epsilon(\text{OR}) \approx 95 \pm 2\%$ (with a threshold > 0.5 phe) and $\epsilon(\text{AND}) \approx 74 \pm 2\%$ (with a threshold > 0.5 phe) have been measured. Increasing the threshold to 1.5 phe we get $\epsilon(\text{OR}) \approx 79 \pm 2\%$, $\epsilon(\text{AND}) \approx 34 \pm 2\%$.

For the double layer configuration we have $\epsilon(\text{AND}) \approx 93 \pm 2\%$ (with > 0.5 phe) and $\epsilon(\text{AND}) \approx 75 \pm 2\%$ (with > 1.5 phe). For the three layer configuration we get $\epsilon(\text{AND}) \approx 96 \pm 2\%$ (with > 0.5 phe) and $\epsilon(\text{AND}) \approx 88 \pm 2\%$ (with $>$

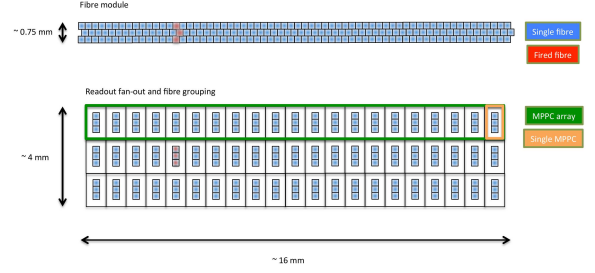


Figure 12.10: Fibre grouping, fibre fan-out and photosensor coupling. In red a typical fibre grouping and how it looks like at the end of the ribbons and in front of the photosensor.

1.5 phe). Table 12.2 summarises the main results.

$\epsilon(\text{AND})$ for Q_{thr}	> 0.5 phe	> 1.5 phe
Single layer [%]	74 ± 2	34 ± 2
Double layer [%]	93 ± 2	75 ± 2
Three layer [%]	96 ± 2	88 ± 2


Table 12.2: Detection efficiency for different layer configurations. For the multi-layer configuration the sum of the charge from several fibres is used, assuming an array readout.

Figure 12.11 shows the typical charge spectra for the three layer configuration, adding the charge collected from several fibres (“array readout”) on the same side grouped offline (Q_{array1} and Q_{array2}) in order to reproduce the array readout. Figure 12.11 (a) and (b) correspond respectively to a threshold of $Q_{array_i} (i = 1, 2) \geq 0.5$ phe and $Q_{array_i} (i = 1, 2) \geq 1.5$ phe (both in “AND” logic).


An example of the timing spectrum for the Large Prototype is given in Figure 12.12 for a single fibre. The variable used is the difference of the photon arrival times at the photosensors divided by 2, $(t_1 - t_2)/2$. Using the leading edge method with a threshold of 0.5 phe we measured a $\sigma = 750 \pm 17$ ps, using a double gauss fit function (core fraction = 68%).

12.2.4 SQUARE FIBRE CONCLUSIONS

In conclusion we have been satisfied with the detector performances achieved with the prototype version 4.1 and confirmed with its big version, the



a) $Q_{array1} \geq 0.5 \text{ phe}$ AND $Q_{array2} \geq 0.5 \text{ phe}$



b) $Q_{array1} \geq 1.5 \text{ phe}$ AND $Q_{array2} \geq 1.5 \text{ phe}$

Figure 12.11: Typical charge spectra for the three layer configuration performing the "array readout" (see text). Two different thresholds on the collected charge has been considered: (a) and (b).




Figure 12.12: The mean time distribution for m.i.p. going through a single fibre of the Large Prototype. A timing resolution of $\sigma = 750 \pm 17 \text{ ps}$ is measured. The time has been extracted using the leading edge with a threshold equal to the half amplitude of one photoelectron.

large prototype. It matches well the expected detector performances providing a full detection efficiency for minimum ionizing particles (m.i.p.) and a good timing resolution ($< 1 \text{ ns}$ in sigma, with a threshold of 0.5 photoelectron, phe) with already two staggered fiber layers. This represents one of our milestones towards the construction of the detector.

12.3 Standalone Monte Carlo Simulation

«««< HEAD Standalone Monte Carlo (MC) simulations based on Geant4 were used to simulate the photon propagating inside the scintillating fibres. In the case of squared fibres SiPM response such as geometrical effect, photon detection efficiency, pixel recovering time, dark current rate, after-pulse, pixel optical cross-talk, pixel inhomogeneity, =====

Standalone Monte Carlo (MC) simulations based on Geant4 [45] were used to simulate the photon propagating inside the scintillating fibres. In the case of squared fibres SiPM response such as geometrical effect, photon detection efficiency, pixel recovering time, dark current, after-pulse, pixel optical cross-talk, pixel homogeneity, »»»>




Figure 12.13: Scintillating photon path (green lines) through a $250 \times 250 \mu\text{m}^2$ multi-clad scintillating fibre. The first and second clad are visible.




Figure 12.14: The number of scintillating photons reaching one end of the fibre as a function of the energy deposit by a m.i.p. at the middle of a 25 cm long fibre.

origin/simon time response were included and the simulation was validated with the data [46])

12.3.1 SQUARE FIBRE

MC simulations were used to understand the detector response and optimize it.

Figure 12.13 shows the path of scintillating photons through a $250 \times 250 \mu\text{m}^2$ multi-clad scintillating fibre. The scintillating photons reaching one end of the fibre as a function of the energy deposit by a m.i.p. is given in Figure 12.14. To obtain this numbers m.i.p are shot through the middle of a 25 cm long fibres in a MC simulation.

The effect of the coupling of the fibre with the SiPM is taken into account: a layer of optical grease and the epoxy window of the SiPM




Figure 12.15: The distribution of the scintillating photons (ph) reaching the end of the fibre is shown in black, the fraction reaching the sensitive area of the SiPM in red (fill factor), the distribution of photoelectrons (phe) taking also QE into account in green and the number of photoelectrons finally providing a signal in blue (discharge probability and SiPM time recovering effect).

is added in the simulation. The fibre is centered to the SiPM. Also the SiPM PDE is taken into account. The effects of the different contributions of the PDE ($\text{PDE} = \text{geometrical fill factor} \times \text{QE} \times \text{discharge probability}$) on the number of photons/photoelectron are shown in figure 12.15.

The spatial distribution of the photoelectrons at the active pixel surface is shown in Figure 12.16. It is important to note that even if the fibre has a size of $250 \times 250 \mu\text{m}^2$ the photoelectron spatial distribution extends up to $750 \times 750 \mu\text{m}^2$. This has to be taken into account in the Fibre Detector design.

To fully study the detector response the electronics chain and the data acquisition system are added to the simulation. Examples of waveforms associated with signals of one, two and three photoelectrons are shown in Figure 12.17. All the parameters for the waveform description (such as rise time, decay time and noise level) are extracted from the data. The comparison of the data collected with the prototype V4.1 with the ^{90}Sr source and the associated MC simulation spec-




Figure 12.16: Spacial distribution of the photoelectrons at the SiPM surface. The dead zones and the pixel size are clearly visible.

trum is given in Figure 12.18. The simulated spectrum and the data are in a good agreement.




Figure 12.17: Typical simulated waveforms for one (black), two (red) and three (blue) photoelectrons.




Figure 12.18: The charge spectrum collected with the prototype V4.1 using the ^{90}Sr source (black) and the associated MC simulation spectrum (red).



12.3.2 ROUND FIBRE

Figure 12.20 shows that the photons leave the round fibres predominantly in the cladding.

Figure 11.7 shows the photon distribution in planes at different distances to the fibres assuming a epoxy layer in between. This photon spread has to be taken into account for a efficient photon detection.

The simulated photon yield is roughly in agreement with the simulations. Figure 12.19 shows the comparison between simulation and measurements in the case of round fibres.




Figure 12.19: Comparison of the number of detected photons between simulation and measurements of round fibre ribbons.




Figure 12.20: Simulation of photon propagation inside round fibres. They leave predominantly in the cladding.



FIBRE READOUT

This chapter is divided into three parts. The first part describes the general requirements a readout scheme has to fulfil to be used in the Fibre Detector if fibres are grouped in columns of SiPM arrays. The second part describes the MuSTiC baseline solutions and in the third possible alternative designs with focus on different readout electronics are discussed.

13.1 Requirements

Grouping the fibres to columns of SiPM arrays has the advantage that it reduces the number of channels; one readout channel per column. On the other side it increases the rate per channel. The corresponding numbers can be found in table 13.1.

The stated occupancy denotes the number of firing channels per frame (per side). If this number approaches values around 20% (more studies are ongoing) matching between fibre hits/clusters and tracks gets more tedious.

As shown in chapter 12, the number of detected photons is rather low (in single fibres the mean number of photons is ~ 5 , see section 12.6). To provide an efficient detector, thresholds down to single photons are required. On the other hand, operating with thresholds above single photons is desirable to exclude dark counts.

The energy information, which is given by the number of firing SiPM cells, is not mandatory for the Fibre Detector, but would be highly beneficial for calibration, redundancies and better clustering algorithms. The Fibre Detector provides time

information by matching fibre hit times to reconstructed tracks. Nevertheless, energy information could increase the timing resolution slightly if used in clustering algorithms or time walk corrections. Furthermore, it would be a very welcome tool for calibration and performance monitoring.

13.2 Baseline Design

The base-line design consists of a STiC based Fibre Detector readout concept. STiC3.1 as described in sections 13 and 18 and its successor MuSTiC, an ongoing development especially for the Mu3e experiment, are mixed mode ASIC chips in UMC 0.18 μm CMOS technology. The idea is to use the same ASIC for both timing detectors.

In contrast to the PET applications (usually LYSO crystals as scintillator) for which STiC was originally designed and the Tile Detector where signals consist of ~ 1000 photons, the Fibre Detector readout has to operate with very few photons. Triggering at the level of single photons is required. STiC can fulfill this requirement in a mode, where the second threshold (*energy threshold*) is switched off. In this mode the chip operates with only one threshold (*timing threshold*), the time is given by a leading edge discriminator. Only a time information in 50 ps^1 bins is provided; no energy information is available.

As shown in [47] a single pixel time resolution of 120 ps can be achieved with STiC. This time resolution and the timing bin size ($\sim 50 \text{ ps}$) of the

¹mean bin width = $\frac{1}{32\text{clk}_{\text{ref}}}$, STiC PLL is designed for $\text{clk}_{\text{ref}} = 622.5 \text{ MHz} \pm 50 \text{ MHz}$



	single	array 250 μm	500 μm	1000 μm
number of channels per side	4608 ¹	1536	769	384
min active area [mm^2]	5184	3840	3840	3840
rate per channel [kHz]	347 (375)	864 (1044)	1274 (1564)	1914 (2501)
occupancy % (50ns frames)	2.7 (2.8)	4.3 (5.2)	6.2 (7.8)	9.6 (12.5)
number of needed STiCs	144	48	24	12
needed space	large	moderate	moderate	little

Table 13.1: Mu3e Fibre Detector requirements overview for different fibre designs. Values for 3 layers of round fibres are stated, values for 3 layers of squared fibres are in parentheses.

¹ 6 modules \times 4 ribbons \times 3 layers \times 64 fibres/layer.

STiC well exceed the Fibre Detector’s requirement of a few 100 ps.

STiC supports differential as well as single-ended readout schemes. The baseline Hamamatsu S10943 arrays have a common cathode, therefore STiC is used in single ended mode.

The possibility to tune the input terminals with a 6 bit DAC in a range of about 700 mV allows to align the bias voltages within of all columns in a SiPM array and correct for local temperature fluctuations.

The implemented zero pole cancellation circuit reduces the dead time of the STiC (down to 40 ns) and base line shifts due to high (dark) rates.

STiC3.1 was extensively used for testing and evaluation. Results are shown in section 13.2.1.

The main changes in MuSTiC, which is currently under development, is a faster LVDS link to overcome the rate limitation caused by that link. An important feature of both versions are the internal L1/L2 FIFO buffers before the chip serializer. This feature allows to handle true random rates up to the rate limit given by the LVDS links. A more detailed description of the digital control logic can be found in [48, p. 71]. In addition a *time mode* allows to transmit only the time information, which decreases the number of bits per event from 47 to 26. This results in another 81 % increase in event rate. Furthermore, the new version will only have half the channels (32). Table 13.2 shows a comparison of the two chips.

13.2.1 STiC3.1 MEASUREMENTS

The performance of the available STiC, version 3.1, is confirmed by measurements independent of the developing group (KIP Heidelberg). Furthermore, it was shown that the theoretical event rate limit of this version of almost 2.6 Mevents /chip,

	STiC3.1	MuSTiC
number of channels	64	32
LVDS speed [Mbit/s]	160	1250
8b/10b encoding	yes	yes
event size [bit]	48	47
<i>time mode</i>	-	26
event rate / chip [MHz]	~ 2.6	~ 20
<i>time mode</i>	-	~ 38
event rate / ch [kHz]	~ 40	~ 650
<i>time mode</i>	-	~ 1200
power per channel [mW]	30	?
size [mm x mm]	5x5	5x5
number of PLLs	2	1

Table 13.2: Comparison of the two STiC versions.

set by the 160 Mbit/s LVDS link, is reachable with true random events. A detailed description of this reference measurements can be found in [49].

The same setup as described in 12.1.1 was used to prove that it is possible to measure timing in round fibre ribbons with STiC3.1. In a first step, threshold scans of dark counts as shown in figure 13.1 are used to determine proper working points. The measured timing results (see figure 13.2) are in good agreement with measurements done with the pre-amplifier and DRS based readout, described in section 12.1.

13.2.2 ADDITIONAL CONNECTIONS

Besides the main fast LVDS links for data transmission and the Serial Peripheral Interface (SPI) for slow control there are a few connections to the chip which can be used for debugging and performance monitoring.



Energy/Performance Monitor

Even though STiC does not provide any digitized energy information for few photon signals, it is possible to look at the direct comparator output before the internal TDC. This provides a Time over Threshold (ToT) information, which corresponds to an energy measurement. Unfortunately, the resolution of this ToT signal of the order of 10 photons is very poor compared to a typical light yield of 5 photons. The ToT of exactly one channel can be multiplexed to one differential pad pair, which is denoted as *digital monitor*. All digital monitors from all STiCs of one detector side will again be multiplexed to one DRS board to read out the waveforms. This feature could give a very useful handle for debugging, performance monitoring and reveals also time-dependent variations.




Figure 13.1: STiC threshold scan of SiPM Hamamatsu S12571-050P, (active area: $1 \times 1 \text{ mm}^2$, pixel size: $50 \times 50 \mu\text{m}^2$) dark counts.

PLL injection

The STiC chip uses a 16-stage voltage controlled ring oscillator to generate its time base. The uniformity of the timestamp distribution of the (random) events gives one handle to check that the Phase Locked Loop (PLL) is in a proper working state. Triggering the TDC directly with a fixed period adds a second complementary handle for STiC performance monitoring. A continuous triggered channel monitors the proper locking of the PLL over the full run time. The possibility to configure the chip to trigger each channel with this pulse allows further debugging, even of channels attached to the SiPM arrays. This direct PLL injection uses the differential *digital monitor* pad pair.




Figure 13.2: Delta time ($T_1 - T_2$) in one round fibre of a ribbon (without titanium dioxide in glue) read out with single SiPMs and STiC3.1 triggered on 0.5 photons.

13.2.3 READOUT BOARD

Figure 13.3 shows a possible integration of the Fibre Detector readout inside the experiment. The fibre ribbons are connected to SiPM arrays (LHCb-style) hosted on support PCBs. These boards are connected with 1 or 2 flexprints to the 24 readout boards², alternately placed on the top and bottom of the cooling and support plate.

The board hosts the readout electronics and on one side the analog connections to the SiPM arrays and on the other side the digital interface to the experiment DAQ system. In the baseline design there is one readout board per side per ribbon, in total 24 boards per side. The SiPM

²4 boards per module, 6 modules, 24 boards per side.




Figure 13.3: Possible integration of the Fibre Detector readout inside the experiment. This design accounts for the available $25 \times 25 \text{ mm}^2$ STiC3.1 MCM.

arrays are attached via a very short kapton flexprint. Since a single ended readout scheme is used, one single line per channel is needed besides the common HV distribution. The board hosts 2 MuSTiCs per ribbon (per side), provides an interface between the 125 MHz system clock and the readout electronics, provides the appropriate HV supply of the SiPMs, hosts various slow control sensors (listed below) and possibly LEDs used for calibration.

The clock distribution on this board has to provide three different clocks:

1. The reference clock for MuSTiC. The STiC PLL is designed for a reference clock of $625 \text{ MHz} \pm 50 \text{ MHz}$. The used frequency determines the mean bin width by $w_{\text{bin}} = \frac{1}{32 \text{clk}_{\text{ref}}}$. Since the pixel timestamps run at 125 MHz a 5 times faster clock is well suited as STiC reference. The use of the same base clock allows an easy integration of the 20bit STiC timestamps into the experiment's readout frame structure.
2. A reference clock for the fast LVDS links. MuSTiC uses double edge detection to generate the fast clock for the 1.25 Gbit/s LVDS links. In principal it is possible to use the same clock for the PLL and the LVDS links.
3. A configurable clock for PLL injection in the order of 100 kHz.

Slow Control

Beside a SPI bus to configure STiC chips, clock distributions and multiplexers the following slow

control measurements are provided by the readout board:

- temperature at STiC
- temperature at SiPM arrays
- common SiPM array voltage
- common (summed) SiPM array current
- STiC current/power consumption

STiC Interfaces

STiC3.1 has the following connects:

- 1 pair LVDS for data
- 1 pair LVDS 160 MHz clock for digital part
- 1 pair LVDS625 MHz clock for PLLs
- 5 SPI (rst, clk, cs, do, di at 10 MHz)
- 1 pair *digital monitor*
- 1.8 V digital
- 1.8 V analog
- 3.3 V

Interfaces

The following interfaces to the experiment DAQ, slow control and power distribution are needed:

information	lines	rate [Mbit/s]
SPI	3	10
SPI select <i>STiC</i>	2	10
SPI select <i>clk dist</i>	1	10
LVDS data links	2 diff	1250
reference clock	1 diff	125
reset	1 diff	125
slow control		
power	large	-
ground	large	-

13.3 Alternatives

The STiC ASIC is developed for PET applications where usually scintillators of around 100 mm^3 give around 1000 photons. The Mu3e Tile Detector signals are very similar to this. Since the Fibre Detector provides much smaller signals (~ 10 photons), it is somehow operated outside its design goals. In particular the missing energy information for small signals is an issue. Therefore, possible alternative readout electronics are discussed in the following.



13.3.1 DRS

In the research proposal the DRS chip developed at PSI was mentioned as a possible readout option for the Fibre Detector. The DRS4 uses a switched capacitor array to digitize waveforms up to 5 GSamples/s with a resolution of 12bit [42]. It was heavily used in the MEG experiment for many years where it provided a timing accuracy in the order of 40 ps [50]. The physics properties (such as time, energy and maybe pile-up suppression) would have been extracted from the waveforms already on a FPGA level. The limitation in event rate of about 100 kHz prevents it from use. Its successor, the DRS5 chip, which is currently under development will use an internal analogue memory (FIFO) to work in a dead-time less fashion up to an event rate of about 5 MHz. The chip's dimensions and power consumption obstruct the possibility to have the readout electronics right inside the detector. A moderate pre-amplification and cabling to the outside would be needed. Due to the very tight space constraints of the whole experiment this solution is not feasible.

13.3.2 OTHERS

Table 13.3 lists possible alternative ASICs and their drawbacks.



chip	description	problems
PETA [51, 52]	same TDC as STiC, focus on PET applications	in developing state, focus on hybrid concept with squared SiPM arrays
KLauS [53]	charge integration, would provide excellent single photon spectrum, time resolution of a few 100 ps probably possible	still in development, timeline \sim 2 years behind STiC
TOFASIC [54, 55]	very similar to STiC, maybe possible to get energy information, much less settings possible, rate limited, (next generation in development)	commercial product, PETsys sells whole solutions no bare ASICs, expensive
TRIROC [56]		rate
BASIC [57] FLEXTOT [58] PETRIROC [59, 60] NINO [61, 62] CITIROC [60]	only front end	very compact custom digitization solution needed, large effort
VATA citeVATA SPIROC [63]	developed for energy	time resolution
any preamp and TDC/ADC		custom solution in tight space needs a lot of development efforts & man power

Table 13.3: Overview of state of the art ASICs for SiPM readout.



FIBRE MECHANICS AND COOLING

14.1 Cooling

The Fibre Detector read out is located right below the outlets of helium for the local flow of the outer pixel layers and inside the global helium flow. It is crucial not to heat up the local helium too much since it is used to cool the pixel detectors. Therefore the Fibre Detector cooling needs enough cooling power to dissipate its own heat.

Independent of the read-out concept the power dissipated by the HV generation and distribution has to be cooled. If MuSTiCs are used ~ 30 mW per channel accumulate. Assuming the 64 channel per ribbon concept, this sums up to 7.7 W per module, ~ 50 W per side.

A water flow inside the massive solid base-plate, on top of which the readout board is located, provides the needed cooling power.

14.2 Mechanics

The drawings 14.1 and 14.2 illustrate the space requirements set by the experiment.

14.2.1 EXAMPLE SOLUTION

In the following, a possible mechanical integration of the Fibre Detector and its readout inside the experiment is shown. This solution is to be seen as the snapshot of the ongoing R&D.

Figure 14.3 shows a possible fibre readout support plate, figures 14.4 one fibre module inside the




Figure 14.3: Drawing of the experiment support plate, corresponding to one module.

detector. In the front view (a) and top view (c) a possible staggering of the fibre ribbons is shown. This staggering in r - as well as in z -direction provides enough space to place the (LHCb-type) SiPM arrays and possibly some ribbon holding structure. Since the SiPM arrays have to be slightly larger than the ribbons, staggering the ribbons in both directions is essential.

The exemplary fibre ribbon holding structure, as shown in figure 14.5, shows the feasibility with respect to the tight space constraints. A 0.5 mm thick and less than 3 mm long structure is directly glued to the ribbon. This is very similar to the holding structure used for the round fibre ribbon prototypes. Springs are used to press the fibres




Figure 14.1: Drawing of support structure and beam pipe to illustrate the available space for the fibre readout.




Figure 14.2: Drawing of the experiment support ring on which the fibre modules are mounted.

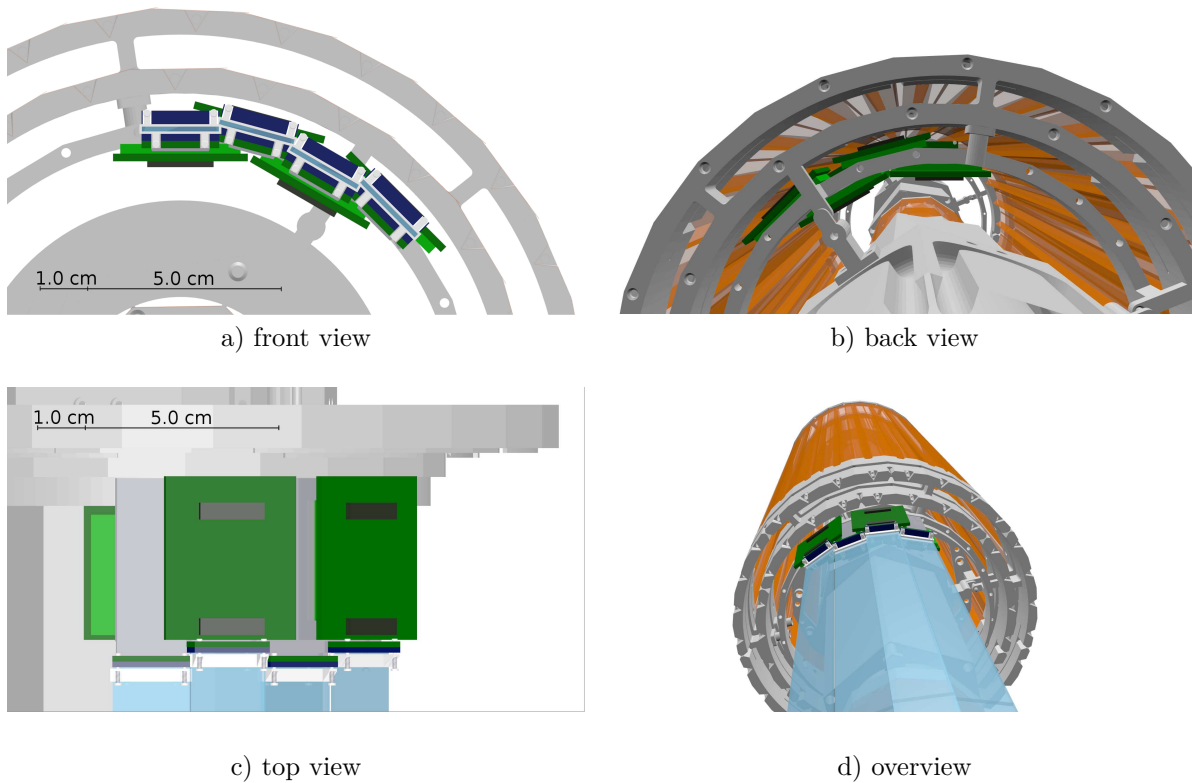


Figure 14.4: Renderings of the Fibre Detector. One fibre module inside the detector. Fibres in light blue, (LHCb-type) SiPM arrays in dark blue, green RO boards on gray cooling and support structure.




Figure 14.5: Possible solution of a fibre ribbon holding structure directly glued on the ribbons. Springs press the fibre ends against readout SiPMs.

against the SiPM array surface. If the fibres are directly glued to the SiPM arrays such a structure may be unnecessary.

The currently available version of STiC is bonded on a BGA¹-board of $25 \times 25 \text{ cm}^2$ and supports differential inputs. These dimensions are used for the presented drawings (light green in figures 14.4). The MuSTiC has only 32 instead of 64 channels, but no bond pads on one side. This allows to place two $5 \times 5 \text{ mm}^2$ chips right next to each other. Since the SiPM arrays are read out single ended, only half the signal bonds are needed. This allows to have again 64 channels on a BGA-board of the same or smaller size.

14.3 SiPM array

Sandro: anything? remove this?

Figure 14.6 shows the schematic of a Hamamatsu S10943 SiPM array.

¹ball grid array



(P 3 / 1 1)




Figure 14.6: Layout of the currently used Hamamatsu LHCb SiPM array in setup with round fibre ribbons.

PART V



TILE DETECTOR



MILESTONES FOR THE TILE DETECTOR

1. Construction of a *technical prototype* consisting out of a ring of 7 submodules with MuSTiC or STiC 3.1 readout: Test of SMD SiPMs, BC418 scintillator, readout PCBs, mechanical support structures and cooling
2. Development and verification of a tile manufacturing strategy and quality assurance scheme
3. Proof of SiPM radiation hardness
4. Development and testing of MuSTiC chip



TILE DETECTOR

The Tile Detector is a hodoscope, which aims at providing precise timing information of the particle tracks. The system is planned to be installed in phase Ib of the experiment, where due to the increased muon rate of about 100 MHz, the precise timing information is needed to suppress accidental background. As shown in figure 3.5, there are two identical Tile Detector segments, which are located in the two re-curl stations inside of the pixel layers; thus complementing the Sci-Fi tracker. The detector consists out of plastic scintillator, which is segmented into small tiles. Each tile is read out with a Silicon Photomultiplier (SiPM), which is directly attached to the scintillator. The main goal of the Tile Detector is to achieve a time resolution of better than 100 ps and a detection efficiency close to 100%, in order to efficiently identify coincident signals of electron triplets and suppress accidental background.

To a large extent, the following sections represent a summary of [64].

15.1 Detector Design

This section presents the baseline concept of the technical design of the Tile Detector. The design is based upon the results of the optimisation studies presented in the next section. Moreover, the design has to respect the physical space requirements in the experiment, as well as the requirements of the readout electronics.

The Tile Detector in the experiment phase Ib is subdivided into two identical segments - one in each re-curl station. In phase II, two additional

segments will be added. Each Tile Detector segment has the shape of a hollow cylinder, which encloses the beam-pipe. The length of a segment is 36 cm in beam direction (z direction), thus covering the same length as the pixel layers. The outer radius is 6.3 cm, which is limited by the surrounding pixel sensor layers. The detector in each re-curl station is segmented into 60 tiles in z direction and 56 tiles along the azimuthal angle (ϕ direction). This is the highest channel density which seems feasible, considering the space requirements for the readout electronics. The high granularity is essential, in order to achieve a low occupancy, as well as a high time resolution. This aspect is discussed in more detail in section 15.2.

The technical design of the Tile Detector is based on a modular concept, i.e. the detector is built up out of small independent detector units. The base unit of the Tile Detector, referred to as *Submodule*, is shown in figure 15.1. It consists out of 32 channels arranged in two 4×4 arrays. The tiles are made out of BC418 plastic scintillator and have a size of $6.5 \times 6.0 \times 5.0 \text{ mm}^3$. The choice of the scintillator material and tile geometry are the result of the simulation studies presented in chapter 16. The edges of the two outer rows of an array are bevelled by 25.7° , which allows for seven base units to be arranged in a circle. This seven-fold symmetry matches the heptagonal structure of the beam-pipe (see figure 15.4).

The individual tiles are coated with reflective TiO_2 paint, in order to increase the light yield and optically isolate the channels. Every tile is read out by a $3 \times 3 \text{ mm}^2$ MPPC with 3600 pixels, which is glued to the bottom $6.5 \times 6.0 \text{ mm}^2$ side




Figure 15.1: Submodule (32 channels).




Figure 15.3: Full detector (exploded view).




Figure 15.2: Module (480 channels).




Figure 15.4: Full detector (front view).

of the tile. At this interface, the scintillator is not painted. The SiPMs are soldered to a printed circuit board (PCB), which is connected via a flexible PCB (flex print) to a readout chip.

A stack of 15 Submodules constitute a *Module*, which comprises 480 channels. The drawing of such a Module is shown in figure 15.2. The Submodules are mounted on a metal support structure, which is also used for the heat dissipation for the readout chips and the cooling of the SiPMs.

15.3 and 15.4 show a drawing of a full Tile Detector segment, which consists out of seven Modules. The individual Modules are connected to a PCB mounted on the beam-pipe. This PCB hosts several FPGAs, which collect the data from the readout chips. The subsequent data transmission and processing is discussed in chapter 24.

15.2 Simulation

The Tile Detector design and response characteristics have been studied and optimised using a comprehensive simulation framework. This framework includes a simulation of the SiPM response [65], a combined scintillator-SiPM simulation and the full detector simulation described in chapter 27. The simulation has been used to study the detector performance in terms of time resolution and efficiency and optimise the tile geometry, as well as the scintillator material and SiPM type.

15.3 16-Channel Prototype

A 16-channel prototype of a Tile Detector Submodule utilising the STiC2 readout chip has been developed and tested. The time resolution and detection efficiency ε_d of this prototype has been




Figure 15.5: ...




Figure 15.8: MPPCs attached to the tiles via optical grease.




Figure 15.6: ...




Figure 15.7: 4×4 scintillator tile array.




Figure 15.9: Test-beam setup of the prototype. The sensor array is connected to a STiC read-out chip via a flex circuit board. The tile array is orientated parallel to the beam, such that the electron beam traverses four tiles in a row.




Figure 15.10: Number of hits in the 4×4 tile matrix. The beam traverses the tile array from the right to the left and is centered around the second tile row. The channels 4, 11 and 12 are not working due to bad electrical connections.

measured with an electron beam at the DESY test-beam facility. A picture of the setup is shown in figure 15.9.

The prototype consists out of an array of 4×4 scintillator tiles (BC408), which is shown in figure 15.7. The tile array is oriented parallel to the beam, such that the incident particles traverse four tiles in a row. Each tile has a size of $7.5 \times 8.5 \times 5.0 \text{ mm}^3$. These tile dimensions correspond to an earlier version of the detector design with $N_z = N_\phi = 48$ tiles arranged in a dodecagon¹ structure, instead of the current tetradecagon² geometry with $N_z = 60$ and $N_\phi = 56$. A diamond milling cutter has been used for the machining of the tiles, yielding a high surface quality. The tiles are coated on five sides with reflective TiO_2 paint³. The scintillation light is read out with a $3 \times 3 \text{ mm}^2$ MPPC (S10362-33-050C) with a pixel size of $50 \mu\text{m}$, which is coupled to the uncoated $7.5 \times 8.5 \text{ mm}^2$ side via optical grease. Figure 15.8 shows the MPPCs attached to the tiles. The MPPCs are connected to a STiC2 readout chip via a flex-rigid circuit board (see figure 15.9). The STiC chip setting, e.g. the individual thresholds, hysteresis and TDC parameters, have been tuned manually. Due to the large amount of parameters, a systematic optimisation of the chip settings could no be carried out. It is expected, that with fully optimised chip settings, the performance can be further enhanced.

Patrick: Say once what CC bins are




Figure 15.11: Energy deposition in a scintillator tile. The spectrum is composed of the Landau peak (blue), a plateau arising from edge effects (green) and a peak from optical cross-talk (red).

Figure 15.10 shows the number of recorded hits for the 16 tiles, as well as the numbering scheme which is used for the different channels. Only hits with an energy larger than $ToT > 50 \text{ CC Bins}$ are shown, in order to reject cross-talk events, which are discussed in the next paragraph. The beam is centered around the second row of tiles, which shows the highest hit rate, and traverses from the left to the right. The channels 4, 11 and 12 are not working due to faulty electrical connections. This problem can easily be solved in future prototypes by a more careful assembly of the circuit board.

Figure 15.11 shows a typical ToT spectrum, exemplary for channel 2. The spectrum shows several distinct features. The most prominent feature is the peak at around $ToT = 130 \text{ CC Bins}$,

¹12-fold symmetry.

²14-fold symmetry.

³EJ-510.



in the following referred to as *Landau peak*. This peak originates from electrons which fully traverse the tile. Such events can be selected by requiring the neighbouring tiles in the row (channels 1 & 3) to also have a signal in the Landau peak. The energy spectrum with this selection is shown by the blue curve in figure 15.11. The shape of the peak is given by a convolution of a Landau distribution, describing the energy deposition in the tile, and a Gaussian distribution, describing the energy resolution of the detector.

The second peak at $ToT \approx 30$ CC Bins originates from cross-talk between neighbouring scintillator tiles. This can be shown by selecting hits where at least one direct neighbour in the rows above or below the tile (channels 5, 6, 7) has a large signal with an energy deposition in the Landau peak. The corresponding events are shown by the red curve in figure 15.11. There is no correlation between these signals and the positioning of the corresponding signal lines, which excludes electrical cross-talk via capacitive coupling as a possible origin. Therefore, this effect is attributed to optical cross-talk, which can be explained by the fact that the tile surface facing the SiPM is not painted, allowing the scintillation light to diffuse to the neighbouring tile. This can easily be prevented in future prototypes by coating the tile surface around the SiPM.

The plateau region to the left of the Landau peak is attributed to edge effects, where the electron enters or exits the tile at the side, indicating that the detector array is not perfectly aligned with the beam. These events can be roughly selected by requiring at least one of the neighbouring tiles in the row to have no signal. In addition, it is required, that no tiles in the neighbouring rows are active, in order to suppress optical cross-talk events. The corresponding energy spectrum of these events is shown by the green curve in figure 15.11.

15.3.1 DETECTION EFFICIENCY

The detection efficiency was determined using particles which traverse all four tiles of a certain row. Such events are selected by requiring at least three hits in the row, with one hit in the first and last tile. The detection efficiency ε_d is then given by the probability to detect a hit in the remaining channel of the row with an energy deposition above the cross-talk level.




Figure 15.12




Figure 15.13

Due to the large light yield, which guarantees the signal to be well above the detection threshold, the efficiency is expected to be $\varepsilon_d \approx 100\%$. In the measurement, an efficiency of up to $\varepsilon_d = 99.7\%$ is achieved. The remaining inefficiency of 0.3% can presumably be attributed to the STiC2 chip and might originate from non-optimal chip settings. For the chip version STiC3 it has been shown, that in principle an efficiency of $\varepsilon_d = 100\%$ can be achieved [66].

15.3.2 TIME RESOLUTION

The coincidence time resolution was measured for different pairs of tiles within one row. The signal in both channels is required to be in the Landau peak, in order to assure that the particle fully traverses both tiles. Only events are selected where the fine counter information can be recovered for both hits.



Figure 15.12 shows a measured distribution of the timestamp difference δt , exemplary for channel 5 and 8. The width of the distribution is about $\sigma_{\delta t} = 100$ ps. Assuming the response of both channels is identical, the time resolution of a single tile is given by: $\sigma_t = 100 \text{ ps}/\sqrt{2} \approx 70$ ps. The time resolution is improved by applying a two dimensional time-walk correction, which is determined from the energy dependence of the measured time difference $\delta t(E_1, E_2) = t_1(E_1) - t_2(E_2)$. For the time-walk correction, the data set is divided into two independent subsets. One subset is used to obtain the correction function and the other subset is used to determine the time resolution applying the correction. The distributions of the corrected time difference is shown by the blue curve in figure 15.12. Depending on the channel combination, the time-walk correction yields an improvement between 13 % and 23 %, which is consistent with simulation results.

Figure 15.13 shows the time resolution with time-walk correction measured for various combinations of channel pairs, including directly neighbouring tiles (bin 1 to 9) and pairs with one (bin 10 to 14) and two (bin 15 & 16) tiles in between. The average time resolution over all combinations is $\bar{\sigma}_t = 66$ ps; the best achieved resolution is $\sigma_t = 56$ ps. The channel-to-channel variations can be explained by different chip settings and variations in the SiPM over voltage. In addition, the quality of the optical coupling between the scintillator and the SiPM is expected to vary for the different channels, due to a non-optimal assembly procedure. For future prototypes, the stability of the coupling between the scintillator

and the SiPM can be improved by using optical cement instead of silicon grease.

Simulation Results

The measurement of the time resolution has been simulated with the framework described in section 15.2. The simulation results should be interpreted as a rough estimate, since the response characteristics of the used setup in terms of noise and frequency response are not precisely known. Furthermore, the over-voltage of the individual SiPMs, which is influenced by the STiC settings, could not be determined during the measurement. Therefore, the SiPM parameters are approximated by the typical values for these devices. For the modelling of the bandwidth, a cutoff frequency of 500(100) MHz is assumed, which is the design value for the STiC chip. The electronic noise at the analogue input of the STiC chip is assumed to be 0.5(25) pe, and the timing threshold is estimated to be roughly 3(2) pe.

The simulated timestamps are randomised according to a Gaussian distribution, in order to take into account the intrinsic jitter of the STiC chip, which is assumed to be $\sigma_{stic} = 30$ ps. A time-walk correction is applied to the simulated data in exactly the same way as for the measured data. The time resolution obtained with this simulation setup is $\sigma_t^{sim} = 58(15)$ ps, which is shown by the green curve in figure 15.13. Within the uncertainties, which arise from the imprecise input parameters, the simulated resolution is consistent with the measurement; thus supporting the validity of the simulation model.



TILE SCINTILLATOR

16.1 Material

An essential component of the Tile Detector is the scintillator material. Achieving the best time resolution requires a scintillator with a fast response time and high light yield. In order to find the best suited scintillator for the Tile Detector, several scintillator materials have been compared using the simulation framework described in section 15.2.

Figures 16.1 and 16.2 show the simulated signal waveforms and time resolution of the scintillator-SiPM system for different scintillator materials. The values for the time resolution do not include the jitter of the readout electronics. Typical values for the MPPC parameters, analogue bandwidth and electronic noise are assumed in the simulation. Five different commonly used plastic scintillator materials, as well as an inorganic scintillator (LYSO) and a Cherenkov detector material (Sapphire) are compared.

The Cherenkov detector inherently has the fastest response. However, the light yield of about 20 detected photons is relatively low, which results in a poor time resolution. In contrast, the inorganic scintillator provides a high light yield, corresponding to a signal amplitude of about 1500 detected photons. However, the relatively slow response time of several tens of nanoseconds significantly limits the time resolution. Furthermore, the slow response leads to a higher pileup rate, which reduces the signal efficiency. Plastic scintillators offer both, a relatively high light yield, corresponding to about 900 detected photons, and

a fast response. This results in an excellent time resolution for the application in the Mu3e Tile Detector. The differences between the studied plastic scintillators are rather small.

The best resolution is achieved with BC418, which has both a high light yield and fast response time and therefore is chosen as the baseline material for the Tile Detector. This scintillator has a nominal light output of about 10 200 photons per MeV, a rise time of 0.5 ns and a decay time of 1.4 ns. The emission spectrum of the scintillator peaks at a wavelength of 391 nm, which roughly matches the maximum spectral sensitivity of the SiPM. This allows to directly read out the scintillation light without the need of an additional wavelength shifter.




Figure 16.1: Simulated signal waveforms for different scintillator materials.




Figure 16.2: Simulated time resolution of the scintillator-SiPM system (not including contributions from the readout electronics) for different scintillator materials. Typical values for the MPPC parameters, analog bandwidth and electronic noise are assumed.

16.2 Radiation Hardness

Given the maximum electron flux of about 50 kHz per tile in phase 1b (see figure 15.6), the integrated electron flux in one year of operation is about $10^{10} \text{ e}^+/\text{mm}^2$. The resulting total radiation dose is about 300 Gy or 30 krad. It has been shown in [67], that the scintillation yield of BC408 decreases by about 5% for a radiation dose of 300 krad, using a 2.5 MeV electron beam. Since the radiation level in the Mu3e experiment in phase 1b is smaller by one order of magnitude, the degradation of the scintillation yield caused by radiation damage in the Tile Detector is expected to be negligible.

16.3 Manufacturing

The scintillator tile manufacturing process involves three main steps. The first step is to mill the tiles out of the scintillator cast sheets and machine the tile surface to achieve the desired surface quality. In the second step, the tiles have to be coated with reflective paint, in order to optically isolate the tiles. An alternative approach is the wrapping of the tiles with reflective foil, which is expected to yield a higher light output. However, the wrapping process is more complicated and a wrapping procedure suitable for large scale production has yet to be developed. In the third step, the individual tile sub-modules are constructed by gluing together an array of 8×4 tiles. The detailed manufacturing procedure for the three steps is currently being developed.

The most time-consuming step is the milling of the tiles and the surface finish. The actual requirements for the surface quality are currently investigated. The baseline for the surface finish is to achieve a high surface smoothness, which assures a high optical reflectivity. The high smoothness can be achieved by using a diamond milling cutter, or using a standard milling cutter with subsequent extensive polishing of the surface. Both approaches are relatively time-consuming. The production of tiles with a more rough surface finish is less complex and would significantly simplify the production process. The impact of such a rough surface on the time resolution is currently being studied.



TILE SILICON PHOTOMULTIPLI- ERS




Figure 17.1: Hamamatsu MPPC S13360-3050PE

Silicon Photomultipliers are well suited for the application in the Mu3e Tile Detector, due to the compact dimensions, insensitivity to magnetic fields, high photon detection efficiency (PDE) and intrinsically fast time response. The most important SiPM characteristics for this application are a high PDE and gain, in order to obtain a large signal amplitude, which is essential for a good time resolution. Besides the good time resolution, a short signal decay time is desirable, in order to minimise pileup of consecutive signals, which reduces the signal efficiency. Usually, there is a

tradeoff between a large signal amplitude and a short decay time, since both parameters scale with the pixel and sensor size. The dark-rate of the SiPM is desired to be as low as possible, in order to achieve the best time resolution and avoid noise triggering.

17.1 SiPM Choice

Different SiPM types have been compared in simulation studies using the framework described in section 15.2, in order to find the best suited device for the Tile Detector. One of the main distinguishing parameters is the size of the sensor. Commonly available sizes are $1 \times 1 \text{ mm}^2$, $3 \times 3 \text{ mm}^2$ and $6 \times 6 \text{ mm}^2$. Figure 17.2 shows the simulation result for the number of photons hitting the sensor, as well as the time resolution of the SiPM-scintillator system. The simulation of the time resolution is simplified by assuming identical SiPM properties for the different sensor sizes. An exception is the dark-rate parameter, which is scaled proportional to the sensor area.

The simulated time resolution improves with the SiPM size, due to the increasing number of detected photons. However, it should be noted, that the rise time of the SiPM signal is expected to decrease for larger sensors, due to the increasing detector capacity. This effect is not taken into account in this simplified study and is expected to slightly degrade the time resolution for larger sensor sizes. The increasing detector capacitance also




Figure 17.2: Time resolution and number of photons hitting the SiPM as a function of the SiPM area. Commonly available SiPM sizes ($1 \times 1 \text{ mm}^2$, $3 \times 3 \text{ mm}^2$, $6 \times 6 \text{ mm}^2$) are marked with larger data points.

results in a significantly longer decay time of the SiPM signal. Consequently, there is a tradeoff between a large photon signal and a short dead-time connected with the sensor size. For the Tile Detector, a $3 \times 3 \text{ mm}^2$ SiPM is chosen. This sensor size yields a sufficient photon signal and time resolution, while exhibiting a moderate signal decay time.

Another main SiPM characteristic is the number of pixels of the device. Typical pixel sizes are $100 \mu\text{m}$, $50 \mu\text{m}$ and $25 \mu\text{m}$ ¹. For a device with a size of $3 \times 3 \text{ mm}^2$, this corresponds to 900, 3600 and 14 400 pixels, respectively. Assuming a typical PDE of 35%, the number of detected photons is expected to be about $3800 \cdot 0.35 \approx 1300$ (see figure 17.2). A 900 pixel device is hence not suited for this application, since the signal would strongly saturate at this light intensity. Smaller pixel sizes usually come along with a lower PDE, due to the reduced geometrical fill factor, which results in a degradation of the time resolution. Therefore, a SiPM with $50 \mu\text{m}$ pixel size is preferred over a device with $25 \mu\text{m}$. A drawback of the larger pixel size is the larger pixel capacitance, which results in a longer signal decay time.

Based on the simulation studies, a $3 \times 3 \text{ mm}^2$ SiPM with $50 \mu\text{m}$ pixel size is chosen as the base-

line photo-sensor. A respective SiPM from Hamamatsu (MPPC S10362-33-050C) has been success-

¹Recently, devices with even smaller pixels size down to $10 \mu\text{m}$ have become available. fully tested in a first Tile Detector prototype (see section 15.3). In the most recent SiPM series from Hamamatsu (MPPC S13360), a $25 \mu\text{m}$ pixel sensor is available which features a similar PDE as the $50 \mu\text{m}$ pixel device. This is a promising alternative to the baseline proposal of a MPPC with $50 \mu\text{m}$ pixel size. The performance of this device has to be verified in future detector prototypes.

17.2 SiPM Radiation Hardness

Ionising radiation can have a large impact on the SiPM characteristics and performance. The most prominent effect caused by irradiation is a strong increase in the SiPM dark-rate. Furthermore, there are several studies (e.g. [1]), which have observed a slight decrease in the detection efficiency after exposure of the SiPM to radiation. A possible explanation for this effect is the progressively larger amount of pixels in a permanent off-state [2]. Both, an increasing dark-rate and a reduced signal amplitude directly influence the time resolution of the sensor. The exact amount of signal degradation caused by radiation depends on the particle energy and type, as well as the specific SiPM device.

During the data taking period of phase Ib of the Mu3e experiment, the SiPMs will be exposed to a total radiation dose of about $10^{10} \text{ e}^+ / \text{mm}^2$. So far, no conclusive experimental data of the SiPM signal degradation is available for the given irradiation dose, particle type and energy. First studies of the radiation damage in SiPMs using a Sr90 source indicate, the degradation of time resolution during the Mu3e phase Ib is in the order of a few percent. However, more detailed studies are required in order to precisely predict the SiPM performance over the phase Ib runtime and obtain a comprehensive picture of the radiation effects.

17.3 SiPM Quality Assurance



TILE READOUT

Patrick: Describe tile RO cooling, cabling, mechanics

18.1 STiC Chip

The baseline solution for the readout of the SiPMs is the STiC¹ chip, which is a mixed-signal ASIC developed for the readout of SiPMs with the focus on Time-of-Flight applications [68]. A detailed description of the chip can be found in [48, 69, 70].

The core piece of the analog input stage is a fast discrimination unit with two tunable thresholds. The lower threshold, referred to as *timing threshold*, is used to determine the timestamp of the input signal via leading edge discrimination. The higher threshold, referred to as *energy threshold*, is used to determine the signal charge via the Time-over-Threshold (ToT) of the signal. The chip is designed in such a way that the ToT is directly proportional to the charge in a large range. The time and energy information is digitised via an integrated TDC with a bin size of 50 ps. The intrinsic time resolution of the chip has been measured to be about 30 ps [70]. In section 24.3.2, the details of the chip are discussed in depth.

An important feature of STiC is the capability to tune the voltage at each input terminal within a range of about 0.7 V. This allows to adjust the SiPM bias voltage and thereby compensate fluctuations in the detector response, which for example can arise due to temperature changes and channel-to-channel variations.

¹SiPM Timing Chip

The chip version STiC2 has successfully been tested in a 16 channel prototype of a Submodule, which is discussed in 15.3. The latest version of the chip (STiC3) features an additional zero pole cancellation circuit, which is designed to reduce the signal tail by a factor of three. This results in a significantly reduced channel dead-time. In addition, this is expected to reduce the baseline fluctuations related to the SiPM dark-rate and thus slightly improve the time resolution. Further development of the chip is planned, in order to optimise the performance for the application in the Tile Detector; including for example an increased data rate of about 1 MHz per channel, which is required for phase II.

18.2 DRS Sampling Chip

An alternative readout solution considered for phase II of the experiment is the DRS chip [42, 71], which digitises the waveform of the SiPM signals. An advantage of the DRS readout is the flexible analysis of the signals, which for example allows for a better pileup detection, compared to the STiC chip. However, the realisation of this readout scheme is quite challenging, in particular due to the high data rate which has to be processed. In section 24.3.2, the functionality of the DRS chip is discussed in more detail.

The chip version DRS4, which features a sampling rate of up to 5 GSPS, has been successfully tested with single scintillator tiles. The results of these tests are presented in section 15.3. However, the event rate for this chip is limited to roughly 100 kHz. Therefore, the successor chip




Figure 18.2: Alternative tile detector readout scheme. The Tile front end board (FEB) is divided into a PCB hosting the MuSTiC chip and two PCBs hosting the SiPMs. The SiPM boards are connected to the MuSTiC board via flex cables. The Tile FEB is placed on a mezzanine board, connecting the Tile FEB to the FPGA readout board.

DRS5 is currently being developed, and is designed to match the rate requirements of the Mu3e experiment. The new version is supposed to provide a dead-time less readout with an enhanced sampling rate of 10 GSPS.

In contrast to the STiC readout, the DRS chips would be located outside the detector in special crates, due to the significantly higher space requirements. The SiPM signal would have to be amplified and sent to the readout chips via a few meters of cable. The signal waveforms would then be analysed on FPGAs, in order to extract the timing and energy information.

18.3 Readout Scheme

18.3.1 MECHANICAL SUPPORT

18.3.2 CABLING

18.3.3 COOLING




Figure 18.1: Tile detector readout scheme. The Tile front end board (FEB) hosts 32 SiPMs and a MuSTiC chip for the processing of the analogue signals. The Tile FEB is mounted on a metal support structure with a cavity for liquid cooling of the ASIC and SiPMs. The Tile FEB is connected via a flex cable to a mezzanine board, which connects the Tile FEB to the FPGA readout PCBs mounted on the beampipe.

PART VI




SERVICES



MILESTONES FOR THE DETECTOR SERVICES

1. X
2. Y



COOLING




Figure 19.1: Schematic of the different helium gas flows of the cooling system.

The cooling system for the pixel detector must be capable of keeping the temperatures at a reasonable level ($< 70^\circ\text{C}$) and should at the same time add very little extra material to the active detector volume. Cooling with gaseous helium has been chosen since it offers a reasonable compromise between cooling performance and radiation length. The power consumption per surface area P/A of the pixel detector depends highly on the chip settings. $P/A = 224 \text{ mW/cm}^2$ was measured for a MuPix6 which was operated with settings optimized for low power consumption. Different settings that offered a compromise between low power consumption and good sensor performance lead to $P/A \approx 400 \text{ mW/cm}^2$. Table 19.1 summarizes the different total power consumptions P of the individual pixel detector components. The detector will be cooled by two different types of helium gas flows, a global helium flow that runs openly through the full inner volume of the magnet and a so-called local flow. For the latter the v-shaped folds in the module support structure serve as cooling channels. Figure 19.1 illustrates the helium gas flow inside the detector. Having the global and the local flow run in opposite di-

rections allows for cooling the detector from both sides simultaneously. The local flow is supplied by a distribution system which incorporates the detector station end rings and the module end pieces. Through the same distribution system, helium is also guided directly into the gap between layers 3 and 4 of each individual detector station and similarly into the gap between layers 1 and 2. After passing through a detector station, the warm helium has to be extracted. It is planned that all helium lines will be integrated in the beam pipe.

19.1 Experimental Helium Cooling Tests




Figure 19.2: Heatable layer 4 module prototype with temperature sensors

The feasibility of the above helium gas cooling system was studied in experimental tests and



DETECTOR PART	SIZE [cm ²]	POWER FOR 250 mW/cm ² [W]	POWER FOR 400 mW/cm ² [W]
layer 1	192	48	76.8
layer 2	240	60	96
vertex layers (1+2)	432	108	172.8
layer 3	1728	432	691.2
layer 4	2016	504	806.4
outer layers (3+4)	3744	936	1497.6
central detector (1-4)	4176	1044	1670.4
detector with 2 recurl stations	11664	2916	4665.6

Table 19.1: Absolute pixel detector power consumption for 250 mW/cm² and 400 mW/cm².




Figure 19.5: Experimental results for helium cooling of a layer 3 and a layer 4 module with only global flow and with additional local flow.




Figure 19.6: Comparison of simulation and experimental results of the cooling of a layer 3 and a layer 4 module




Figure 19.3: Heatable module prototypes placed inside the flow channel



Figure 19.4: 3D model used to simulate the helium cooling test setup

computational fluid dynamics (CFD) simulations¹ [72, 73]. The experimental cooling tests were performed inside a cylindrical flow channel with a diameter of 22 cm and a length of approximately 1 m that was placed in a helium environment of constant temperature. The tests were performed with detector module prototypes featuring a layer of aluminum Kapton laminate, which allowed for heating them ohmically. PT1000 temperature sensors were glued on top of the modules. To distribute the produced heat uniformly over the

¹Autodesk® Simulation CFD software was used.




Figure 19.7: Results of the beam pipe cooling simulations

modules, a heating pattern was cut into the aluminum layer with a laser. A photo of a module prototype is shown in figure 19.2. Two modules were placed inside the flow channel on top of each other, such that they resembled a segment of the detector layers 3 and 4 of one detector station. Figure 19.3 shows a photo of several module prototypes inside the flow channel. A global helium cooling flow was realized with fans at the ends of the cooling channel and an oppositely directed local flow was distributed via the module end pieces to the v-shaped folds in the module support structure. Figure 19.5 shows the improvement of the cooling performance due to the local flow as opposed to cooling with only global flow. The x -axis shows the position on the modules parallel to the beam axis. The y -axis shows the difference between measured temperature and the temperature of the incoming helium (20°C). For comparison, CFD simulations of the above cooling test setup were performed. The according 3D model is shown in figure 19.4. Both, the experimental and the simulated results are shown in figure 19.6. The helium flow velocity values $v_{global} = 2.5 \frac{\text{m}}{\text{s}}$ and $v_{local} = 20 \frac{\text{m}}{\text{s}}$ are not fully optimized but were found to be reasonable [72].

19.2 Water Cooling

In addition to the helium gas cooling, several parts of the detector will be water cooled. Water cooling channels will be integrated in the beam pipe, the support structure of the tile detector, (see chapter 15) and the end ring of the fiber tracker (see chapter 11) to cool the front-end FPGAs, the power regulators, and the STiC chips. CFD simulations of the water cooling of the beam pipe were performed. A 1 mm thick layer of silicon was added on top of the 3D model of the beampipe to simu-

late the front-end FPGAs. A heating power of 200 W was specified for each of the seven segments. Inside each beam pipe segment, the water runs once back and forth in a U-shaped channel. The flow velocity at each inlet was set to 1.36 m/s. This value equals the flow velocity in an according experimental cooling test, which was carried out with a prototype of one beampipe segment. During the experimental test the temperature was measured only near the turning point of the cooling channel. The result was $\Delta T = 2.9^\circ\text{C}$ [73]. The result of the simulations are shown in figure 19.7. The maximum temperature is $\Delta T_{max} \approx 7.4^\circ\text{C}$.

19.3 CFD Simulations of Phase I

CFD simulations of the phase Ia and phase Ib detector were performed. The according 3D models are shown in figures 19.8 and 19.9. The phase Ia model includes:

- the target
- all four pixel layers of the central detector
- the pixel end rings with integrated helium distribution system
- the beampipe (upstream and downstream) with integrated water cooling channels
- a 1mm thick layer of silicon on top of the beam pipe to resemble the front-end FPGAs
- pipes to guide helium to the end rings (which don't represent the final design)

To create a helium flow between layers 1 and 2, a helium channel is integrated in the beam pipe. The phase Ib model additionally includes two pixel recur stations and additional barrels underneath pixel layer 3, which resemble the timing detectors. Their thickness is $750 \mu\text{m}$ in the central detector and 5 mm in the recur stations. The boundary conditions were the following:

- $P/A = 250 \text{ mW/cm}^2$ and $P/A = 400 \text{ mW/cm}^2$ for the pixel layers
- $P = 200 \text{ W}$ per beam pipe segment for the front-end FPGAs
- Helium flow velocity $v_{he} = 100 \frac{\text{m}}{\text{s}}$ at each of the helium pipe inlets, which leads to a local flow velocity of $v_{local} \approx 16 \frac{\text{m}}{\text{s}}$ and a flow velocity between layers 3 and 4 of $v_{gap} \approx 3.5 \frac{\text{m}}{\text{s}}$
- Helium flow velocity between pixel layers 1 and 2 $v_{inner} = 4 \frac{\text{m}}{\text{s}}$
- Global helium flow velocity $v_{global} = 3.5 \frac{\text{m}}{\text{s}}$
- Water flow velocity at the inlets $v_{water} = 1.36 \frac{\text{m}}{\text{s}}$




Figure 19.8: 3D model used for the cooling simulation of the phase Ia detector




Figure 19.9: 3D model used for the cooling simulation of the phase Ib detector




Figure 19.12: Simulated temperature profiles of the phase Ib detector for $P/A = 250 \text{ mW/cm}^2$




Figure 19.11: Simulated temperature profiles of the phase Ia detector for $P/A = 400 \text{ mW/cm}^2$




Figure 19.13: Simulated temperature profiles of the phase Ib detector for $P/A = 400 \text{ mW/cm}^2$




Figure 19.10: Simulated temperature profiles of the phase Ia detector for $P/A = 250 \text{ mW/cm}^2$

- Inlet water and helium temperatures $T_{he} = T_{water} = 0^\circ\text{C}$

The results of the simulations are shown in figures ?? to 19.13. One can see that the current cooling scheme is sufficient assuming a power consumption per surface area of the pixel detector of $P/A = 250 \text{ mW/cm}^2$ but not for $P/A = 400 \text{ mW/cm}^2$.



19.4 Possible Design Optimizations

One design option to increase the cooling performance is to increase the cross-section area of the v-shaped folds in the module support structure, i.e. the local cooling channels. Figure 19.14 shows the results of a simulation of the pixel layers 3 and 4 of one detector station with increased size of the cooling channels (height of the triangle increased from 5.2mm to 6mm and the base from 6 mm to 7.5 mm) and for comparison the results for the current design.




Figure 19.14: Comparison of local cooling with current design and with increased cross-section area of the cooling channels

For layers 1 and 2, which do not feature v-shaped folds, increasing the helium flow inside the gap between the two layers should improve the cooling performance.

19.5 Stability against Flow-induced Vibrations

One factor that possibly limits the range of the applicable helium flow velocity is the occurrence

of flow-induced vibrations, which could significantly worsen the spatial resolution of the pixel tracker. Therefore, a setup has been built to measure vibration amplitudes in normal direction to the pixel layers and the frequency response of detector module prototypes exposed to helium flow via Michelson interferometry. Two Bachelor theses have been carried out in this context [74, 75]. A layer 3 module prototype was equipped with a mirror by sputtering aluminum on one of the glass plates before it was placed inside one of the two laser arms of the interferometer and local helium flow was applied. An algorithm was developed to determine vibration amplitudes from interferometry data samples, which were taken with a photodiode and an oscilloscope. For a local flow velocity of approximately $20 \frac{\text{m}}{\text{s}}$, the maximum observed vibration amplitude was below $10 \mu\text{m}$ and the average amplitude was below $2 \mu\text{m}$. At a flow velocity of about $30 \frac{\text{m}}{\text{s}}$, the maximum observed vibration amplitude was $39 \mu\text{m}$.

The frequency response was measured by placing a speaker in front of the module to excite the latter with a known frequency. Several resonance frequencies were observed within the range of 50 Hz to 1000 Hz, but the highest vibration amplitudes only occurred at very low frequencies ($< 100 \text{ Hz}$).

It is planned to improve the test setup by putting another beam splitter into the recombined laser beam and measure each of the two resulting beams with separate photo-diodes after creating a phase difference of $\frac{\pi}{2}$ between the two. This allows for determining not only the magnitude of the module displacement in the laser beam direction but also the sign, as discussed in [76]. This would significantly reduce the uncertainty in the amplitude measurement caused by the superposition of different oscillation modes (see [75]). It is also planned to include global helium flow in the tests.



MECHANICAL INTEGRATION

Dirk: Describe cage inside magnet and how everything attaches to cage

Dirk, Felix, Peter-Raymond: Describe procedure for accessing detector

The Mu3e detector, comprising pixel detector, fiber detector and tile detector together with the muon stopping target is located inside the warm bore of the Mu3e solenoid magnet. The detector has to be mounted in such a way, that electrons and positrons produced by muons decaying at rest on the target can freely move in the magnetic field in order to precisely determine their momentum. It is however important to fix the detectors firmly to achieve a good spatial resolution with the pixel detector. Access to the detector has to be provided for detector installation and for exchanging faulty modules.

To fulfill the above requirements a mechanical scheme has been developed in which all detectors are mounted to the up- and downstream beam-pipes. The beam pipes are supported by the detector cage, which consists of large end-wheels on either end of the Mu3e detector connected by rods running outside the active volume. This cage rests on rails on which it can be moved into the Mu3e magnet.

20.1 Support of Detector Stations

The Mu3e detector is subdivided into a central station, an upstream and a downstream station, see figure 20.1. Each station is barrel shaped (prismatic) and has several concentric detector layers.

All three stations have two outer pixel layers, the central station in addition consists of the target, the vertex pixel layers and the fiber tracker, the upstream and the downstream station have a tile detector layer. The mechanical design of all detector stations is based on detector modules mounted to end-wheels, see figure 20.2. Size and geometry of the detector modules are described in the corresponding sections 8, 14 and 15.1.

20.1.1 SUPPORT OF PIXEL VERTEX LAYERS

The vertex layers of the pixel detector and the target are mounted to two small end-wheels. These small end-wheels are directly fastened to the corresponding up- or downstream beam-pipe. As for all detector modules the vertex pixel detector modules are fixed to one end-wheel while they are mounted with sliders on the opposite side. One reason to fully fix the detector modules only on one side is the thermal expansion, the other reason is the gravitational sag of the two beam-pipes which are not supported in the middle.

20.1.2 SUPPORT OF THE CENTRAL STATION

The central station consists of the vertex layers mentioned above, the fiber tracker and the central outer pixel detector layers. The end-wheels supporting the outer pixel detector layers and the fiber detector are mounted on additional feet to the upstream and downstream beam-pipes. Three feet connect each end-wheel of the central station to the up- or down-stream beam pipe, so as for the vertex layers special care must be taken to accommodate the beam-pipe sag.




Figure 20.1: Screenshot of the Mu3e detector mechanics 3-D model. The Mu3e detector is subdivided into three stations.




Figure 20.2: Upstream and downstream beam-pipes and station end-wheels.



20.1.3 SUPPORT OF THE RECURL STATIONS

The recurl stations each have a barrel made of one tile detector layer and two outer pixel detector layers. The barrels are built by mounting tile and pixel detector modules to end-wheels, which are similar to the ones for the fiber detector and the central outer pixel layers. For the up- and down-stream barrels it is foreseen to slide them pre-assembled onto the support feet facing the target and then fix them to the feet accessible from the outside. The up-stream station end-wheels and the down-stream station end-wheels are mounted only to the corresponding beam-pipe. In addition to the detector modules readout PCBs are attached to the beam pipes.

20.2 Detector cage and rail system

A detector cage mechanically supports the up-stream and the downstream beam-pipes inside the solenoid magnet, see figure 20.3. It has four large end-wheels with the upstream and the downstream beam-pipes supported on the axis. On the outer rim these large end-wheels are connected via long rods outside the active volume of the Mu3e experiment to each other. If necessary the cage can be further re-enforced between these rods by adding sheets on the circumference. The cage slides on aluminum rails into the magnet warm bore. It is foreseen to have two double rails running on the left and on the right side of the detector cage. During data-taking the detector will be locked in the central position. The rail system is supported on bracers at both ends of the warm bore, so the rails are not touching the warm bore inner surface. The warm bore inner surface has little flatness and moves with changing operation conditions, so it cannot be used as a rail support. In addition to the rails spanning the inside of the warm bore, an extraction tool will continue the rails to the outside of the magnet. This extraction tool can be connected to the rails inside the magnet and help to shift the full detector cage inside or outside.

20.3 Access to the Mu3e Detector

The Mu3e detector must be installed into the solenoid magnet and access is required for detector debugging, repair or upgrading of detector components. To this end, the Mu3e detector is mounted into a large support cage which is inserted into the solenoid magnet on a rail system, as discussed above in section 20.2.

Before having access to the warm bore of the magnet, the magnet has to be moved from the measurement position to the detector installation position. In a first step the beam-line has to be decoupled from the solenoid, which implies decoupling the beam vacuum. Most likely some beam elements have to be moved to make space for the following operations. Once it is free, the Mu3e solenoid has to move and turn in order to create space in front of the down-stream magnet door. The movement of the 30 t magnet is alleviated by using air-cushions to lift the magnet. In a next step the two halves of the shielding iron on the downstream end of the solenoid must be lifted off using a local crane system. These shielding door-wings have cut-outs, so that the cables and pipes connecting the Mu3e detector on the downstream side with the outside world don't have to be removed beforehand. Now these cables and pipes connected on the downstream side have to be removed and the end-flange which is sealing the Mu3e solenoid on the downstream side has to be taken off.

Since the detector is fairly delicate a trolley will be designed to take the detector over. The trolley has a rail system fitting the one inside the solenoid magnet. After coupling the rail system on the trolley to the rails inside the magnet the detector cage will be unlocked either to be moved from the trolley into the magnet or the other way around. Special care must be taken to ensure the alignment of the rails on the trolley and the ones inside the magnet and to avoid gaps between the rail segments. When the movement is done the detector cage must be locked to the rails.

After the access the end-flange, the cables, the pipes and the iron shielding door-wings have to be put back in place. Then the magnet will be pushed back to the position for physics data taking, marked on the ground. Finally the beam elements must be put back in place.




Figure 20.3: Mu3e detector in cage and rail system (old versions).



POWER DISTRIBUTION AND CABLING

Dirk, Stefan: Describe power supplies, cabling and power delivery to subsystems

The power distribution of the Mu3e detector has to face some major challenges. Power has to be distributed to the MuPix chips inside the active volume of the experiment, which limits the use of material to a minimum in order to avoid multiple coulomb scattering. The power distributed to the front end ASICs MuPix and MuSTiC has to be extremely clean in order to achieve high analog performance and avoid extra noise or crosstalk. Further on the Mu3e experiment is located inside a strong superconducting magnet, which will quench potentially inducing large currents. Apart from the above constraints a robust and modular power distribution system should be implemented.

21.1 Power distribution to the front end ASICs

The power distribution to the front end ASICs MuPix and MuSTiC requires special care since these chips process delicate analog signals on the inputs and drive high speed electrical links at the outputs. In addition the MuPix chip is located in the active area, so in order to avoid extra multiple coulomb scattering external capacities cannot be placed near the MuPix and the supply lines have to be made from low z material, which de-

facto means aluminum. Within the active region the power distribution to the MuPix chips is integrated in the flex-prints also carrying the signals from and to the MuPix, see section 8.2. These flex-prints are connected to rigid PCBs at the fringe of the detection volume, a location where at least compact capacitors for the power supply line are located. The MuSTiC readout chip for the tile detector and the fiber detector is placed outside the active detection volume on a rigid PCB housing filter capacitors and fanning out the signals to the MuSTiC chips, see section 13.2.3.

From the rigid PCBs at the edge of the pixel detector modules, the scintillating fiber detector modules and from the tile detector modules it is planned to run flexible cables to the front-end PCBs. Depending on the outcome of current studies these flexible cables will be copper-Kapton flex-prints or stranded copper cables, which will carry the power lines and the signals. These cables are plugged to the front end boards attached to the metal upstream and downstream beam-pipes. Robust DC-DC converters are located on the front-end boards, generating the low voltages for the front end ASICs MuSTiC and MuPix (1.5 V, 1.8 V) and for the other integrated circuits inside the magnet from a supply voltage in the range of 10 to 20 V. The maximum output current of these DC-DC converters can be as high as 20 A, enough to supply the power for more than 20 MuPix chips of $2 \times 2 \text{ cm}^2$.



CLOCK DISTRIBUTION

Stefan: Describe generation and distribution of global clock and reset

Stefan, Simon, Nik, Patrick: Describe clock interfaces to sub-detectors

The clock distribution for the Mu3e experiment has to be well designed and implemented since the experimental concept is based on reducing accidental background with the help of precise time measurement. This requires an experiment-wide phase stability of the clock distribution of better than 100 ps. On the other hand the data acquisition system has a very high throughput using 1.25 Gbit/s LVDS links and multi Gbit/s optical links. The operation of these links requires very stable reference clock frequencies. It is also highly desirable to run all data links with the same reference clock, otherwise links contributing fragments to the same event would drift apart. The clock frequency for the clock distribution system is 125 MHz, other frequencies can be derived locally.

22.1 Clock Distribution scheme

Since a single clock should be used by all front end ASICs, all FPGAs on front-end boards on the detector, the switching FPGAs and the FPGAs in the farm PCs receiving fragments and assembling events, an experiment wide clock fan out must be provided. The clock will be generated by an oscillator on a pcb with 16 or more synchronous clock outputs. From there it will be distributed via optical fiber to the switching boards, which will likely be the PCIe40 PCB developed by the Marseille group for LHCb upgrade []. The clock signal is

distributed from the switching boards via optical fibers to the front-end boards. On each front-end board the clock signal is received and distributed to the connected front-end ASICs. Extra clock filter chips are foreseen to improve the quality of the clock in terms of jitter, typically the remaining jitter is below 5 ps [].

22.2 Reference clock generation

The reference clock for the Mu3e experiment is generated by a commercial-off-the-shelf clock oscillator. The frequency of the reference clock is 125 MHz, which is well suited for clock distribution systems. In order to ensure the ease of use, the reference oscillator should sit on a dedicated PCB with very stable power supply, which itself is in a rack compatible department like a 1U box. On this clock distribution board there are clock fan-out chips, in order to provide at least 16 clock outputs. Preferable the clock outputs are all in phase and output via optical fiber compatible with the SFP+ optical input on the PCIe40 board used as switching board.

22.3 Clock distribution through the switching boards

The switching board has an optical SFP+ input for the clock signals and up to 48 fast optical outputs towards the front-end boards. The total number of switching boards will be adapted to the number of optical data links coming from the front-end PCBs. Since by design there will be one



optical link towards the detector for every data link from the detector, this link to the front-end boards can be used for clock distribution. This has to be implemented in such a way, that the clock phase on the receiving side is un-ambiguously recovered.

22.4 Clock distribution on the front-end boards

Since the clock is distributed over an optical fiber, there will be a dedicated optical receiver on each front-end board. The clock has to be recovered with the correct phase, ensuring the synchronicity of Mu3e after each power cycle. The clock will be filtered and distributed to the front end ASICs by commercial clock multiplier chips. The clock multiplier chip chosen can be configured to shift the clock phase in tiny steps of 0.28 ps, to use a minimum delay of 100 ps or have the default delay of 2 ns. From these clock multiplier chips one LVDS pair of clock signals is sent to each array of front-end ASICs (MuSTiC or MuPIX).

22.5 Clock distribution to the ASICs

The clock is sent from the front-end board over LVDS pairs either to the MuPIX ASICs or the MuSTiC readout boards.

22.5.1 TO THE MuPIX

It is foreseen to supply one differential pair of clock signals to each flex-print. All MuPIX chips sharing one high density interconnect (flex-print) use the same clock pair, with the termination on the last chip.

22.5.2 TO THE (FIBER DETECTOR) MuSTiC

Similarly, one LVDS reference clock is routed to each of the Fiber Detector RO boards, where the same clock multiplier and filter ICs are used to generate all the clocks needed locally: a 625 MHz reference clock for the ASIC internal PLL, a reference clock for the fast LVDS links and MuSTiC's digital part and a configurable and much slower clock, in the order of 100 kHz, for PLL injection.



SLOW CONTROL

The slow control system deals with all “slow” data such as high voltages for the SiPMs and silicon sensors, ambient temperatures and pressures, the beam line magnet settings and parameters of the cooling system. The silicon pixel system will use a separate system described later in this chapter.

For the slow control parameters it is important to have all data and control functionality in a homogeneous single system. This makes the maintenance of the system much simpler, since only a limited number of different hardware standards have to be taken care of. The integration of all data enables us to define control loops between otherwise completely different subsystems. Examples are to switch off the muon beam line in case of some overheating of the pixel system, or the control of the cooling system based on data derived from signals from the SiPMs of the fibre and tile sub-detectors.

The integration of all systems will be done through the MIDAS DAQ system and its associated MIDAS Slow Control Bus (MSCB) system [77]. Both systems have been used in dozens of experiments in several laboratories over the past decade and have proven to be very flexible and reliable. The MSCB system uses a serial differential bus for communication, with two data lines (positive and negative polarity) and a common ground. Over long distances, such as between crates, the physical standard for this bus is RS-485, running at a relatively low speed of 115.2 kBaud in half-duplex mode. The slow speed makes this bus highly immune against improper termination or electrical interferences, while the short commands

of the MSCB protocol still allow the readout of many hundred nodes per second. Over short distances, such as between several FPGAs on a single PCB, the bus uses the LVDS standard. The FPGAs have a serial interface (UART¹) implemented to interface to this bus. The MSCB bus uses a single-master, multiple-slave architecture, where all slave nodes on the bus only have to reply on requests sent by the master node, thus making the bus arbitration very simple. Different bus adapters exist already to bridge between LVDS and RS485 segments as well as between RS485 and optical fibers. We plan to connect about ten to twenty FPGAs via LVDS lines to an individual bus segment, which then is connected to an optical fiber leaving the detector vessel. Each optical fiber will be interfaced to the Ethernet based control network of the experiment. A small microcontroller implements a simple UDP protocol, so that the MSCB bus segments can be accessed from any computer connected to the experiment Ethernet network.

The MSCB nodes inside the experiment will be either dedicated 8-bit microcontrollers or soft-core microcontrollers instantiated inside each FPGA. The micro controllers perform local control loops, such as high voltage stabilization, and send measured values to the central DAQ system for monitoring. Many devices already exist for this system, such as the SCS-2001 unit shown in figure 23.1. Since the system was developed at PSI, it can be quickly adapted to new hardware. The high voltage control for the SiPMs can for example be directly integrated into the carrier boards holding

¹universal asynchronous receiver/transmitter




Figure 23.1: SCS-2001 unit as part of the MSCB slow control system. This unit has 64 input/output channels, which can be configured via plug-in boards as digital or analog channels. Many plug-in boards exist already such as PT100 temperature sensor readout cards, analog high resolution inputs (24 bit resolution), valve control outputs and many more.

the SiPMs, thus eliminating the need for high voltage cables. The optimized protocol of the MSCB system allows the monitoring of many thousand channels with repetition rates in the 100 ms range, which will be more than enough for this experiment.

Since the MSCB protocol is very simple, using only a few bytes for addressing, data and redundancy, it can be implemented in less than 700 lines of C code. This makes it possible to run the MSCB core inside some soft-processor inside every FPGA used in the experiment.

Test implementations showed that this needs only a few percent of the available FPGA resources, which can be easily accommodated. Having a dedicated slow control link to all FPGAs in the experiment is a powerful tool for debugging and configuration, since this allows the management of the FPGAs in case the serial data links are down.

In addition to the MSCB nodes inside our detector, we will have dedicated slow control crates next to the detector and at the infrastructure platform using several SCS-2001 units. These crates have analog and digital in- and outputs to interface all kinds of sensors and actuators around the experiment. They are also connected via Ethernet to the experiment network.

In addition to the MSCB system, the MIDAS slow control package contains interfaces to the PSI beamline elements via the EPICS system [78]. This allows monitoring and control of the beamline from the main DAQ system, which has been proven very versatile in other experiments using this scheme.

All slow control data will be stored in the history system of the MIDAS system, so that long term stabilities of the experiment can be effectively verified. The slow control data is also fed into the main event data stream, so that any of-line analysis of the event data has this data available.

All data fed into the MIDAS system are accessible by the MIDAS distributed alarm system. This system allows to set upper or lower limits on all slow control data in a flexible way through the MIDAS web interface. In case of an alarm, shift crew can be notified through spoken alarm messages and contacted via mobile phones. Scripts can be triggered which put the whole experiment in a safe state in order to avoid damage such as from too high temperature.

The following list briefly shows the different sub-systems in the Mu3e experiment and corresponding parameters which will be monitored and/or controlled through the slow control system:

- Mu3e magnet
 - currents
 - temperatures (ambient and cryogenic)
 - quench system status
 - helium compressor status
 - field (Hall probes)
- Separator
 - high voltage
 - current
 - X-ray emission
- Beamline
 - magnet currents
 - slit settings
 - accelerator proton current
- SiPMs
 - temperatures
 - bias voltage
 - bias current
- Detector cooling system
 - temperatures
 - differential and absolute pressures
 - humidities
 - oxygen content

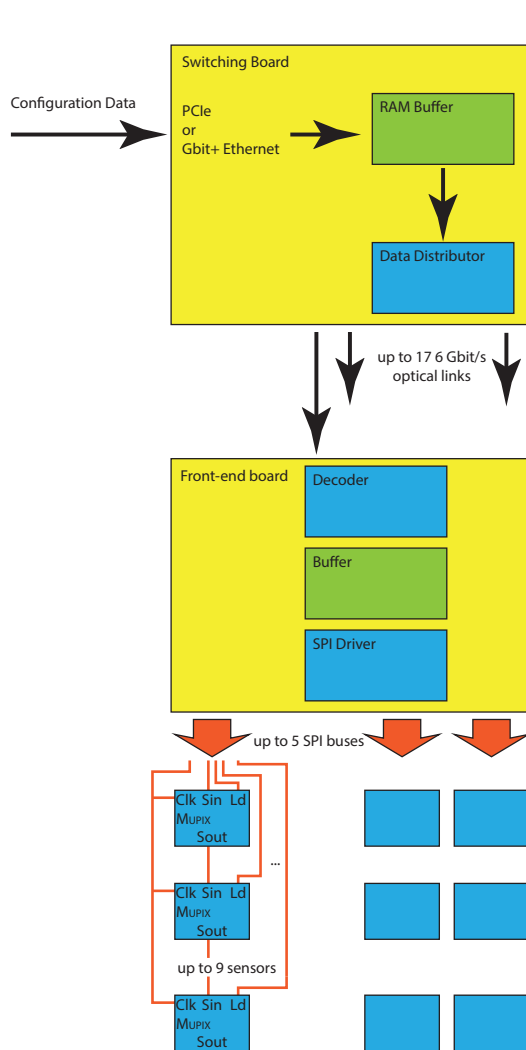


Figure 23.2: Data flow for the pixel configuration.

- ASICs (MuPIX and MuSTic)
 - temperatures
 - power consumption
- FPGAs
 - temperatures
 - power supply voltages
 - link status
 - PLL locks
- Computer farm
 - temperatures
 - humidity
- Pumps
 - temperatures

- speed
- power consumption
- Data acquisition
 - event sizes
 - event rates

23.1 Pixel Configuration and Control

A special case is the configuration of the pixel detectors, which require many million parameters, like the tune-DAC values for each pixel. Since the amount of data here is considerably larger than for all other systems (approximately 120 MB for the phase Ib detector), an extension of the slow control system is planned. A dedicated program manages, visualizes and exchanges the pixel detector configuration parameters between an optimized database and the pixel hardware. In this way the time required to configure the pixel detectors can be minimized, while this program is still connected to the main DAQ system. It can be synchronized with run starts and stops, and can inject pixel monitoring data periodically into the event data stream for offline analysis. The regular slow control data stream will contain a pointer to the relevant state of the pixel configuration database.

The configuration of the individual pixel sensors is written via a serial programming interface (SPI) with up to 9 sensors connected in series. This corresponds to approximately 20 Mbit of configuration data, which in turn dictates the need for SPI clock speeds above 10 MHz in order to guarantee fast run starts. Sensors on different ladders (different SPI buses) can and have to be programmed in parallel. As the on-chip memory of the front-end FPGA is too small to hold the complete configuration data for all connected sensors, it has to be delivered just in time from the switching boards. Here bandwidth is not an issue (we have a 6 Gbit/s optical downlink at our disposal), but the data stream has to be synchronized such that no buffering on the front-end is required, which necessitates a careful interplay between the software driving the data into the switching boards, the switching board firmware and the front-end firmware. An overview of the data flow for the pixel configuration is shown in figure 23.2.

PART VII



DATA ACQUISITION AND EVENT FILTER



MILESTONES FOR THE DATA ACQUISITION

1. Set up a slice of the readout chain using development boards, develop, and test the firmware.
2. Produce and test the front-end board prototype.
3. Develop, produce and test the final small front end board.
4. Acquire and test the switching board.
5. Run the full selection algorithm on a GPU.
6. Integrate the readout chain and the selection algorithms.
7. Integrate the farm PCs with the MIDAS DAQ system.
8. Scale readout system to full phase Ia capability.
9. Scale readout system to full phase Ib capability.



DATA ACQUISITION

24.1 Overview

The Mu3e data acquisition system works without a hardware trigger on a push basis, i.e. the detector elements continuously send hit information to the data acquisition (DAQ) system. The DAQ consists of three layers, namely front-end FPGAs, switching boards and the filter farm. The topology of interconnects is built such that every farm PC gets to see the complete detector information for a select time slice. See Figure 24.2 for an overview of the readout scheme and Figure 24.1 for the scheme at detector start-up (phase Ia). Complete overviews of all optical readout links can be found in the appendix on pages 250 and 251.

Hits in all subsystems are timestamped and the front-ends ensure that time ordered information is forwarded to the rest of the read-out system. On the interfaces to the farm PCs, data from several timestamps are merged to overlapping reconstruction frames, see figure 24.3.

24.2 Occupancy

The bandwidth requirements of the data acquisition are largely determined by the expected detector occupancy, as all the Mu3e sub-detectors produce zero-suppressed output.

Occupancies are obtained with the full simulation running at a muon stop rate of $1 \cdot 10^8$ Hz and pessimistically estimating the beam related background by losing another $0.9 \cdot 10^8$ Hz of muons along the last metre of beam line.

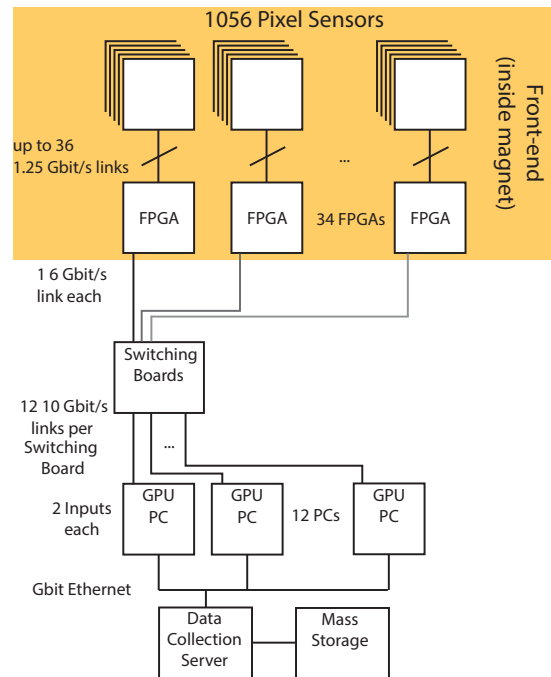


Figure 24.1: Mu3e readout scheme for the start-up detector.

The distribution of the occupancy over the pixel sensors is shown in Figures 24.4 and 24.5. The average occupancy of the fibre detector is shown in figure 24.6. Occupancy maps for the tile detector are shown in figure 24.7.

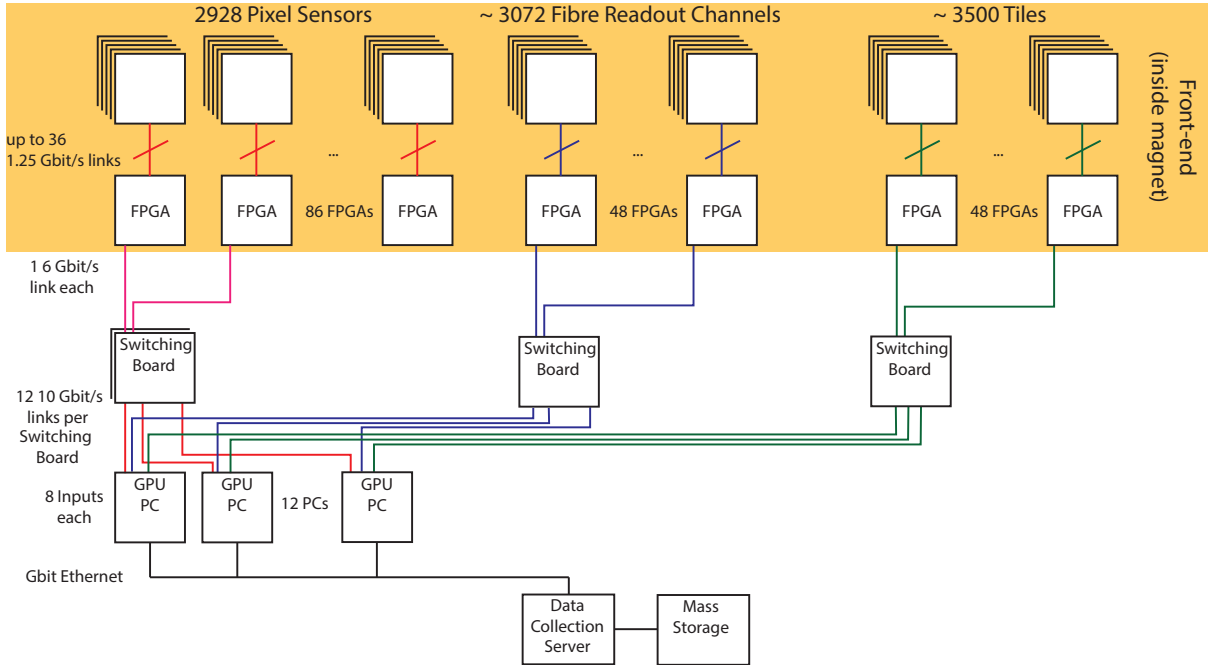


Figure 24.2: Overall Mu3e readout scheme

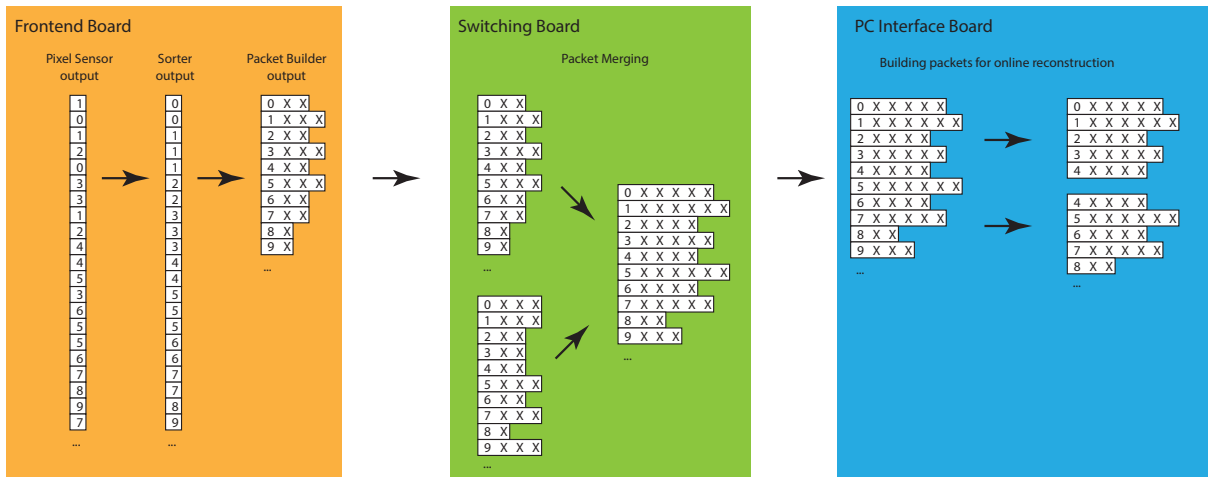


Figure 24.3: Schematic flow of pixel time information through the Mu3e readout system.




Figure 24.4: Hit rate in the central pixel sensors for 10^8 muon stops/s. The axes enumerate sensor numbers. From top left to bottom right, the four central pixel layers are shown from the inside out. Note that the azimuthal structure in the innermost layer is caused by the elliptical beam spot.

24.3 Front-end requirements

24.3.1 PIXEL DETECTOR

The pixel sensors contain electronics for hit detection and time as well as address encoding. All hits are assigned to a time-stamp (62.5 MHz or 125 MHz).

The hits are then serialized and sent off chip via an 1250 Mbit/s low-voltage differential signalling (LVDS) link (see section 10.4 for details of the readout state machine and the protocol).

The protocol implemented in the MUPIX 7 prototype uses at maximum 68% of the available time slots for sending hit information. Together with the 8bit/10bit encoding, this leads to a maximum hit bandwidth of 685.7 Mbit/s or $28.6 \cdot 10^6$ hits per sensor and second, giving a safety factor of three to four even for the busiest sensors (see figure 24.8).




Figure 24.8: Distribution of sensors according to hit rate at 10^8 muon stops/s.




Figure 24.5: Hit rate in the recurl pixel sensors for 10^8 muon stops/s. The axes enumerate sensor numbers. From top left to bottom right, the maps show the occupancy in the inner and outer, up- and downstream recurl pixel layers.

The total bandwidth requirements for phase Ia and phase Ib are shown in tables 24.1 and 24.2. The system will be built to run at phase Ib rates

(see figure 24.8) with a large safety factor for noise and diagnostic data. For the later phase II, the inner detector part needs to be replaced, whilst




Figure 24.6: Average occupancy of the Fibre Detector per fibre at 10^8 muon stops/s. The three plots show the three fibre layers inside out. The red histograms show the downstream end, the blue ones the upstream end. The missing fibres are at the edges of ribbons.

	Sensor Chips	Max Hits /Chip/s 10^6	Average Hits /Layer/s 10^6	Chip→FPGA link capacity Mbit/s	Chip→FPGA total in Layer Gbit/s	Front-end FPGAs	FPGA→RO capacity Gbit/s per FPGA
Layer 1	48	0.76	28	18.2	0.67		
Layer 2	72	0.46	26	11.0	0.62	8	0.22
Layer 3	432	0.16	38	3.8	0.91	12	0.10
Layer 4	504	0.12	34	2.9	0.81	14	0.08
Total	1056		126		3.01	34	4.08

Table 24.1: Pixel readout requirements (Phase Ia, 10^7 muon stops/s).

the outer pixel readout should be able to cope with the 20 times higher rates.

Besides the bandwidth requirements, also occupancy peaks are of importance, as they determine maximum latencies and thus required buffer sizes

further downstream. Figure 24.9 compares average and maximum sensor occupancies.

The mapping of sensors to FPGAs and the overall addressing scheme can be found in the appendix on page 252; 30 bit are sufficient to address every single pixel in the complete detector system.




Figure 24.7: Maps of the occupancy of the tile detector at 10^8 muon stops/s. The top plot shows the upstream, the bottom plot the downstream detector.

	Sensor Chips	Max Hits /Chip/s 10^6	Average Hits /Layer/s 10^6	Chip→FPGA link capacity Mbit/s	Chip→FPGA total in Layer Gbit/s	Front-end FPGAs	FPGA→RO capacity Gbit/s per FPGA
Layer 1	48	7.6	280	182	6.7		
Layer 2	72	4.6	260	110	6.2	8	2.2
Layer 3	432	1.6	380	38	9.1	12	1.0
Layer 4	504	1.2	340	29	8.1	14	0.8
Recurl IU	432	0.22	60	5.3	1.4	12	0.16
Recurl OU	504	0.20	61	4.8	1.5	14	0.14
Recurl ID	432	0.16	39	3.8	0.9	12	0.10
Recurl OD	504	0.14	39	3.4	0.9	14	0.10
Total	2928		1459		34.9	86	47.28

Table 24.2: Pixel readout requirements (Phase Ib, 10^8 muon stops/s).

For simplicity, hits are assumed to be of 24 bit size on the flexprint links and 32 bit size¹ on all optical links and for storage.

24.3.2 TIMING DETECTOR

In the baseline design, the two timing detectors use the same ASIC. Therefore, the general re-

quirements arising from the common ASIC are discussed here. The sub-detector specific requirements are listed in the readout chapters of the respective detectors.

¹The time information is carried separately at this point

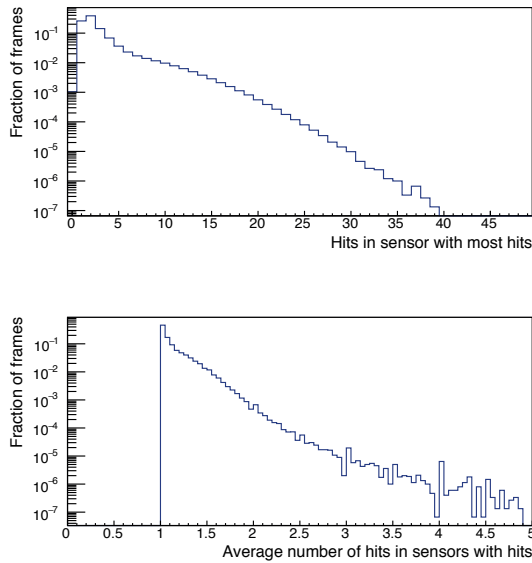


Figure 24.9: Number of hits per 50 ns frame in the busiest sensor of the frame for 10^8 muon stops/s (top) and average multiplicity in sensors with hits at the same rate (bottom).

STiC readout

The baseline solution for the readout of the SiPMs is the STiC² chip, which is a mixed-signal ASIC developed for the readout of SiPMs with the focus on Time-of-Flight applications [68]. A detailed description of the chip can be found in [48, 69, 70].

The core piece of the analogue input stage is a fast discrimination unit with two tunable thresholds. The lower threshold, referred to as *timing threshold*, is used to determine the timestamp of the input signal via leading edge discrimination. The higher threshold, referred to as *energy threshold*, is used to determine the signal charge via the Time-over-Threshold (ToT) of the signal. The chip is designed in such a way that the ToT is directly proportional to the charge over a large range. The time and energy information is digitised via an integrated TDC with a bin size of 50 ps. The intrinsic time resolution of the chip has been measured to be $\sigma_{stic} \lesssim 30$ ps [70]. In section 24.3.2, the functionality of the chip is discussed in more detail.

An important feature of the STiC is the capability to tune the voltage at each input terminal within a range of about 0.7 V. This allows to ad-

²SiPM Timing Chip




Figure 24.10: Sketch of the functionality of the MuSTiC chip, the successor of the STiC3. The time and energy information of the analogue input signal is obtained via two discriminator units. The discriminator output is processed by a TDC with a 625 MHz coarse counter and a fine counter with a bin size of 50 ps.

just the SiPM bias voltage, compensating fluctuations in the detector response, e.g. due to temperature changes, and channel-by-channel variations.

The chip version STiC2 has successfully been tested in a 16 channel prototype of a tiles submodule, which is discussed in section 15.3. The latest version of the chip (STiC3) has been tested with fibre prototypes as described in section 13.2.1. It features an additional zero pole cancellation circuit, which is designed to reduce the signal tail by a factor of three. This results in a significantly reduced channel dead-time. In addition, this is expected to reduce the baseline fluctuations related to the SiPM dark-rate and thus slightly improve the time resolution. Further development of the chip is planned to optimise the performance for its use in the tile and fibre detector, including the ability to handle an increased data rate to about 1 MHz per channel, which is required for the Fibre Detector and phase II of the Tile Detector.

Data Format The data format from the two timing detectors has the same structure. The 50 ps timestamps consists of 20 bit and 1 bit flag.



This is further expanded by 5 bits for the channel number.

Since the Tile Detector also provides an energy information, another supplementary timestamp for the ToT information is added.

The in-phase 50 ps timestamps from the timing detectors divide the common frames, a 16 ns timestamp from the pixel detector, into 320 sub-bins. To store the additional information 9 bits are required. In case of the Fibre Detector 12 bits are needed to address each readout channel: 1 bit for up- or downstream, 5 bits for the ribbon number, 6 bits for the SiPM array channel. In total the information of one hit can be stored in less than 24bit.

24.4 Front-end FPGA boards

As both the MUPIX and STiC chips provide digital output on a LVDS link, a common FPGA based readout board is foreseen. Both the hardware and the firmware for this board are currently being prototyped. For an overview of the firmware components common to all detector subsystems, see page 253 in the appendix.

24.4.1 FRONT-END HARDWARE

The front-end board has to collect the data sent from either STiC or MUPIX on LVDS links, sort and package it and then forward it to the switching boards on a fast optical link. In addition, it has to provide the sensors with control signals, monitor the environment and in case of the pixel sensors also provide high- and low voltage supply. The severe space constraints inside the recur layers necessitate small, highly integrated boards incorporating FPGAs and optical modules with a small footprint (also thermally).

A prototype board based on an Altera Stratix IV FPGA (EP4SGX70HF35C4) was developed in Heidelberg and is currently being manufactured. It has the link and processing capabilities of servicing one quarter of the vertex detector (5 half-ladders with three chips with three high speed LVDS links each). The board has a wide range of testing facilities built in and thus a larger form factor than required for the final version, which we will design as soon as we have ascertained that all the necessary functionality is working on the prototype. There is also an option to change to the next generation of FPGAs (Altera Arria V,

5AGXBA1D4F31C5), which would bring advantages in terms of the number of LVDS links per device, power consumption and cost.

For the optical data transmission we use the MiniPOD™ by Avago (AFBR-811VxyZ transmitter and AFBR-821VxyZ receiver), which combines twelve optical channels with up to 10 Gbit/s in a very small footprint of $19 \times 22 \text{ mm}^2$ [79].

Concurrently with the board development, we have created a first version of the firmware in order to ensure correct pin wiring and get a first estimate of the power consumption.

24.4.2 FRONT-END FIRMWARE

The front-end firmware contains both a set of common functionality for all sub-detectors and blocks which are specific to the pixel, tile and fibre detectors. Many of the firmware blocks are already under development and can be tested with particle beams in the MUPIX telescope. Figure 24.11 shows an overview of the data flow and firmware blocks in the pixel front-end FPGA.

We currently develop the firmware on Altera FPGAs; for most components, no or very little changes are required between Stratix and Arria models or even for a vendor change to Xilinx.

Pixel front-end firmware

Data path

Receiver For receiving the LVDS data streams from the MUPIX and also the STiC chips, we use the built-in LVDS receiver and deserializer circuitry of the FPGAs. It turns the 1250 MHz incoming serial bit stream into a series of 10 bit wide words at 125 MHz. It also synchronizes the data for a whole block of receivers (up to 28) to a single clock using dynamic phase alignment (DPA) - note that the frequency needs to be locked to the LVDS frequency, which is however straightforward, as the FPGA also provides the clocks to the front-end chips.

8bit/10bit decoder We use the open source 8bit/10bit decoder from OpenCores [80], which is closely modelled after the original IBM implementation [35, 36] whose patent has expired. It turns the 10 bit data stream into a decoded 8 bit data stream plus a K-marker, indicating comma words.



Mu3e Pixel Front-End Firmware
Inner detector version (15 links)

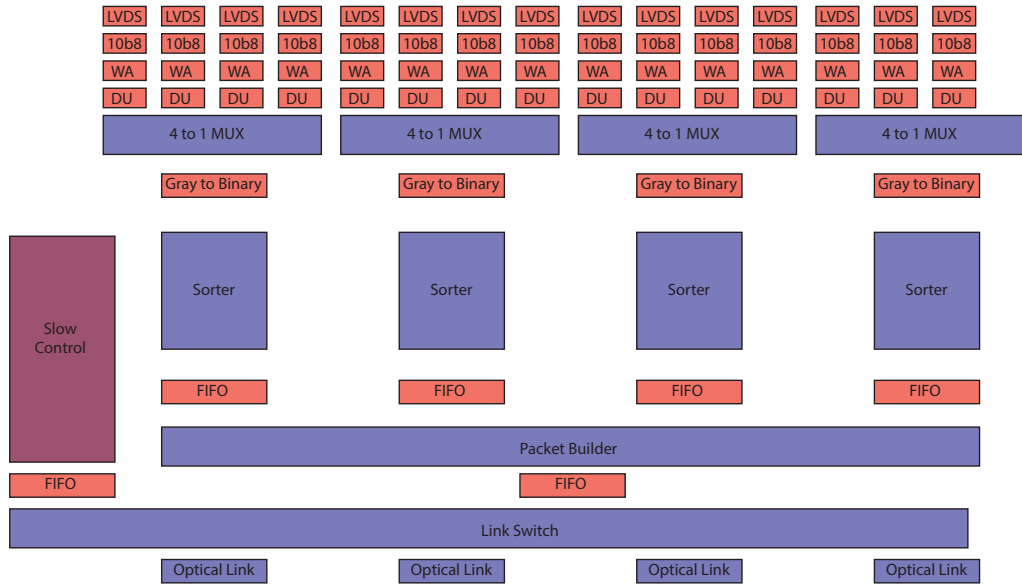


Figure 24.11: Components in the main data flow of the inner detector pixel firmware. *10b8* stands for the 8bit/10bit decoder, *WA* is the word aligner and *DU* the data unpacker. For the outer detector, the layout is similar, but with 36 inputs and 9 sorter chains.

Word aligner This standard piece of firmware is used to insert bit slips into the incoming data stream until a comma word is repeatedly detected, ensuring correct word alignment. In order to facilitate this, the MUPIX chip generates a long string of comma words before starting the readout. Once locked, this unit monitors that comma words appear frequently in the data stream and that the alignment has thus not been lost.

At the moment, the tasks of the above described firmware are undertaken by Stratix GX transceivers as provided by Altera [81] in the MUPIX telescope. No errors attributable to this part of firmware have been observed in test beam campaigns or in the lab. Implementations using the LVDS receivers foreseen for these tasks in the detector setup exist as well and are currently tested in the lab.

Data unpacker This piece of firmware removes comma words and separates the data stream into hit and counter information, i.e. in-

verts the data protocol used by the MUPIX. The output are 24 bit hits and 24 bit counters. In addition, errors are detected when the incoming data-stream is not consistent with the expected data protocol from the MUPIX.

The data unpacker has been successfully tested in simulations, the lab and in test beams with the MUPIX telescope.

4-1 multiplexer In the current MUPIX protocol, one hit (and also one set of timing information) occupies four 8 bit words; a full hit can thus only arrive once in every four 125 MHz clock cycles. This in turn allows for multiplexing up to four input data streams into one 125 MHz hit stream, reducing the number of required sorters (and thus the memory footprint of the firmware) by a factor of four.

The multiplexer has been extensively simulated and tested both in the lab and with four MUPIX chips in the MUPIX telescope.



Address	TS 7-5	TS 4-0	counter 3-0
Time-stamp	Block	Time bin bits 4 down to 0	Hit number
0	Block 0	0	
1		1	
2		2	
3		3	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	
12		12	
13		13	
...		...	
31		31	
32	Block 1	0	
33		1	
34		2	
35		3	
36		4	
37		5	
38		6	
39		7	
40		8	
41		9	
42		10	
43		11	
44		12	
45		13	
...		...	
63		31	
Blocks 2-6			
224	Block 7	0	
225		1	
226		2	
227		3	
228		4	
229		5	
230		6	
231		7	
232		8	
233		9	
234		10	
235		11	
236		12	
237		13	
...		...	
255		31	

Table 24.3: Layout of the sorting memory

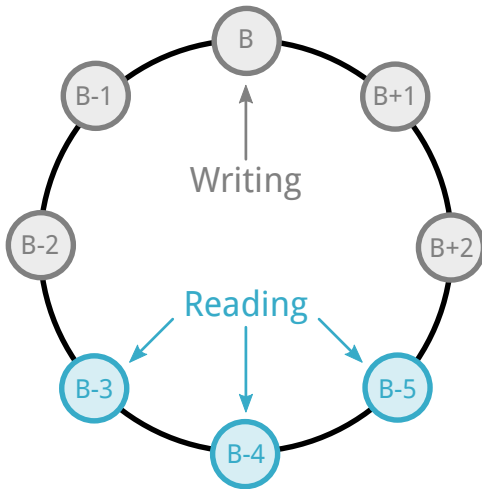


Figure 24.12: Ringbuffer of the sorter indicating which blocks B are occupied by the write and read process.

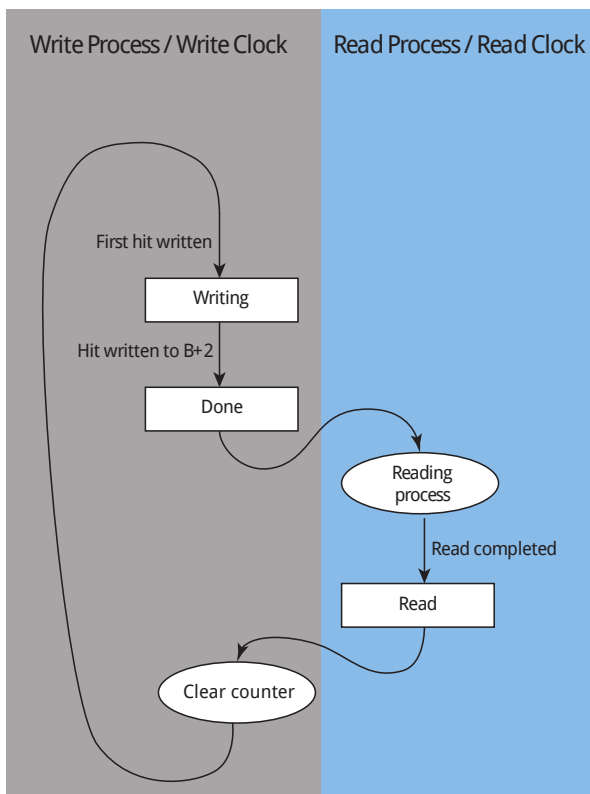


Figure 24.13: The life of a memory block between the write and read domains with status signaling via flags.

Sorter The readout sequence in the MUPIX chip does not strictly conserve the time ordering

of the hits, i.e. a hit with a later timestamp can be streamed out before one with an earlier one. Prior to event reconstruction, the hits have to be sorted by timestamp. This is done early in the readout chain as merging sorted data streams from different parts of the detector further on is then straightforward.

The use in the front-end firmware requires a fully pipelined sorting algorithm. In essence, a memory with bins for every timestamp is used, which sorts hits by storing them in the correct order. In the following the sorting algorithm is described in more detail.

Writing to the sorting memory uses the timestamp as part of the address. The lowest 4 address bits are given by a counter for each timestamp, allowing for a maximum of 15 hits per time slice (the counter value 0 is reserved for no hits in the slice). If more than 15 hits for a slice arrive, the surplus hits will be ignored and an overflow flag for this specific time slice is set. The layout of the memory is shown in table 24.3. Each block of 32 timestamps will end up in a readout frame. Whenever a hit belonging to a new block B arrives³, the block B is marked as *writing* and the blocks $B - 3$ to $B - 5$ marked *writing* are re-marked as *done*, which signals to the memory reading process that reading can start (see figure 24.12). Once the read process is finished reading a block, it marks it as *read*, which in turn prompts the writing process to clear the associated counters and overflow bits as well as the *done* flag. This memory block is now ready to be written to again. The complete process is illustrated in figure 24.13.

If a block is marked as done, the reading process starts building the output by first constructing a header. At the same time, the counters containing the number of hits per timestamp bin are copied. The memory read address is constructed from the current block, the current bin and the current counter value. As soon as a bin is empty (the counter reaches 0), the algorithm switches to the next bin skipping bins without data. If all bins are empty, the end of block marker is output, the block is marked as *read* and the algorithm moves to the next non-empty block or to the idle state if no blocks are ready.

Note that the read process can be run at a different (higher) clock frequency than the write process. Output is zero suppressed; blocks without

³The block arithmetics are understood to be modulo the total number of blocks, eight.



hits will not appear in the data stream (this can easily be changed).

The complete data sorter as currently implemented (i.e. without diagnostics or debug facilities) requires 2025 ALMs and 73728 Bits of memory. This implies that a sufficient number of these entities can be put on a 70K FPGA to serve as many chips as can be connected via LVDS.

We have a very extensive simulation of the sorter (which is one of the most complex blocks of firmware in the data stream) and have also tested it in a first beam test at CERN in July 2015, which led to a series of improvements to the algorithm. Subsequent tests at PSI and DESY in October 2015 both delivered the desired performance with hit losses consistent with overflows past the 15 hits per time bin and occasional very late hits.

Packet builder Depending on the number of chips connected to one front-end FPGA, the output from between four and nine sorters has to be collected in a single, synchronized data stream. We intend to do this via an array of FIFO memories at the sorter outputs and using a time sorted tree reading mechanism. In addition, slow control data has to be mixed into the data stream at opportune moments.

This algorithm is currently in the prototyping stage and will be tested as soon as possible.

Link switch In view of the phase II experiment, every front-end FPGA will be connected to four full-duplex optical links. As this bandwidth is not required in phase I (and the required switching boards will not be present), we can use the extra links as spares in case of failures or for debugging and monitoring purposes. A switch is required to choose the link used for data transmission and links used for debugging.

This simple piece of firmware needs to be written and tested.

Gbit/s transmitter The data will be sent out via a 6 Gbit/s link. Most of the required firmware is provided directly by Altera. We add a parity controller to ensure a maximum run length of consecutive zeroes or ones on the link. The link settings and the corresponding maximum bit error rates have been studied in detail in a master- [82] and a bachelor thesis [83] and fulfil the requirements of the experiment.

Slow-control path For slow control data, two paths are foreseen; on the one hand, surplus bandwidth on the optical links to and from the switching board can be used, which is especially useful for large volume data such as pixel tune values. On the other hand, a separate differential line for use of the MCSB protocol (see chapter 23) is foreseen, which is especially useful for monitoring the status of optical links, switching power, etc..

The interface to MCSB and the slow control related tasks on the FPGA will be implemented in a NIOS II processor core [84] on the FPGA. We have successfully implemented and tested such cores on FPGAs of the target families and are currently adapting the MSCB interface code to the architecture.

Timing front-end firmware

The firmware for the timing systems (both fibres and tiles) will be very similar to the pixel firmware, with the exception of details in the hit data format (24.3.2) and the sorting algorithm. Development of the detector specific code will start as soon as the readout solution and data transfer protocol for the fibres (*time mode*) and tiles are finalized.

24.5 Read-out links

In total there are three different types of read-out links in the Mu3e data acquisition system. The same type of links can be used for a (smaller) number of slow control links in the opposite direction.

The data from the MUPIX chips will be transmitted to the front-end FPGAs via LVDS links at 1250 Mbit/s. The fast serializer and LVDS drivers for the MUPIX chip is implemented in the MUPIX 7 prototype and working as expected, see section 10.2.

The link is physically implemented as a matched differential pair of copper traces on the sensor flex-print, see section ??.

There are two types of optical high speed data links. The first one is going from the front-end FPGAs to the switching boards, the second from the switching boards to the FPGA PCIe boards in the event filter farm PCs. The optical links from the front-end FPGAs to the read out boards have a bandwidth of 6 Gbit/s, which fits well with the FPGA specifications. Each FPGA has at least nine fast transceiver blocks, which connect to the



MiniPOD optical driver (see section 24.4.1). The data laser has a wavelength of 850 nm and the optical fibre is of 50/125 multi-mode type, since this is a standard both in industry and in particle physics detector readout. The fibres from the 12-way emitter will go to a splitter and patch panel in order to combine fibres going to the same switching board.

The second type of link connects the switching cards to the FPGA PCIe boards in the event filter farm PCs. This optical link will be implemented as a 10 Gbit/s high speed link. Since each switching board is connected to every sub-farm PC with one of the high speed links, these are point to point single fibre links. The fibres are of 50/125 multi-mode type operated at 850 nm to stay compatible with the links from the detector. The PCIe FPGA board is fitted with a quad small form-factor pluggable (QSFP+) optical module.

All the links have been tested using the development hardware and were found to have bit error rates low enough for stable and consistent running of the experiment [82, 83].

24.6 Switching cards

The main task of the switching cards is to act as switches between the front-end FPGAs on the detector and the on-line reconstruction farm, thus allowing each farm PC to see data from the complete detector. The board design and choice of FPGAs is dominated by the number of fast links required.

We plan to use the *PCIe40* card currently under development for the LHCb and ALICE upgrades at LHC [85,86]. These boards provide up to 48 full duplex optical links at up to 10 Gbit/s, plus two eight lane PCIe 3.0 interfaces bundled to a sixteen lane interface by a switch. The FPGA is an ALTERA Arria 10, which is currently available as an engineering sample. All functionality essential for Mu3e has been successfully tested with two prototypes available in Marseille [87]. We have the framework for firmware development in hand and should receive a prototype board in early 2016. Due to the large number of links, a single board is sufficient for phase Ia operation; phase Ib will require four boards (plus spares). The PCs hosting the cards can be used to store and transmit the extensive pixel slow control and tuning data and

link the cards to the outside world via standard Ethernet.

Already well established boards like the LHCb TELL1 cards [88] or PANDA compute nodes [89] in larger numbers would also fulfil our needs, but complicate the receiver on the reconstruction PC side and require a crate infrastructure.

24.6.1 SWITCHING BOARD FIRMWARE

The firmware for the switching board will be based on components already produced for either the front-end or the event filter interface; the task is to receive several data streams in parallel, merge them synchronously and then forward them to the event filter. The merging is very similar to the packet builder in the front end (except with higher performance requirements).

The PCI express interface developed for the event filter will be used to inject and extract large volume slow control data (pixel tune values).

An extra piece of firmware is needed in case of the fibre tracker, where hits from both ends of the fibre have to be matched in order to suppress dark counts. This firmware will be developed once the fibre data protocol is fixed.

24.7 Event filter interface

The filter farm PCs will be equipped with FPGA cards in PCIe slots and optical receivers. We currently foresee to use Altera Straix V evaluation boards here (see figure 24.14) [90,91]. They provide one QSFP slot for a fourfold optical receiver, which is sufficient for phase I. This FPGA will perform the event building and buffering and also allows to run simple clustering, sorting and selection algorithms. The event data are then transferred via Direct Memory Access (DMA) over the PCIe 2 bus⁴ to the main memory of the filter farm PC and are subsequently copied to the memory of a GPU, where the fitting and vertex selection algorithms are run. The GPU then posts selected events to the main memory of the PC, from where the CPU ships these events and some monitoring data via Ethernet to the central data acquisition computer running the MIDAS software. At that computer, the data streams from the farm PCs

⁴Note that PCIe is actually not a bus protocol, but offers switched point-to-point connections. The *bus* designation is due to the software-side backwards compatibility to the original PCI bus interface.




Figure 24.14: ALTERA Stratix V PCIe development board.

are combined into a single data stream, combined with various slow control data, compressed and stored. The connection to the PC is via an eight lane PCIeExpress interface, which can be operated both at PCIe 2.0 and PCIe 3.0 speeds. We have working firmware for the PCIe 2.0 case, which is described in section 24.7.1.

The maximum data rate over an eight-lane PCIe 2.0 bus is 4Gbyte/s, amply sufficient for phase I. For the full phase II rate, the raw link speed of PCIe 3.0 (16 Gbyte/s for 16 lanes) is still sufficient, would however have to be fully and efficiently used. The PCIe 4.0 standard, doubling this rate, should become commercially available at the end of this decade, compatible with phase II running; alternatively, the number of farm PCs could be increased.

The GPU boards will be obtained commercially as late as possible in order to profit from the fast developments and sinking prices. As far as raw floating point throughput is concerned, current high-end GPUs already pack enough power for high rate running [92, 93]. Newer cards are however expected to offer higher memory bandwidth and better caching. For example, between the GTX 680 and the GTX 980 GeForce GPUs, both the compute power and the copy speed increased by 30 %. The latter is due to the two copy engines present on the GTX 980 which allow for concurrent copying to and from the device. Also the performance of the driver software, especially as far as the PCIe 3 bus is concerned, and the GPU compilers is expected to improve.

We currently plan to host the farm PCs in individual rack-mounted tower casings, ensuring

enough space for the FPGA board and the high end GPU whilst allowing for air cooling. Each tower will consume around 0.5KW, so adequate cooling of the racks and the counting house will be necessary.

24.7.1 PCIe INTERFACE

Altera provides the low level PCIe interface; our efforts concern the transaction layer packets (TLP) and higher layers. The interface exposes four types of bus addressable regions (BARs), namely blocks of 64 32 bit registers, either writeable from the host or from the FPGA side and memories of 256 KB, also writeable either from the host or the FPGA side. In addition there is a DMA engine for block data transfers from the FPGA to any other bus addressable memory. The register banks as well as the memories support reads and writes of single aligned 32 bit words.

DMA Engine

Data to be transmitted via DMA is written to a 256 kB RAM. Every time 4kB have been written, one DMA block is prepared and sent off via the PCIe interface. As shown in figure 24.15, each block is split up into 16 PCIe blocks since the maximum payload possible with the typical chipsets in use on our motherboards is 256 B (the FPGAs could also go to the PCIe maximum of 2048 B). A three- (four-)word header in each PCIe block contains the destination address in the CPU RAM for memory below (above) 4 GB, the data length and format and traffic control flags [94]. If the PCIe transmitter is ready, the data is sent to the specified address. Otherwise, the next block of data is written and when the transmitter is ready the first block which has not yet been sent is transmitted. Since the RAM can hold 64 DMA blocks, this is the upper limit of blocks waiting to be transmitted. The size of this buffer limits the rate of data the DMA engine can sustain, as the bus is occasionally busy with housekeeping tasks for several hundred cycles. As DMA happens without interference of the CPU, the processor needs to be notified about the copy progress. Every 64 DMA blocks, the current memory position is written to a control buffer in CPU RAM and an interrupt message is sent. This tells the processor that something important happened and action is required by the driver to deal with the interrupt. In




Figure 24.15: Scheme of PCIe and DMA blocks.

our case, the software is informed that 256 kB of memory can be read.

24.8 Data transmission to the GPU

Every time an interrupt is received from the FPGA, the read function of the user-space program is called. It contains information about the location last written to in the ring buffer as well as the amount of data written since the last interrupt. This takes care of any missed interrupts.

Data arriving via DMA from the FPGA is stored in a ring buffer on the main memory (RAM) of the PC and then copied to the GPU memory. The copy from RAM to GPU memory also occurs via DMA, initiated by the NVIDIA driver of the graphics card. For this purpose, page-locked memory is needed and the NVIDIA driver must know that it is usable for DMA. NVIDIA's computing platform and Application Program Interface (API) CUDA provides a memory allocation function which allocates exactly this type of memory on the RAM. If memory from a normal user-space memory allocation is used to copy data from RAM to GPU memory, a step in between is needed where the data is copied to DMA-capable memory and only after that it is copied to GPU memory. To avoid this extra copy within RAM, we use the CUDA API to allocate a ring buffer with a size in the order of GB.

The user space pointer to this memory is communicated to the kernel module handling the FPGA via input/output control (ioctl) which is a device specific function handling input/output communication that is not possible with a normal system call.

As seen from virtual memory in user space, the ring buffer is one large chunk of contiguous memory. On the level of physical memory however, we deal with many smaller pieces scattered across

the available memory. Within the kernel module, the pages for the ring buffer memory are obtained. For each contiguous piece of memory, the virtual address is mapped to a physical and finally to a bus address using the Linux kernel's DMA layer. The latter mapping is needed since an I/O memory management unit is installed on the motherboard which translates physical addresses to bus addresses as seen from external devices. The mapping process provides the bus address and length of each contiguous piece of memory in the ring buffer. These are written into a RAM on the FPGA and are used as destination addresses by the DMA engine.

With this setup, we obtain a bit error rate $\leq 4 \times 10^{-16}$ at a data rate of 1.5 GB/s. Theoretically, 4 GB/s are possible with eight lanes of PCIe 2.0. However, there is an overhead due to PCIe packet headers and communication packets on the TLP layer and below. We are further limited by the size of RAM on the FPGA buffering the data to be transferred. When increasing the data rate, losses due to an overflow in the RAM are observed. However, since data arrives at a maximum of 10 Gbit/s from the switching boards, this is fully sufficient for phase one of the experiment.

24.8.1 EVENT FILTER FIRMWARE

On the event filter FPGA cards, several steps have to be taken in order to prepare the arriving hit data for reconstruction on the GPU. Here we report the required pieces of firmware and their current status.

Receiver/Synchronizer

In phase Ia, each PC will receive data over a single 10 Gbit/s optical link, greatly simplifying the receiver firmware. For phase Ib, four high-speed data streams have to be aligned and merged, similar to the same task on the switching boards. The corresponding firmware will be derived from the switching board firmware.

Coordinate transformation

The addresses of the incoming pixel hits have to be transformed to 3D Cartesian coordinates in floating point format. Assuming plane sensors, this can be done by scaling and adding three vectors:

$$\vec{x} = \vec{c} + r \cdot \vec{u} + c \cdot \vec{v}, \quad (24.1)$$



where \vec{x} is the hit position obtained from the position of the lower left corner of the chip \vec{c} and the row r times the row direction \vec{u} normalized to the pixel size and the column c times the column direction \vec{v} . In total, six multiplications and six additions are required, together with memory space for storing the \vec{c} , \vec{u} and \vec{v} for every sensor.

A fully pipelined version of this firmware using the DSP blocks on the Stratix V FPGA has been written and tested in simulations [82].

Triplet preselection

In order to reduce divergence in the GPU code (see section 29.2), a pre-selection of hit triplets from the first three layers potentially belonging to a track candidate has to be performed on the FPGA and the result stored in a GPU friendly memory layout.

The corresponding firmware will be developed when the requirements of the GPU selection algorithm are fixed.



RUN CONTROL, DATA COLLECTION AND STORAGE

25.1 The MIDAS system

The filter farm will output selected events at a data rate in the order of 50 MBytes/s in total. This data rate is low enough to be collected by a single PC connected to the filter farm by common GBit Ethernet and written to local disks. Then the data will be transferred to the central PSI computing centre, where it is stored and analysed. For the central data acquisition the well established MIDAS (Maximum Integrated Data Acquisition System) [95] software package will be used. This software is currently used in several major experiments such as the T2K ND280 detector in Japan [96], ALPHA at CERN and the MEG experiment at PSI [97]. It can easily handle the required data rate, and contains all necessary tools such as event building, a slow control system including a history database and an alarm system. A web interface allows controlling and monitoring the experiment through the Internet. The farm PCs will use MIDAS library calls to ship the data to the central DAQ PC. The framework also offers facilities to send configuration parameters from a central database (the “Online DataBase” or ODB) to all connected farm PCs and to coordinate common starts and stops of acquisition (run control).

For the purpose of monitoring and data quality control of the experiment the MIDAS system offers taps to the data stream for connections of analysis and graphical display programs. The output of graphical user interface programs can be fed back into the web interface of the MIDAS system

so the experiment can be monitored also remotely with just a web browser.

25.2 Run start/stop synchronisation

In traditional data acquisition systems, starting and stopping of the data acquisition is controlled by enabling and disabling trigger signals. In a streaming system such as in Mu3e, this is not an option and great care has to be taken in order to synchronize data across the complete detector at run start and ensure that the frame numbers in all subsystems are in lockstep.

To this end, a global reset signal is distributed together with the global clock. At the front-end, the reset signal is forwarded to the pixel sensors and there sets the timestamp counters to zero as long as it is on (note that the pixel sensors cannot be inactivated, so even during a reset they will still collect, process and send hits, however all with timestamp zero, see also section 10.4.3). At run start, the reset signal is released synchronously for all sensors, which then start counting timestamps. The front end firmware will ignore all hits with timestamp zero at the beginning of the run and start sending packets into the switching network as soon as non-zero timestamps arrive. All subsequent stages in the network then synchronize on the first packet and from then on stay in sync using consistent packet numbering.



The currently available STiC version supports two reset modes. A single clock (the digital part is operated at 160 MHz) cycle signal resets the coarse counter values of the TDC; a synchronisation pulse. A longer pulse resets also the finite-state machine of the sequential digital circuits. All present hits are removed, the TDC gets ready and the coarse counter (timestamp) are set to an initial value. As long as reset is active no hits enter the hit logic, therefore no hits are transmitted. Furthermore, the frame number (6.4 μ s) of the serial data transmission is also reset. The first not reset frame number signalizes run starts. After a

reset, the timing detector hits in this frame can be matched directly to the first non-zero time stamps from the pixel detector. This ensures the synchronisation of the two systems (with respect to the reset signal). It is planned to keep the current reset scheme in MuSTiC unchanged.

Similarly, at run end, the timestamp reset goes high and the frontend continues forwarding packets until timestamp zero is detected. Unfortunate timing of the run end can lead to partial events arriving at the filter farm, which then have to be rejected either on firmware or on software level.

PART VIII



SOFTWARE



MILESTONES FOR THE MU3E SOFTWARE

1. Fully include fibres and tiles in the reconstruction and analysis.
2. Prepare reconstruction for detector raw data.
3. Prepare analysis code for the search for $\mu \rightarrow eee$.
4. Prepare analysis code for other measurements.



SOFTWARE DEVELOPMENT AND TOOLS

The Mu3e experiment specific software comprises simulation, reconstruction and analysis programs, which are described in detail in the following chapters. Here we would like to introduce the chosen development process and both self-developed and external software tools and libraries.

26.1 Development Process

At the time of writing, simulation and reconstruction software are fairly mature. The development process has smoothly transitioned from essentially a single developer to a core team of about five very regular contributors supported by a larger circle of collaboration members occasionally writing code. So far, we have not discovered major design flaws requiring redesigns; small, targeted refactorings improve the code in areas where it does not keep up with demands.

26.1.1 ORGANISATION

The milestones and main directions of the software development are regularly discussed at the Mu3e collaboration meetings. The software coordinator then breaks down the tasks into smaller work packages, which are distributed among the developers, mostly according to the subdetectors concerned. A brief weekly meeting (typically conducted on Skype [98]) discusses progress, re-

distribution of tasks as well as technical and design issues.

The host of our repositories (BITBUCKET.ORG [99]) provides functionality to directly comment committed code as well as a chat room [100] displaying all commits and comments which allows for brief informal discussions of technical issues between meetings.

26.1.2 VERSIONING AND QUALITY CONTROL

The use of the distributed version control system GIT [101] allows for easy branching and merging. Typically, every developer works on his own branch, with the complete freedom to merge in developments from colleagues at any time. Once a development is deemed complete, the developer creates a *pull request*, which is then reviewed by the software coordinator and/or a subsystem expert. These reviews allow us to spot problems early and educate developers in good coding practices without having to enforce a rigid set of rules. After requested changes have been implemented, the code is merged into the main branch and a build as well as a few simple tests are performed to ensure the functionality. After the addition of significant new functionality or in preparation of large Monte Carlo production, the state of the main branch is subjected to extensive simulation and reconstruction tests and then tagged as a release, for which also the complete API documentation as well as release notes are built.



26.2 Tools and Libraries

The selection of tools and libraries for the Mu3e experiment was guided by the desire to keep external dependencies to a minimum, whilst at the same time avoiding the re-implementation of existing and well-tested functionality. We also encourage the use of up-to-date versions of external software and try hard not to depend on particular versions.

We are committed to use open source, community supported software wherever possible and also try to contribute fixes, extensions and our own developments back to the community.

26.2.1 OPERATING SYSTEM

The Mu3e software compiles and has been tested on a variety of Linux [102, 103] operating systems as well as on Mac OS X [104]. In principle, it should also compile and run on Windows systems, however no collaboration member has so far been willing to set up the corresponding build environment including the root and Geant4 libraries.

The simulation clusters in Heidelberg and Mainz as well as most development workstations run on a recent version of OpenSuse [105], which is our suggested flavour of Linux. Problems with other distributions are typically quickly identified and fixed.

26.2.2 PROGRAMMING LANGUAGE

The Mu3e programming language of choice is C++ [106], as it seamlessly interfaces to the libraries used and can easily be reduced to plain C [107] for use in micro-controllers or on GPUs. For offline usage, we encourage the use of the new features introduced in C++11 and C++14 [108]. Typically, we employ recent versions of the GNU compiler collection [109] and try to avoid compiler warnings, however the code should compile equally well with e.g. the Intel compilers [110].

For programming the GPUs used for the online reconstruction, we use the CUDA [111] extensions to C provided by NVIDIA, as this is currently by far both the most powerful and most user-friendly framework available. Unfortunately, the open and vendor-independent framework OpenCL [112, 113] is far less developed, which makes it very hard to access low-level functions of the cards and also has a much smaller user community.

26.2.3 BUILD SYSTEM

The Mu3e core software employs CMake [114] as a build system, which automatically checks dependencies and creates makefiles.

26.2.4 EXTERNAL LIBRARIES

Geant4

Geant4 [45, 115–119] is the de-facto standard for the simulation of particle interactions with matter. We employ the library for the description of the detector geometry and for the simulation of all physics processes; in order to avoid the use of external generators, we have implemented several new processes within Geant4, see section 27.3.2. We strive to validate important parts of the simulation in beam experiments, which has led to a new model for multiple Coulomb scattering [120]. The continuous beam and streaming data taking required some changes to the Geant4 event structure, which are described in section 27.4.

Root

The root histogramming package [121] has, despite its unfortunate name and some questionable design choices, become the standard analysis tool in high energy physics. We use it for simulation and analysis data storage (for raw data we will use a well documented binary file format based on MIDAS [95]) as well as for histogramming. We thus try to minimize the dependence on root whilst at the same time allowing students to gain root experience, which is nowadays a core skill in high energy physics. Whenever possible, we use tools improving on root, e.g. PyRoot or self-filling histograms (see below).

Self-filling Histograms

The self-filling histogram package [122] was developed as an add-on to root, allowing for a more object-oriented and declarative programming style in analysis. It is especially useful if very large numbers of similar histograms need to be declared and filled, e.g. in systematic studies or data quality tasks.

Qt

For parts of the code that require a graphical user interface (GUI), namely interactive Geant sessions



and the event display, we use the Qt library [123] for the GUI parts of the code. The event display requires version 5 of Qt, as this has much improved OpenGL [124] support.

Qt is also extensively used for GUIs for detector tests and the MUPIX telescope.

Wherever possible, we plan to replace Qt based visualisations with implementations that can be directly displayed in a web browser and do not need any locally installed code.

26.2.5 MU3E LIBRARIES

Versioned Files

The Mu3e core software uses a subclass of root files which store the configuration with which they

were generated, namely all configuration files as well as the git hashes of the code. Every additional processing step adds its own configuration data. This allows to reconstruct all the settings that were used to generate a file without external bookkeeping.

Watson tracking library

The Watson library was developed as part of a Ph.D. project in Mu3e [28] and offers a fast and unified implementation of all the core tasks in the tracking, namely propagation of particles in the field, track- and vertexfits (triplets, fast helix [125] and broken lines [126, 127]) and full error propagation. The toolkit is now also used in the P2 experiment at Mainz.



SIMULATION

This chapter describes the Geant4 [115, 116] based simulation used to study and optimize the detector design, to develop the reconstruction code and to estimate signal efficiency and background rates.

27.1 Detector geometry

The simulated detector geometry closely follows the planned detector geometry described in earlier chapters. The structure of parts of the mechanics, electronics, services and cabling outside of the sensitive regions are represented by equivalent blocks and rings of material until we know the exact geometry. The simulated volume extends for three meters in all directions from the target center; the magnet metal and the surrounding volume is however only used in the cosmic ray simulation. Figures 27.1 and 27.2 show the simulated detector in the phase Ia configuration and figures 27.5 to 27.7 do so for the phase Ib set-up.

27.1.1 BEAM DELIVERY

In the detector simulation, the beam is started 1 m in front of the target inside the beam pipe. Beam particles are generated with a profile and momentum spectrum taken from the beam simulation at the same spot. The beam passes a thick lead collimator and shortly before the target it is slowed down by a 600 μm Mylar moderator, followed by another collimator removing particles undergoing large angle scattering in the moderator and then exits the beam vacuum through a 25 μm vacuum window.

27.1.2 TARGET

The target is simulated as a hollow Mylar double cone supported by three nylon strings at each end and a nylon string along its axis, see also chapter 6.

27.1.3 PIXEL DETECTOR

The simulated geometry of the pixel detector includes the sensor, the flexprint (with an average trace density assumed and represented as thinner metal layers) and the Kapton support structure. The plastic end pieces and aluminium support wheels are also simulated. Currently not in the simulation is the data and power cabling and the piping for the helium distribution.

27.1.4 SCINTILLATING FIBRES

The fibre ribbon simulations accounts for the fact that many different fibre, ribbon and readout configurations are still under discussion. Optionally round or square fibres of adjustable diameters are placed in ribbons in such a way that the whole Fibre Detector is located at a radius of around ~ 6 cm. Parameters such as fibre type (round or square), fibre diameter, thickness of possible Aluminium coating, dead width between fibres (filled with glue), number of fibre layers, number of fibres per ribbon or number of ribbons can be freely chosen. Depending on the fibre shape, cladding thickness and staggering is adjusted.

Different fibre matchings to SiPM schemes are available. In a single fibre scheme each fibre is matched on both sides to exactly one SiPM. In




Figure 27.1: Perspective view of the simulated phase Ia detector cut open at $x = -19.5$ mm.




Figure 27.2: Side view of the simulated phase Ia detector cut open at $x = -19.5$ mm.




Figure 27.3: Perspective view of the simulated phase Ib detector cut open at $x = -19.5$ mm.




Figure 27.4: Side view of the simulated phase Ib detector cut open at $x = -19.5$ mm.




Figure 27.5: Perspective view of the central part of the simulated phase Ib detector cut open at $x = -19.5$ mm.




Figure 27.6: Side view of the central part of the simulated phase Ib detector cut open at $x = -19.5$ mm.

an array-like scheme, SiPM array cell width, cell pitch, number of cells and misalignment can be configured.

The baseline consists of 24 ribbons with 3 layers of which each consists of 64 round $250\ \mu\text{m}$ fibres read out by SiPM arrays with $250\ \mu\text{m}$ cell width.

27.1.5 TILE DETECTOR

In the simulation, the Tile Detector consists out of the scintillating tiles, SiPMs, two layers of PCBs hosting the SiPMs and the readout chips, respectively, as well as the support structure. The length of the Tile detector in beam direction is 36 cm for each recurl station, matching the active length of the pixel tracker. The material of the scintillating tiles is Polystyrol. The geometrical arrangement of the tiles is parameterised by several input parameters. The default configuration is as follows: the scintillator is segmented in 60 tiles along the beam direction. In the $R - \phi$ plane the detector is segmented in 56 tiles arranged in a tetra-decagon shape (14-fold symmetry) with 4 tiles per side. The first and last tile of each side is bevelled by 12.86° , which results in a smooth transition between neighbouring sides. The SiPMs underneath the scintillator layer are modelled by a 1.4 mm continuous layer of Silicon. Since the scintillator is the last active layer along a particles track, the simplified SiPM geometry is fully sufficient to obtain accurate simulation results. The SiPMs sit

on a steel support structure, which is covered by a layer of PCBs on both the inner and outer surface.

27.2 Magnetic field

The magnetic field in the simulation is taken from a azimuthally symmetric field map with 10 mm step size calculated in the magnet design and tendering process. Linear interpolation is used between the field map grid points. It can easily be replaced by a measured map once the magnet is available.

27.3 Physics Processes

27.3.1 MULTIPLE COULOMB SCATTERING

Multiple Coulomb scattering is the main limiting factor for the resolution of the experiment; an accurate simulation is thus crucial. The best results are obtained by simulating each individual scattering, which however results in prohibitively large computing times. A large selection of multiple scattering parameterizations are available in Geant4; in a test setup they were compared to the single scattering model, see Figure 27.8. The best overall description is obtained from the Urbán-Model [119] at large step widths, which also has the shortest computation times. In the helium gas




Figure 27.7: View along the beamline of the central part of the simulated phase Ib detector cut open at $z = -55.5$ mm.

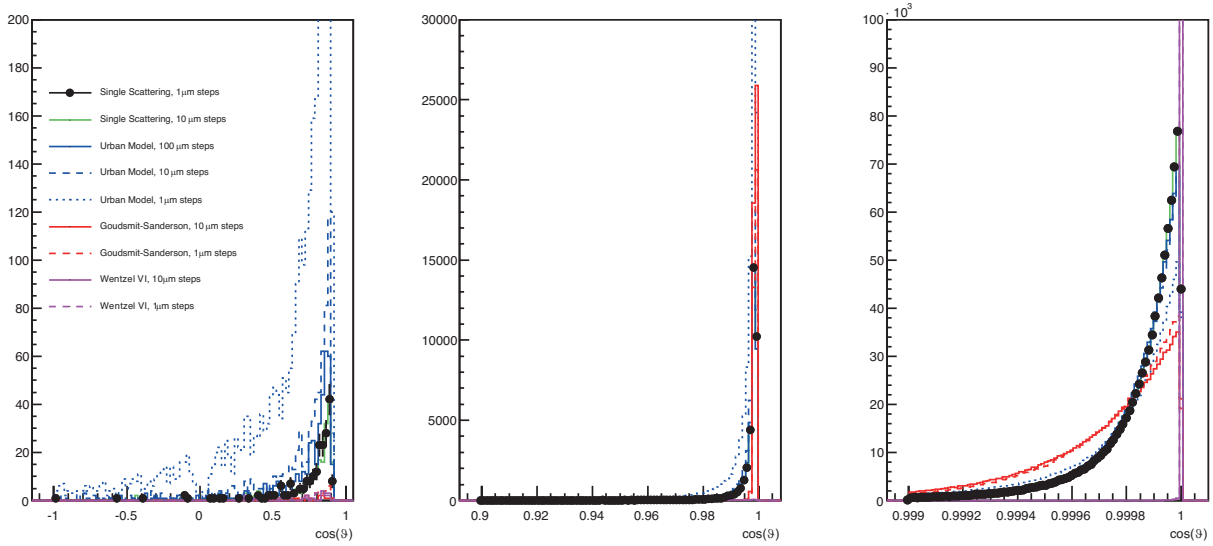


Figure 27.8: Comparison of multiple coulomb scattering models in different scattering angle ranges. The scatterer is a single silicon-Kapton assembly shot at a right angle with 30 MeV positrons. The black dots and the green line show the single scattering model which serves as a reference; as expected, the single scattering model is not affected by the Geant step size. Of all the parameterizations, the Urban model with a step size that treats each bit of material as a single volume performs best.




Figure 27.9: Comparison of multiple coulomb scattering models in 1 m of helium gas for different scattering angle ranges. The test particles are 30 MeV positrons. The black dots and the green line show the single scattering model which serves as a reference. All the parameterizations are unfortunately inadequate.



on the other hand, none of the parameterizations performs adequately, see Figure 27.9.

We have studied multiple Coulomb scattering in thin silicon at the DESY electron test beam (see figure 27.10) and found that the scattering models in Geant4 describe the core of the scattering distribution well but fail to account for the tails, especially at high momenta. This led us to develop our own scattering model based on Student's t distribution [128], which leads to a better description, see figures 27.11 and especially figure 27.12. The measurement and the model are described in detail in a publication [120]. The set-up at DESY does however not allow for measurements at electron energies much below 1 GeV $/c^2$. We will attempt to further validate the model and also study scattering in gas at upcoming test beam campaigns using the MUPIX telescope.




Figure 27.10: Horizontal scattering angle distribution for 1, 3 and 6 GeV $/c$ electrons with 50 μm silicon target and an incidence angle to the beam of 15° in the device-under-test position. As fit function, a convolution of the shape obtained from a fit to the angular distribution without target and a Student's t -distribution is used. From [120].

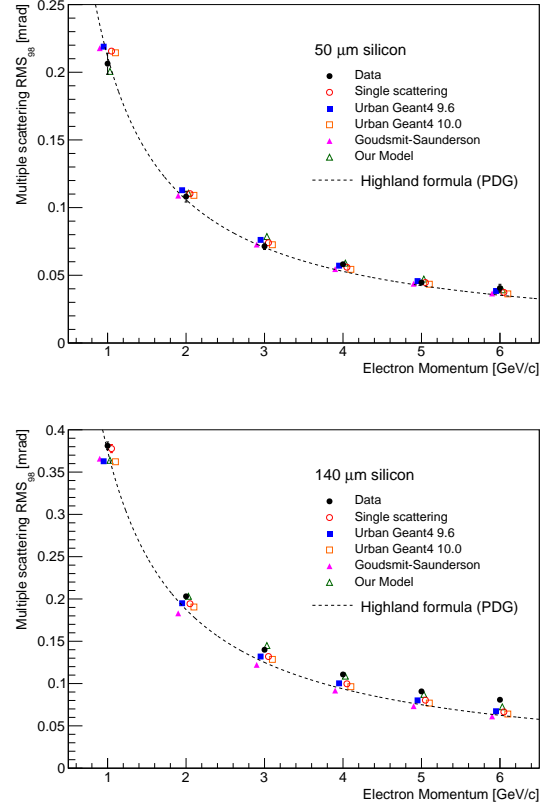


Figure 27.11: Comparison of the RMS of the central 98% of the fitted Student's t distribution versus momentum for various scattering models in Geant4 with our data obtained with a 50 μm silicon target perpendicular to the beam (top) and a 100 μm silicon target at an angle of 45° to the beam (bottom). The data points are slightly offset from their horizontal positions at multiples of 1 GeV $/c$ for better visibility. The Highland parametrisation is also shown for reference. The error bars represent the statistical uncertainty of the fit. From [120].

27.3.2 MUON DECAYS

Michel decay

Geant4 implements the Michel decay including polarization of both the muon and the positron based on Scheck and Fischer ([129] and [130]). The spectra of the neutrinos do not follow the physical distribution, this does however not affect the simulation for Mu3e. Somewhat more worrying is the fact that the Michel matrix element contains




Figure 27.12: Comparison of the Student's t tail parameter ν versus momentum for various scattering models in Geant4 with our data obtained with a 50 μm silicon target perpendicular to the beam (top) and a 100 μm silicon target at an angle of 45° to the beam (bottom). The error bars represent the statistical uncertainty of the fit. Smaller ν corresponds to a larger tail fraction. From [120].

radiative corrections but is not clearly separated from the radiative decay matrix element.

Radiative decay

The radiative decay of the muon in Geant4 was implemented by the TWIST collaboration [131] based on Fronsdal et al. [132]. This code has been adapted for the simulation of the Mu3e experiment using the differential branching fraction provided by Kuno and Okada [3]. As the branching fraction has an infrared singularity, a cutoff value at low photon energies is applied. The value of the cutoff can be chosen by the user. Then, the branching fraction for the radiative de-




Figure 27.13: Integrated branching fraction of the radiative muon decay with internal conversion as function of the sum of the decay particles energies for unpolarized muons, using the matrix element by Djilkibaev et al. [4] (blue line and red dots) or by Signer et al. [133] (green line).

ay with this specific cutoff is scaled with respect to the known branching fraction for photons with 10 MeV or more by the means of Monte Carlo integration. A collinear singularity does not occur as the electron is treated as a massive particle. Because of the infrared singularity, using the accept-reject method to generate events would be very inefficient if the photon energy is generated evenly over the whole range. Therefore, the transformation method is exploited in the generation of the photon energy. The code does not describe the neutrino spectra.

Radiative decay with internal conversion

The radiative decay with internal conversion is simulated using the hit and miss technique on events generated evenly in phase space using the RAMBO code [134] and applying the matrix element. The matrix element has been updated from Djilkibaev et al. [4] to Signer et al. [133] taking now into account the polarization of the muons from the beam. In the case of unpolarized muons, no deviations have been observed applying either the matrix element by Djilkibaev et al. or Signer et al. (see figure 27.13).

The hit and miss technique is very expensive in terms of computation time if the complete phase space is to be simulated (as the matrix element



varies by more than 16 orders of magnitude). This can however be overcome by restricting the simulation to regions of particular interest, e.g. high invariant masses of the visible particles.

Signal

The signal kinematics are highly model-dependent, see chapter 1. If not otherwise noted, we have used three particle phase space distributions in the simulation, following the practice of SINDRUM and earlier experiments.

Special decays

The simulation allows the simulation of overlap decays, where we force more than one muon decay to happen at a single vertex. Thus we can simulate the accidental backgrounds arising e.g. from the overlap of an internal conversion decay and a Michel decay without having to generate in excess of 10^{16} frames.

Cosmic muons

As the detector alignment will rely in part on the high momentum tracks of cosmic ray muons, we have also implemented a cosmic muon generator. For this, we generate a flat distribution of intersection points with a horizontal plane going through the centre of the detector and slightly larger than the detector itself, then create muons according to the cosmic ray parametrisation of Biallass and Hebekker [135]. These muons are then back-extrapolated to 1 m above the detector centre and passed to Geant4 for simulation.

27.4 Time structure

As the Mu3e experiment operates with a quasi continuous beam, the paradigms of *bunch crossing* and *event* from collider experiments do not apply; they have however informed the design of the Geant4 package. In our simulation, particles are assigned a 64bit ID, which is unique over a run and thus conserves identities and mother-daughter relationships across the boundaries of time frames. Before each step of tracking the particle through the detector, it is checked whether the particle has just crossed into the next time slice. If so, its information is stored, its time relative to the time slice adjusted and tracking deferred to the

next slice. Thus we ensure that we correctly treat muons stuck in the target before decaying and decay products crossing frame boundaries while traversing the detector. In order to simulate a steady state, where approximately the same number of muons enter the target and decay, the first 5 ms of simulation running, during which the target is loaded, are thrown away and not used in occupancy or efficiency studies.

Currently not simulated are effects of the 50 MHz structure of the primary proton beam on the time structure seen in the detector; if this would be needed, a measured structure could easily be superimposed on the generation of muons in the simulation framework.

27.5 Detector response

27.5.1 PIXEL DETECTOR RESPONSE

The response of the pixel detector can be simulated by either setting a threshold on the charge deposited or by defining a single hit efficiency, which is then applied by randomly discarding hits. Noise is simulated by randomly creating extra hits at an adjustable rate.

The simulation does currently not include effects of charge sharing between pixels and will thus predominantly generate single pixel clusters. δ -electrons are simulated if they have a range above $50 \mu\text{m}$, the associated (large) clusters should thus be correctly simulated. We plan to implement a more detailed pixel response simulation as soon as corresponding data are available from the production sensor.

Pixel Readout simulation

The readout of the pixel detector is not strictly in order of time stamp and very large clusters of hits can lead to overflows in the sorting algorithms. These effects are partially simulated by treating each column as a queue, into which hits are pushed at creation. A fixed number of hits is then removed from these column queues for every time slot. Hits are time-sorted in a separate program and those that are too far out-of-time are dropped.

We do currently not simulate the dead-time caused by hits stored in the pixel cell - several hundred million pixel cells represent a huge buffer




Figure 27.14: Number of photons arriving at square $250\ \mu\text{m}$ fibre ends (before SiPM coupling) parametrized in deposited energy (E_{dep}) and hit position.

memory, which would increase the memory footprint of the simulation by a large factor and thus preclude parallel execution on standard multi-core machines.

27.5.2 FIBRE DETECTOR RESPONSE

In a first step the response of the scintillating fibres to an incident particle is simulated. Since simulating single photon propagation inside fibres is not feasible in the main simulation, the response of the scintillating fibres is parametrized. The number of arriving photons at both fibre ends can either be parametrized in deposited energy (E_{dep}) and hit position or simply generated according to measured efficiencies. The parametrization (see 27.14) in E_{dep} and hit position is aggregated in a stand alone Geant4 simulations described in section 12.3.

In a second step the SiPM response to the arriving photons is simulated. As described in section 27.1.4, different SiPM configurations are possible. Depending on the chosen configurations, the distribution of photons into the different SiPM cells is modelled. The main parameter for this process is the photon distribution at the fibre ends and propagation in the epoxy layers before the SiPM active layer as well as optical cross-talk. Figure 11.7 shows the photon distribution for round $250\ \mu\text{m}$ fibres. The SiPM response depends on an adjustable photon detection efficiency (PDE) and is mixed with a constantly present

dark rate and its own pixel to pixel cross-talk. The time distribution of the detected events bases on the measured time resolution of fibre ribbons and photon time of flight in the fibres. In a last step pile-up events are merged.

27.5.3 TILE DETECTOR RESPONSE

The tile detector will record the time stamps of the scintillation signals, as well as the energy deposition in the tiles, which is proportional to the number of scintillation photons. In the simulation, the scintillation process and photon propagation is not simulated, in order to maintain a reasonable computation time. Instead, the response characteristics of the tile, including the readout electronics, is parametrised, using the true timestamp and energy deposition of a hit as an input. The response is described by the following parameters: time resolution, energy resolution, jitter of the readout electronics, channel dead-time and energy threshold.

In order for a signal in the tiles to be detected by the readout electronics, a minimum energy deposition is required. This corresponds to the energy threshold of the STiC chip, which is assumed to be roughly $E_{thresh} = 0.1\ \text{MeV}$. Due to the linearised ToT method implemented in the STiC chip, the digitised energy information is approximately proportional to the energy deposition in the tile. The non-linearity of the STiC response is neglected in the current version of the simulation. The energy deposition of consecutive hits (pileup events) which occur within the dead-time of the channel is assigned to the original hit. This reflects the behaviour of the STiC chip.

The energy resolution is dominated by the intrinsic resolution of the STiC chip, which is given by: $\sigma_E = 12\% \cdot E$. The contribution from fluctuations in the number of scintillation photons is $\sigma_E/E \propto 1/\sqrt{N_{pe}}$, which is below 4% for the expected light yield and thus can be neglected.

The channel dead-time is determined by two parameters: the intrinsic dead-time of the STiC TDC and the dead-time related to the time the ToT of the analogue input signal. The TDC dead-time has been measured to be about 42 ns. The ToT of the signal is proportional to the energy and has been measured to be approximately 220 ns for a 1.4 MeV signal. The total channel dead-time is thus given by $\Delta t_{dt}(E) = 220\ \text{ns} \cdot \frac{E}{1.4\ \text{MeV}} + 42\ \text{ns}$.



The time resolution is parametrised by the intrinsic jitter of the STiC chip and the energy dependent resolution of the tile. This STiC jitter is estimated to be $\sigma_{STiC} \lesssim 30$ ps [70]. The energy dependence of the time resolution of the tile is expected to be: $\sigma_t(E) \propto \frac{1}{\sqrt{E}}$, which has been verified using a combined tile-SiPM simulation. The total resolution therefore is parametrised by: $\sigma_t(E) = \frac{\varsigma}{\sqrt{E}} \oplus \sigma_{stic}$. The parameter ς is determined from a time resolution measurement with a tile detector prototype (see section 15.3). This measurement yields a resolution of about $\sigma_t = 56$ ps for an energy deposition of about $E = 1.4$ MeV, which results in a value of ς of: $\varsigma = (56 \text{ ps} \ominus \sigma_{stic}) \cdot \sqrt{1.4 \text{ MeV}} \approx 56 \text{ ps MeV}^{-\frac{1}{2}}$.

As discussed above, the digitised energy information can also include pileup signals, if a consecutive hit occurs within the dead-time. Therefore, this parameter is not suited for the determination

of the time resolution, which only depends on the prompt energy deposition of the original hit. For this reason, a so-called prompt energy deposition is defined, which includes the energy deposition of all hits in a channel occurring within a short time interval $t_{prompt} \ll \sigma_t$. This approach takes into account that in the Geant4 simulation the interaction of a particle with the scintillator (and therefore the energy deposition) is partitioned into one or more discrete hits. As the discrete hits in a certain tile occur within t_{prompt} , the total energy deposition of all hits contributes to the time resolution.

The simulation allows to store the following truth information of the particle hits: time-stamp, energy deposition, impact position and angle, particle type and origin. The energy deposition of secondary particles produced in the scintillator is assigned to the respective primary particle.



RECONSTRUCTION

Alex: Not ready chapters marked with red.

Alex: Describe track reconstruction

28.1 Introduction

A precise track reconstruction of electrons is of highest importance for the identification of the $\mu \rightarrow eee$ decay with a sensitivity of 1 out of 10^{15} (10^{16}) ordinary Michel decays in phase I (II).

Due to the high rate and the resulting high occupancy, especially at phase II of the project with up to 100 tracks per readout frame, the reconstruction algorithm has to deal effectively with the combinatorial background in order to reduce the fake rate, i.e. the rate of wrongly reconstructed tracks, to an acceptable level. The combinatorial problem is not only due to the high rate but also due to the large bending of the low momentum electrons in the strong magnetic field of $B = 1$ T, which, depending on the position and flight direction can make several turns in the detector (recurlers). Hit combinations can span over distances of more than half a meter. Hits of recurling tracks are found on opposite sides of the detector and still have to be correctly combined by the reconstruction program. This is of particular importance for the determination of the flight direction and therefore charge of the particle. Only for a fully reconstructed track the time information provided by the time of flight system can be correctly applied.

As the full detector readout is triggerless, all muon decays have to be fully reconstructed already on filter farm level, setting high demands on

the speed of the online track reconstruction algorithm. A further complication comes from the fact that the track resolution is dominated by multiple scattering (MS) in the silicon pixel sensors and not by the pixel size, in contrast to most other experiments. Therefore, standard non-iterative circle fits of tracks [125] as used in high energy experiments can not be used here.

In order to reduce multiple scattering, the number of sensor layers are reduced to a minimum in the detector design which, unfortunately, also reduces redundancy for track reconstruction. Therefore, the track reconstruction also has to cope with a minimum of information provided by only four sensor layers.

Two different approaches are followed for the Mu3e track reconstruction, the broken line fit [126,127] and the fast linear fit based on multiple scattering [136]. The broken line fit determines hit positions and scattering angles simultaneously and was implemented in 2D [126,137] and recently also in 3D [28,127]. It is based on linearising a preceding circle fit, works non-iteratively and provides the correlations between all fit quantities. The fit, however, requires knowledge of the assignments of hits to tracks from a previous linking step. Therefore, the broken line fit can only be used in the final step of the track reconstruction, also because a previous track fit is required for the linearisation procedure.

The fast three-dimensional MS fit [136] is based on fitting the multiple scattering angles at the middle hit position in a hit triplet combination, see figure 28.1. In this fit, spatial uncertainties of the hit positions are ignored. This is a




Figure 28.1: Schematic view in xy -plane of the triplet made from hits h_0 , h_1 and h_2 with MS in the middle hit.

very good approximation in the case of Mu3e, as the pixel resolution uncertainty given by $\sigma_{\text{pixel}} = 80/\sqrt{12} \mu\text{m}$ is much smaller than the uncertainty from multiple scattering in the corresponding sensor layer. The MS fit requires a detailed knowledge of the material distribution in the detector for the calculation of the scattering angle uncertainty. It minimizes the azimuthal and polar scattering angles at the sensor corresponding to the middle hit and exploits energy conservation¹.

28.2 MS Triplet fit

The triplet is defined as a sequence of three hits produced by a particle with multiple scattering in the middle hit. Figure 28.1 shows a schematic view of a triplet in the xy -plane.

The trajectory of a particle in this system of three hits can be described with two helices. The first helix is defined between points h_0 and h_1 , and the second helix is defined between points h_1 and h_2 , as shown in Figure 28.1. From the relative distances (d_{01} in the xy -plane and z_{01} along the z -axis) between hits h_0 and h_1 the parameters of the first helix, such as opening (bending) angle φ_{01} and inclination angle to the xy -plane λ_{01} , can be derived from the 3D radius r_{01} of the helix. Similar parameters, φ_{12} and λ_{12} , can be defined from r_{12} for the second helix.

¹ The energy loss in the Mu3e experiment is only about 80 keV per sensor layer and can be neglected for track finding.

Using these parameters, the MS angles φ_{MS} in the transverse plane and λ_{MS} in the longitudinal direction are defined as

$$\varphi_{MS} = \delta\varphi - \frac{1}{2} \cdot (\varphi_{01}(r_{01}) + \varphi_{12}(r_{12})) \quad (28.1)$$

and

$$\lambda_{MS} = \lambda_{12}(r_{12}) - \lambda_{01}(r_{01}), \quad (28.2)$$

where $\delta\varphi$ is the angle between vectors h_0h_1 and h_1h_2 .

Assuming no energy loss, the radius $r = r_{01} = r_{12}$ can be defined such that both MS angles only depend on r . The solution to the MS fit is given by minimization of the χ^2 defined as

$$\chi^2(r) = \frac{\varphi_{MS}^2(r) \cdot \cos^2 \lambda + \lambda_{MS}^2(r)}{\sigma_{MS}^2}, \quad (28.3)$$

where σ_{MS} is the MS angle variance from scattering theory [138] and the factor $\cos^2 \lambda$ corrects for the transformation of the MS angle when it is projected on the xy -plane.

28.2.1 LINEARIZATION

In case of small scattering angles Eq. 28.3 can be linearised and solved in a fast non-iterative procedure [136].

Using Taylor expansion and neglecting quadratic and higher order terms the angles φ_{01} and λ_{01} can be written as linear functions of r

$$\varphi_{01} = \varphi_{01}^{(0)} + \alpha_{01} \cdot (r - r_{01}), \quad (28.4)$$

$$\lambda_{01} = \lambda_{01}^{(0)} + \beta_{01} \cdot (r - r_{01}), \quad (28.5)$$

where $\varphi_{01}^{(0)}$, $\lambda_{01}^{(0)}$ and r_{01} correspond to the initial helix solution, and α_{01} and β_{01} are first derivatives of the opening and inclination angles, respectively.

The derivatives α_{01} and β_{01} can be calculated from the equations of the helix trajectory in the transverse plane

$$r_{T,01} \cdot \sin \frac{\varphi_{01}}{2} = \frac{d_{01}}{2} \quad (28.6)$$

and in the longitudinal plane

$$(r_{01}^2 - r_{T,01}^2) \cdot \varphi_{01}^2 = z_{01}^2, \quad (28.7)$$

where $r_{T,01} = r_{01} \cdot \cos \lambda_{01}$ is the transverse helix radius.



Taking the derivatives of Eq. 28.6 and Eq. 28.7 with respect to r_{01} and solving the resulting system of equations for $d\varphi_{01}/dr_{01}$ and $d\lambda_{01}/dr_{01}$ gives

$$\alpha_{01} = -\frac{\varphi_{01}/r_{01}}{\sin^2 \lambda_{01} + \cos^2 \lambda_{01} \cdot \frac{\varphi_{01}/2}{\tan \varphi_{01}/2}} \quad (28.8)$$

and

$$\beta_{01} = -\frac{\tan \lambda_{01}}{r_{01}} \cdot \left(\frac{\alpha_{01}}{\varphi_{01}/r_{01}} + 1 \right). \quad (28.9)$$

Inserting Eq. 28.4 and Eq. 28.5 in Eq. 28.1 and Eq. 28.2 gives

$$\varphi_{MS}(r) = \varphi_{MS}^{(0)} + \alpha \cdot r \quad (28.10)$$

and

$$\lambda_{MS}(r) = \lambda_{MS}^{(0)} + \beta \cdot r, \quad (28.11)$$

where $\alpha = -(\alpha_{01} + \alpha_{12})/2$ and $\beta = \beta_{12} - \beta_{01}$.

Finally, Eq. 28.3 reduces to

$$\chi^2(r) = \chi_{min}^2 + \frac{(r - r_{min})^2}{\sigma_r^2}, \quad (28.12)$$

where r_{min} is the solution to the χ^2 minimization, σ_r is an uncertainty estimate on r_{min} and χ_{min}^2 is the value of χ^2 at $r = r_{min}$

$$\sigma_r^2 = \frac{\sigma_{MS}^2}{\alpha^2 \cos^2 \lambda + \beta^2}, \quad (28.13)$$

$$r_{min} = -\frac{\sigma_r^2}{\sigma_{MS}^2} \cdot \left(\varphi_{MS}^{(0)} \alpha \cos^2 \lambda + \lambda_{MS}^{(0)} \beta \right). \quad (28.14)$$

The solution to the triplet fit in linear approximation can be solved starting from any known trajectory. The simplest approach is to use a circular trajectory in the xy -plane. In this case the transverse radius r_T is the same for both helices and is given by

$$r_T = \frac{d_{02}}{2 \cdot \sin \delta\varphi/2}, \quad (28.15)$$

where d_{02} is the distance between hits h_0 and h_2 . From the transverse radius the opening and inclination angles can be derived

$$2 \cdot \sin \varphi_{01/12}/2 = \frac{d_{01/12}}{r_{T,01/12}}, \quad (28.16)$$

$$\tan \lambda_{01/12} = \frac{z_{01/12}}{r_{T,01/12} \cdot \varphi_{01/12}}. \quad (28.17)$$

It is possible to improve (update) the fit result by performing the linear fit in several iterations. In this case the updated $\varphi_{01/12}$ are calculated using Eq. 28.4 and the current value of r , then $r_{T,01/12}$ and $\lambda_{01/12}$ are derived using Eq. 28.16 and Eq. 28.17.

The fit that starts with the circular solution is used during normal (initial) track finding, and the update step is used when fitting with energy loss correction, as described in ??.

28.3 MS Track fit

The linear MS triplet fit is used as basis for the reconstruction of tracks in the pixel detector.

The tracks are represented as sequence of triplets, where each triplet is composed of hits from subsequent layers, such that consecutive triplets share hits from two layers, as shown in Fig. ??. The total number of triplets in the track is $n_t = n_h - 2$, where n_h is the number of hits.

The track is fitted by minimizing the total χ^2 , that is defined as the sum of χ^2 of individual triplets

$$\chi^2 = \sum_t \frac{\varphi_{MS}^2(r) \cdot \cos^2 \lambda + \lambda_{MS}^2(r)}{\sigma_{MS}^2}, \quad (28.18)$$

where all variables in the sum correspond to their respective triplet t .

Using the linear approximations for φ_{MS} and λ_{MS} in Eq. 28.18 the solution for the minimization of χ^2 can be represented (exactly) as weighted average of the linear solutions for individual triplets:

$$\frac{1}{\sigma_r^2} = \sum_t \frac{1}{\sigma_{r,t}^2} \quad (28.19)$$

and

$$r_{min} = \sum_t \frac{r_{min,t}}{\sigma_{r,t}^2}, \quad (28.20)$$

where $\sigma_{r,t}$ and $r_{min,t}$ are the solution for the triplet t .

The simplest implementation of the fit does not take into account spatial uncertainties or energy loss. However, it is possible to update the fit to take into account (partly) these effect.

28.3.1 PIXEL SIZE

For the inner and outer double layers, where distances between hits are of the order of 10 mm, the




Figure 28.2: Schematic view of the triplet in the transverse xy -plane (left) and longitudinal zs -plane (right). Hits h_0 , h_1 and h_2 are produced on pixels (thick red lines). The effective pixel size depends on the incident angles τ and λ at which the particle hits the sensors (pixels). The correction to the reconstructed scattering angle in the middle hit h_1 depends on the offsets x from the pixel centers to the real positions at which the particle hits the pixels and on the path lengths between hits.

contribution of the pixel size becomes comparable with average scattering angles. In addition, the effect of pixel size may become significant when calculating uncertainties on the reconstructed angles that are used in the vertex fit.

To take into account the pixel size the σ_{MS} for each triplet is modified separately for φ and λ . Figure 28.2 shows a schematic view of the triplet in the transverse (left) and longitudinal (right) planes. The track hits h_0 , h_1 and h_2 are located at the centers of the pixels (red lines). The effective pixel sizes (blue lines) depend on the incident angles τ and λ at the corresponding hit positions.

In the longitudinal plane the distances x_0 , x_1 and x_2 between the centers of the pixels and the real hit positions introduce an additional variance in MS angle λ

$$\begin{aligned} \sigma_{px,\lambda}^2 &= \frac{\sigma_{px}^2}{s_{01}^2} \cdot 2 \cdot \cos^2 \lambda_{01} \\ &\quad + \frac{\sigma_{px}^2}{s_{12}^2} \cdot 2 \cdot \cos^2 \lambda_{12} \\ &\quad - \frac{2 \cdot \sigma_{px}^2}{s_{01} \cdot s_{12}} \cdot \cos \lambda_{01} \cdot \cos \lambda_{12}, \end{aligned} \quad (28.21)$$

where λ_{01} is the first helix inclination angle at hits h_0 and h_1 , and λ_{12} is the second helix inclination angle at hits h_1 and h_2 . The first and second terms describe contributions corresponding to the first and second helices. The third term is the correlation between the helices due to the same offset x_1 at the middle hit.

Similarly, the variance in the transverse plane is

$$\begin{aligned} \sigma_{px,\varphi}^2 &= \frac{\sigma_{px}^2}{s_{T,01}^2} \cdot (\cos^2 \tau_{01} + \cos^2 \tau_{10}) \\ &\quad + \frac{\sigma_{px}^2}{s_{T,12}^2} \cdot (\cos^2 \tau_{12} + \cos^2 \tau_{21}) \\ &\quad - \frac{2 \cdot \sigma_{px}^2}{s_{T,01} \cdot s_{T,12}} \cdot \cos \tau_{10} \cdot \cos \tau_{12}, \end{aligned} \quad (28.22)$$

where τ_{01} and τ_{10} are the first helix inclination angles at hits h_0 and h_1 , and τ_{12} and τ_{21} are the second helix inclination angles at hits h_1 and h_2 .

The triplet χ^2 can be adjusted to include the effect of pixel size by adding $\sigma_{px,\varphi}^2$ and $\sigma_{px,\lambda}^2$ to Eq. 28.3

$$\chi^2 = \frac{\varphi_{MS}^2}{\sigma_{MS}^2 / \cos^2 \lambda + \sigma_{px,\varphi}^2} + \frac{\lambda_{MS}^2}{\sigma_{MS}^2 + \sigma_{px,\lambda}^2}. \quad (28.23)$$

28.4 Track reconstruction

The number of hits produced by a particle depends on the single hit efficiency. In the current reconstruction procedure the efficiency is assumed to be 100%. There are two types of tracks that can be reconstructed in the detector: short tracks that are reconstructed from hits in four layers, and long tracks that bend in the magnetic field and produce additional hits when they curl back. The long tracks are an improved version of the short tracks, as short tracks are used as seeds to reconstruct long tracks.

Note that there is a big qualitative difference between the short tracks and long tracks. The bending radius of long tracks can be measured with much higher precision due to the larger bending in the magnetic field. This leads to very different average momentum resolutions for the different detector setups.

The long tracks can be further divided in two groups: long 8-hit and 6-hit tracks. Long 8-hit tracks are produced by particles that curl back to



the central part of the detector, effectively producing two short tracks that can be linked together. One particle can produce several 8-hit segments until it finally leaves the central part of the detector. In the current reconstruction program the long tracks corresponding to different turns of the particle are not linked together and tracks with more than 8 hits are not reconstructed (the tracks are kept separate).

Particles that leave the central part of the detector after producing a 4-hit track can be reconstructed as long 6-hit tracks if they hit the two outer layers on their way back to the detector.

Figure ?? illustrates short and long tracks.

In phase Ia without recur stations, most tracks are reconstructed with only four hits and a momentum resolution of the order of 1 MeV. Adding two inner recur stations in phase Ib considerably improves the acceptance and increases the number of reconstructed long 6-hit tracks. The reconstruction of the recurring long tracks, which provide much better momentum resolution of the order of 0.2 MeV, is crucial for the success of the experiment (see 2).

28.4.1 DETECTOR AND SIMULATION

For this study of the track reconstruction software the full Geant4 simulation of the phase Ib detector geometry is used (see chapter 27). The beam intensity is set such that on average 6.6 μ^+ decays are observed in one 50 ns frame, within the target region ($v_z < 60$ mm and $v_r < 20$ mm). The study of the tracking acceptance and efficiency is performed on a sample of Michel decays.

During the simulation each hit produced by a particle is assigned a track ID and a hit ID. This information allows to identify if the track is reconstructed correctly or if the track is made from a wrong combination of hits, such as noise or hits from different particles or from the same particle but in a wrong sequence.

As the particle can make several turns in the magnetic field, such that on each turn it passes through the detector, the number of reconstructed tracks (short or long) corresponding to the same particle may be two or more. These tracks correspond to different segments of the particle trajectory and can be arranged sequentially.

For efficiency studies only tracks corresponding to the first segment of the particle trajectory are used, as these segments have to be used during

the vertex fit. These tracks are selected using MC information. For normalization the sets of reconstructable tracks of the same type (short or long) are used. The reconstructable tracks are formed by dividing a set of particle hits in sequences of consecutive 4, 6 or 8 hits, that correspond to short or long tracks.

28.4.2 TRIPLET SEEDS

As a first step, the triplets that will be used as seeds for subsequent reconstruction of short tracks are constructed. These triplets are made by combining hits from the first three layers. The number of triplet combinations that have to be fitted for each frame scales as n^3 , where n is the average number of hits in one layer. For $n = 10$ this leads to 1000 hit combinations that have to be fitted in each frame. To reduce the number of fits to be performed, first, geometrical selection criteria are applied on the hit combinations.

These criteria include the distance between hits along the z -axis, the angular distance relative to the center of the coordinate system (center of the detector) in the transverse plane, and the alignment of hits when projected on the rz -plane.

First, hits in the first two layers are combined if the distance between hits along the z -axis is $|z_1 - z_0| < 30$ mm and the angle difference $|\varphi_1 - \varphi_0| < 0.8$ rad. Then, a hit from the third layer is added if the angles satisfy $|\varphi_2 + \varphi_0 - 2 \cdot \varphi_1| < 0.8$ rad and the alignment in rz -plane is $|\delta_{rz}| < 1$, where

$$\delta_{rz} = \frac{z_2 - z_1}{r_2 - r_1} - \frac{z_1 - z_0}{r_1 - r_0}. \quad (28.24)$$

The triplets that lie on a circle in the transverse plane with $30 < r_T < 250$ mm are subsequently fitted. The soft cut of $\chi^2 < 100$ is applied and resulting fitted triplets are checked, such that at each hit the tangent to the trajectory would point outward with respect to the corresponding sensor. The cut $\chi^2 < 100$ reduces the number of wrong combinations by a factor 3 with only 0.3% loss in truth triplets.

Figure 28.3 and Figure 28.4 show the distance $z_1 - z_0$ and the alignment δ_{rz} for truth and fake combinations after all selections and the subsequent fit.

Figure 28.5 shows the χ^2 distribution of the fitted triplets. Note that the χ^2 distribution peaks at zero as the number of degrees of freedom $NdF = 1$ for the fitted triplets.




Figure 28.3: Distance between first pair of hits along z -axis for truth (black line) and wrong (red line) triplet combinations.




Figure 28.4: Alignment δ_{rz} in the rz -plane of hits for truth (black line) and wrong (red line) triplet combinations.




Figure 28.5: Distribution of the fit χ^2 for truth (black line) and wrong (red line) triplet combinations.

For wrong combinations the distributions are wider than for truth combinations. With all selections applied the number of selected triplets is about 30 per frame with the fake rate $n_{fake}/n_{truth} \approx 2$. This is a factor 100 smaller compared to the number of all possible combinations.

Figure 28.6 shows the reconstructed triplets for one frame in the transverse view of the detector.

28.4.3 SHORT TRACKS

Short 4-hit tracks are reconstructed by combining triplet seeds with hits from the fourth layer of the central station, thereby using hits from all four layers. To improve performance, the number of triplet-hit combinations is reduced by using the estimated position of the fourth hit. The position of the hits is estimated by the propagation of the triplet trajectory to the cylindrical approximation of the fourth layer. All hits that lie within 50 mrad in the transverse plane and 5 mm along the z -axis from the estimated position of the fourth hit are combined with the triplet and subsequently fitted.

Figure 28.7 shows the χ^2 distribution for truth and fake short tracks. Note that the χ^2 distribution peaks around 1 as the number of degrees of freedom $NdF = 3$.




Figure 28.6: Reconstructed triplets for one frame in transverse view of the detector. Blue dots are sensor hits. Black lines are truth triplets, while red lines are wrong combinations.




Figure 28.8: Reconstructed short tracks for one frame in transverse view of the detector. Blue dots are sensor hits. Black lines are truth short tracks.




Figure 28.7: Distribution of χ^2 for truth (black) and fake (red) short tracks.

To reduce the fake rate the cut $\chi^2 < 32$ is used. In addition the inclination angle at the fourth hit is checked, such that the trajectory points outward with respect to the sensor surface. This selection suppresses combinations of triplets with hits that are produced by particles when they curl back.

Figure 28.8 shows reconstructed short tracks in the transverse plane for one frame.

The geometrical cuts, used to select combinations of hits for subsequent fitting, mainly affect the edges of acceptance. The χ^2 affects the overall normalization. Figure ?? shows the distribution of reconstructed short tracks relative to the total number of particles that produced at least one hit in each layer.

28.4.4 DOUBLE HITS

The geometry of the layers allows a particle to hit two adjacent sensors in the same layer. Consequently, during short track reconstruction, several tracks for the same particle can be reconstructed that differ only in the combination of hits as illustrated in Figure ??.

These tracks have to be combined or otherwise removed. First, a procedure called merging is used to combine tracks. This merging is an iterative procedure that looks at pairs of tracks that share




Figure 28.9: The triplet is modified such that first helix is defined by hits h_0 and $h_{1,1}$, second helix is defined by hits $h_{1,2}$ and h_2 .

all hits but on one layer, where the hits are from adjacent sensors. These hits are combined into doublet hits. The merging can produce tracks with up to twice the number of initial hits.

To take into account doublets in the fit the triplets are modified as shown in Figure 28.9. The helices in triplets are redefined to their corresponding hit pairs, which is possible as the helix is completely defined by the radius and relative hit positions.

The hits are combined in doublets only if they align with the original triplet trajectory. The alignment is checked by comparing pointing from one hit to the other, which should be less than the pixel size.

A particle may produce several hits in the same sensor, or there may be several unrelated hits in the vicinity of a true hit. In both cases several tracks can be reconstructed with one track per hit. Only the track with minimum χ^2 is retained from these tracks. This procedure is applied only to hits that are not the first or last hit of the track, where the spread between hits may be large.

Figure 28.10 shows the distribution of the total number of hits per short track after the merging procedure. Note that doubles in triplets are effectively treated as single hits and no additional helices are added to the track. In particular the




Figure 28.10: Total number of hits per short track, taking into account double hits.

n -hit notation defines the number of single/double hits and not the total number of hits.

28.4.5 TRACK INTERSECTIONS

Due to the small pixel size the probability for a pixel being hit by two particles is negligible. If two tracks share the same hit, then this hit is either noise or a hit produced by only one particle. In a perfectly reconstructed frame each hit is used at most once, however, due to the presence of fake tracks some hits are used by several tracks.

Figure 28.12 shows the distribution of the number n_s of tracks that share at least one hit with a given track. Figure 28.11 shows a schematic view of two truth tracks that share hits h_1 and h_2 with a fake track. For truth tracks this number is expected to be close to zero and larger than zero for wrong combinations. The truth tracks with $n_s = 1$ have usually one intersection with a fake track or are due to a failure to merge the tracks. The tracks with $n_s = 2$ and the tail with $n_s > 2$ are due to a large number of fake combinations reconstructed for tracks with $\lambda \approx 0$, where there is a large chance to reconstruct tracks with a hit from another segment of the same particle track.

This information can be used to reduce the number of wrong combinations. A particular configuration of tracks is used in which one track has an intersection with at least two tracks with




Figure 28.11: Schematic view of two truth tracks which share hits h_1 and h_2 with a fake track. The ... (???)




Figure 28.13: Distribution of n_s after removing tracks that have intersection with at least 2 tracks with $n_s = 1$.

$n_s = 1$, meaning that they intersect only with this one track. In this configuration the track is regarded as a wrong combination and is removed. This procedure allows to reduce the number of wrong combination by 27% with no loss of truth tracks. Figure 28.13 shows the distribution of n_s after the procedure.

In addition to the suppression of wrong combinations, n_s can be used to select a clean sample of tracks. Selecting tracks with $n_s = 0$ produces samples with a fake rate of less than 10^{-3} .

28.4.6 SHORT TRACKS PERFORMANCE

As noted above any track that produced at least one hit in each layer can in principle be reconstructed. The reconstruction “efficiency” in this case is only limited by geometrical and χ^2 cuts. Figure 28.14 shows the number of reconstructed tracks relative to the number of reconstructables ones.

There are several features in this plot. In the top right corner there are tracks with momentum $p > 53 MeV/c$. These are electrons produced by muons decaying in flight. The cut off at low momentum is due to the acceptance, the drop in efficiency close to the cut off is the result of geometrical cuts. The efficiency reaches its maximum at 95% due to the χ^2 cut.




Figure 28.12: Distribution of number n_s of tracks that share at least one hit with a given track.




Figure 28.14: Ratio of reconstructed short tracks to reconstructable as function of momentum p and angle λ .

Figure 28.15 shows the difference between reconstructed and truth momentum for short tracks. The bias in reconstructed momentum is due to the energy loss in the target, fibres and sensor material. However, it can be neglected for short tracks, as the bias is much smaller than the resolution.

Figure 28.16 shows the resolution as function of the truth momentum of short tracks. The relative resolution $\sigma_p/p \approx 4\%$ is approximately constant.

28.4.7 LONG TRACKS

Long 8-hit tracks are reconstructed from pairs of short tracks. First, the inner long tracks, corresponding to particles crossing the center of the detector, are reconstructed. Then, outer long tracks, particles that curl back, are reconstructed from unused short tracks or a combination of unused short tracks with inner long tracks.

As the path length between short tracks for inner long tracks is much smaller compared to outer long tracks, the number of wrong combinations for the inner tracks is small. Subsequently, combining short tracks with inner tracks allows to reduce the number of wrong outer track combinations.

The number of wrongly reconstructed long 8-hit tracks is much larger than for short tracks. This means that the wrong combinations are produced from incorrect matching of short tracks, and




Figure 28.15: Difference between reconstructed and truth momentum for short tracks at $34 < p_{mc} < 36$ MeV.




Figure 28.16: Momentum resolution σ_p as function of truth momentum p_{mc} of short tracks.




Figure 28.17: Distance along z -axis between first and last hit of outer long track for truth (black line) and wrong (red line) combinations. The blue histogram corresponds to reconstructable truth tracks.

the contribution of incorrectly assigned individual hits is small. These wrong assignments are mainly observed for tracks produced by particles which make many turns in the detector. Figure 28.17 shows the distance along the z -axis between the first and last hit of outer long tracks for truth and wrong combinations.

Long 6-hit tracks are reconstructed from short tracks and pairs of unused hits in the outer layers, including the recurl stations. These tracks provide good momentum resolution, similar to long 8-hit tracks, and have a small fake rate.

Out of all reconstructed short tracks 85% are used for long tracks, 41% are long 6-hit tracks and 44% are long 8-hit tracks (inner and outer). The remaining 15% correspond to particles that on the way back to the detector go into gaps between stations, fly past recurl stations ($|\lambda| > 1.2$ rad) or do not pass matching criteria. Figure 28.18 shows the number of reconstructed long 6-hit tracks relative to the number of short tracks.

Figure 28.19 shows the momentum resolution as function of the truth momentum for 6-hit long tracks. For these tracks the best resolution $\sigma_p = 90$ keV is reached at $p = 20$ MeV. At this momentum two hits in the outside layer lie on oppo-




Figure 28.18: Number of reconstructed long 6-hit tracks relative to number of short tracks as function of momentum p and inclination angle λ .

site sides of reconstructed tracks in the transverse plane, which provides the optimal geometry to determine the momentum.

As the momentum resolution for long tracks is determined by the outer part of the reconstructed tracks, the resolution for long 6- and 8-hit tracks is approximately the same.

28.4.8 ENERGY LOSS

For long tracks the momentum resolution becomes comparable with the total energy loss observed as a shift in the reconstructed momentum. To correct for this shift the energy loss have to be taken into account in the track fit.

In the standard fit all triplets have the same radius r , calculated as weighted average of individual triplet fits. In presence of energy loss we are interested in radius at the beginning of the tracks or radius of first helix of first triplet. With each hit along the trajectory this radius have to be adjusted according to energy loss at corresponding hit. Given radius r_i of first helix of triplet i the scattering angles can be written as

$$\varphi_{MS} = \delta\varphi - \frac{\varphi_{12}(r_i - \delta r_i) + \varphi_{01}(r_i)}{2} \quad (28.25)$$

and

$$\lambda_{MS} = \lambda_{12}(r_i - \delta r_i) - \lambda_{01}(r_i), \quad (28.26)$$




Figure 28.19: Momentum resolution σ_p as function of truth momentum p_{mc} of 6-hit long tracks.

where $r_i - \delta r_i$ is the radius of second helix and δr_i corresponds to the most probable energy loss at the middle hit.

By inserting $r_i = r - \Delta r_i$, where Δr_i is cumulative energy loss of previous triplets, both Eq. 28.25 and Eq. 28.26 can be reduced to Eq. 28.10 and Eq. 28.11 with constants $\varphi_{MS}^{(0)}$ and $\lambda_{MS}^{(0)}$ redefined such that

$$\varphi_{MS}^{(0)} \rightarrow \varphi_{MS}^{(0)} - \frac{\alpha_{12}\Delta r_{i+1} + \alpha_{01}\Delta r_i}{2} \quad (28.27)$$

and

$$\lambda_{MS}^{(0)} \rightarrow \lambda_{MS}^{(0)} - \beta_{01}\Delta r_i + \beta_{12}\Delta r_{i+1}. \quad (28.28)$$

Starting from the standard fit, the energy loss corrections δr_i are estimated for all triplets. These corrections allow to calculate updated $\varphi_{MS}^{(0)}$ and $\lambda_{MS}^{(0)}$ which are used for triplets and track solutions.

The energy loss δE is calculated according to material distribution in the sensor ladder and trajectory path length

$$\delta E = \sum_m \frac{dE_m}{dX} \cdot X_m, \quad (28.29)$$

where dE_m/dX is energy loss per rad.length for given material m and X_m is the rad.length corresponding to trajectory path length.




Figure 28.20: Energy loss in a second layer.

Figure 28.20 shows the truth energy loss in the second layer.

The corrections δr are calculated for energy loss dE as

$$\delta r = f_E \cdot \frac{dE}{0.3 \cdot B}, \quad (28.30)$$

where $f_E = 0.75$ is scale factor that corrects for the position of energy loss peak maximum relative to average energy loss \bar{E} and $B = 1$ T is the magnetic field.

Figure 28.21 shows ratio of estimated energy loss δr to truth energy loss in second layer. Note that distribution peaks at one. The tail to left is due to long tail at higher energy losses.

For fibre tracker the energy loss is about factor five larger than typical energy loss in sensor ladders. As the fibre planes are located close to third layers the energy loss in fibres is added to energy loss in third layer.

Figure 28.22 shows the difference between reconstructed and truth momentum at first helix of first triplet for long tracks with (black) and without (red) energy loss correction.

28.5 Vertex reconstruction

A vertex fit will be used to precisely reconstruct the position of the muon decay for signal events. The vertex fit intrinsically checks the consistency of the assumption that all three signal candidate




Figure 28.21: Ratio of estimated energy loss to truth energy loss in a second layer.




Figure 28.22: Difference between reconstructed and truth momentum at first helix of first triplet- for long tracks with (black) and without (red) energy loss correction.

tracks originate from the same vertex. A common vertex fit of the three candidate electrons allows a suppression of the combinatorial background of Michel decays by a factor of about 100 due to the very good pointing resolution of the pixel detector, which is mainly given by the multiple scattering at the first sensor layer. Instead of using a common vertex fit, also the distance between target impact points of reconstructed and extrapolated tracks can be used to discriminate signal against background.

28.6 Fibre linking

The fibre tracker provides timing information that will be used to reduce the fake rate during track reconstruction and help to resolve ambiguities in the sign of track momentum.

In the current reconstruction software fibre clusters are assigned to short tracks. These fibre clusters are composed from individual fibre hits that are combined based on the proximity in space between hits and their time difference: adjacent fibres with a time difference of less than 2 ns are combined into clusters.

Associating clusters to short tracks is based on the distance in the transverse plane from the cluster to the track trajectory. The distance is calculated using the parameterization of the track trajectory and the cluster position. The cluster closest to the track within a 1 mm search radius is assigned to the track, resulting in an assignment efficiency of 95 % to truth tracks. Figure 28.23 shows the distance between tracks and associated clusters.

Of all truth tracks with an associated fibre cluster, 1.5 % are faulty, where the cluster contains hits not produced by the particle.

The time information from fibres can be used to resolve the ambiguity in the flight direction (and thus charge) of long 8-hit tracks. Figure 28.24 shows the time difference between two clusters assigned to short tracks of the long 8-hit tracks as a function of the path length along the track from the first cluster to the second.

The positive difference corresponds to tracks with correct (physical) direction and the negative difference to those with wrong direction.

Figure 28.25 shows the speed of the particle in ns/m estimated from the time difference and flight path differences between fibre hits. For truth




Figure 28.23: Distance from short track to the assigned fibre cluster.




Figure 28.25: Reconstructed speed of the particle in ns/m for truth (black) and fake (red) long 8-hit tracks. The speed is estimated from the time difference and flight path differences between fibre hits.




Figure 28.24: Time difference between clusters assigned to long 8-hit track as function of distance along trajectory.

tracks (black) the flight speed is about 3.3 ns/m which is close to light speed. For wrong combinations (red) the time difference between fibre hits is larger, as short tracks of the long track correspond to different cycles of the particle, resulting in smaller reconstructed speed

The current implementation of the fibre reconstruction does not take into account the position along the z -axis, where a track hits a fibre. Correcting for the light propagation delay in fibres could improve the time resolution. In addition, it is possible to implement an algorithm that will use time information during long track reconstruction, further reducing fake rates.

28.7 Tile linking

The tile detector provides time measurement with resolution of about 100 ps. This information could be used in vertex fit to reduce number of wrong combinations. In addition the vertex fit could be improved by using time information directly in the fit.

The tiles are located in recur stations, consequently, only particles reconstructed as long 6-hit tracks can produce tile hits. The linking of tile




Figure 28.26: Number of linked tile hits for truth tracks that point to at least one tile. 98.5% of tracks have at least one linked tile hit.

hits is performed for each long 6-hit track with last hit in recurv station. The track trajectory is extrapolated to surface of tile detector and all tile hits within 3×3 mask are linked to the tracks.

In the current implementation of tile linking 98.5% of truth long tracks, that point to at least one tile, have linked tile hits. Figure 28.26 shows distribution of number of linked hits.

If at least one tile hit in a cluster is no produced by the particle the assignment is considered faulty and is about 1.3%.

Taking the average cluster time and correcting for the flight time from the first hit the approximate vertex time could be estimated. Figure 28.27 shows the difference between true and estimated vertex time for truth tracks originating from target region.

The resulting vertex time resolution is about 80 ps. The small offset is due to flight time from the vertex to the first hit.

28.8 Summary

The track reconstruction shows good performance.

Effect of noise on fake rate for different classes of tracks. Track reconstruction in presence of sensor




Figure 28.27: Difference between true and estimated vertex time for truth tracks originating from target region. The estimated time is the average time of tile cluster corrected for flight time from first hit.

inefficiency and cross frame reconstruction. Column fibre readout. Improve handling of fibre MS.



ONLINE EVENT SELECTION

The online data rate needs to be reduced such that only physically relevant event candidates are stored to disk. As in the offline analysis, event selection in the filter farm relies on the coincidence of three tracks in time, a common vertex, and on their kinematics. Especially for high rate running in phase II, a timing coincidence is not sufficient to reduce the data rate by three to four orders of magnitude. Thus an online track reconstruction will be required. To this end, a simple version of the fast linear fit based on multiple scattering (see chapter 28) is implemented on GPUs for quick track fitting. In addition, events with two positive and one negative electron track are checked with respect to a common vertex. All these selections will be applied on a frame by frame basis on the individual farm PCs and the selected events will be merged into the global data stream afterwards, see figure 29.1 for an overview. The technical implementation of the event filter is described in sections 24.7 and 24.8.

29.1 Selection Cuts

For the online reconstruction, hits from the central station of the pixel detector only are considered since time information from the tiles and fibres is not available (phase IA) or hard to reconstruct online (phase IB) to match recurlers correctly. Furthermore it is not necessary for a first selection. Combinations of hits from the first three detector layers are matched to form triplets. However, before the actual fitting procedure, a number of selection cuts are applied to reduce combinatorics.

In the following, the indices label the i^{th} layer of the detector, \mathbf{X}_i denotes the vector from the origin to the position of a hit h_i in the transverse plane (see figure 28.1), \sphericalangle the opening angle between two such vectors, and $r_i = |\mathbf{X}_i|$.

- $\sphericalangle(\mathbf{X}_0, \mathbf{X}_1) \leq 0.8 \text{ rad}$
- $\sphericalangle(\mathbf{X}_1, \mathbf{X}_2) \leq 0.8 \text{ rad}$
- $|z_1 - z_0| < 30 \text{ mm}$
- $\frac{|z_2 - z_1|}{r_2 - r_1} - \frac{|z_1 - z_0|}{r_1 - r_0} \leq 1.0$

The first two cuts ensure proximity of the hits in the x-y plane. Figure 28.3 displays the difference in the z coordinate between the first and second hit which is ensured by the third cut. The last cut requires only a small change of direction in the r-z plane depicted in figure 28.4. Applying these selections reduces the number of possible hit combinations by a factor of 50 at a muon stopping rate of 10^8 muons/s. For a first performance test, the cuts are applied on the GPU. In the final setup, they will be simplified to integer value comparisons and will be executed on the PCIe FPGA boards.

29.2 Multiple Scattering Fit

The fitting of triplets is non-recursive and linear and can be done in parallel for all hit combinations. Therefore, it is an ideal candidate for parallelization on GPUs. With their many computing kernels but little memory space they perform well at tasks with high combinatorics where many

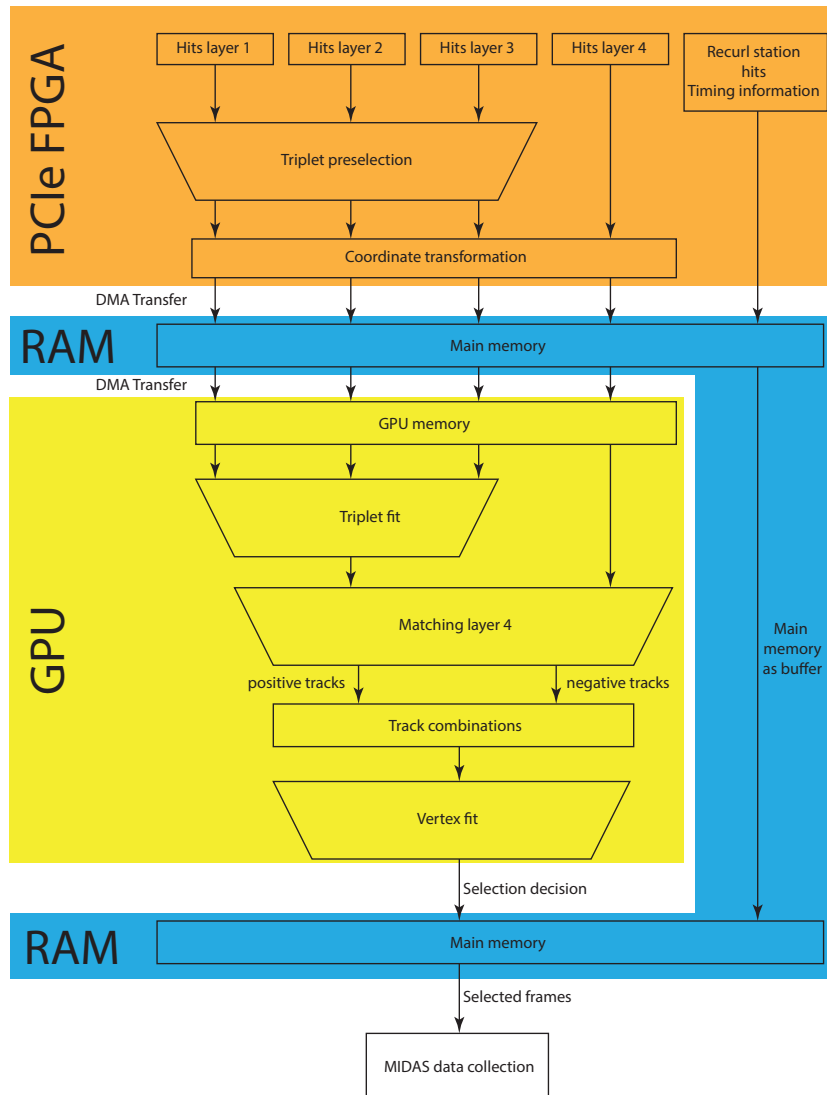


Figure 29.1: Flow diagram of the online reconstruction software and firmware.

computations are performed on the same memory content. On NVIDIA GPUs, tasks are performed in parallel on a grid of compute nodes on the software level. On the hardware level, several streaming multiprocessors (SMs) are available to perform tasks in parallel (the GTX 980 GPU e.g. has 16 streaming multiprocessors with 128 compute cores each). On each such SM, 32 threads are executed at the same time, doing exactly the same calculation. This means that a divergence of the different threads due to a conditional clause in the algorithm executed by the thread causes part of

the threads to stall while waiting for the others to finish. Consequently, divergence should be minimized in code running on a GPU which is one reason to apply the selection cuts already at the FPGA level. Due to the partitioning of the SMs into chunks of 32 threads, the number of threads is best chosen to be a multiple of 32 to avoid inherently idle cores on the SMs. Each of these threads performs the multiple scattering fit (see section 28.2) for one combination of three hits. During the fitting procedure, the 3D radius r of the helix is calculated for a minimal value of χ^2



(eqn 28.3) which is a function of the linearized scattering angles $\Phi_{MS}(r)$ and $\lambda_{MS}(r)$ (see eqns 28.4 and 28.5). Good fits are selected by requiring $\chi^2 \leq 100$ and $30 < r_T < 250$ mm. After this selection, 25 % of the selected hit combinations remain for the next steps.

Depending on the incident muon rate, the number of tracks per time frame and therefore also the performance of the triplet fit on the GPUs varies. For a muon rate of 10^8 Hz, $\mathcal{O}(10)$ hits per layer and time frame lead to $\mathcal{O}(10^3)$ combinations, whereas a rate of $2 \cdot 10^9$ Hz with a few hundred hits per layer results in $\mathcal{O}(10^6)$ combinations. Higher combinatorics present a challenge in processing all triplet candidates quickly, nevertheless, current GPU's compute nodes can be operated close to full load. For example, the compute utilization on a GTX 680 of the GPU's SMs is 3 % at the lower rate, limited by memory accesses, at the higher rate, the compute utilization increases to 80 %. Measuring the time for fitting only (without cuts on χ^2 or r), $3 \cdot 10^7$ fits/s are possible at the low rate and $5 \cdot 10^9$ fits/s were achieved at the higher rate.

29.3 Vertex Selection

For each triplet passing the χ^2 and radius cuts, the track is extrapolated to the fourth detector layer (see figure ??). If a hit exists within a certain

transverse radius and distance in z , it is assigned to the triplet, otherwise the triplet is discarded. This selection reduces the combinatorics by a factor of four. Finally, the charge of the particle is derived from the track curvature and all combinations of two positive tracks and one negative track are examined with respect to a common vertex. For this, the tracks will be extrapolated to the target and the proximity in the x - y -plane as well as in z will define whether or not the three tracks originate from one vertex. The proximity of this vertex with respect to the target location serves as an additional selection criterion. After the vertex selection, the data will be reduced by a factor of 1000, which is enough to store the data on disks for offline analysis.

29.4 Current Status

The implementations of the geometrical selection cuts, the triplet fit and the propagation to the fourth layer are working on the GPU and the fit have been optimized with respect to performance. In terms of GPU code, the remaining tasks are to efficiently store positive and negative track candidates and to implement the vertex fit. The selection cuts need to be ported to the FPGA and the coordinate transformation has to be applied. The DMA transfer has been tested successfully as described in section 24.7.1.

PART IX



CALIBRATION AND ALIGNMENT



MILESTONES FOR THE DETECTOR CALIBRATION

1. Develop, test and run an automatic pixel threshold calibration for the full detector.
2. Develop, test and run a fibre timing calibration.
3. Develop, test and run an automatic SiPM overvoltage calibration.
4. Develop, test and run a tile timing calibration.
5. Prepare the detector alignment software, test it with simulated data.
6. Perform track based alignment of the phase Ia detector.
7. Enhance the alignment software to include fibres, tiles and recurl stations.
8. Perform track based alignment of the phase Ib detector.



PIXEL CALIBRATION




Figure 30.1: The maximum achievable voltage shift depending on VPDAC measured for a single pixel from [20].

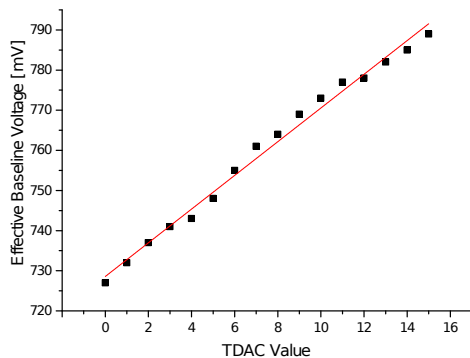


Figure 30.2: Achievable TDAC voltage shifts for a fixed VPDAC measured for a single pixel from [20].

Heiko: Describe Pixel tuning procedures

In a large scale silicon pixel chip there will be small performance variations for each pixel. These can be caused by local variations of process parameters in the production process or by bias voltage differences due to small voltage drops. To counteract these variations the MuPix sensor provides a tuning infrastructure which allows to adjust each pixel threshold individually.

30.1 MuPix Pixel Tuning

The individual pixel tuning is implemented in the pixels discriminator. By introducing an additional current the difference between the global threshold and baseline voltage can be adjusted. The tuning current can be selected with a 4 bit DAC (TDAC), while the maximum current is steered by a global bias voltage (VPDAC).

30.2 MuPix Tuning Procedure

The methods used to find the global tuning bias setting and the individual 4-bit tuning values are discussed in the following.

30.2.1 TUNING PARAMETERS

The goal of the sensor tuning is to equalize the overall pixel response and to maximise the sensors efficiency for ionising particles. To achieve a high




Figure 30.4: Threshold distribution for a 0.1 Hz per pixel noise rate measured after applying the tuning procedure for the MuPix6 chip from [139].




Figure 30.3: Threshold distribution for a 0.1 Hz per pixel noise rate measured for an untuned MuPix6 chip from [139].

efficiency the offset between threshold and baseline has to be minimised in order to be sensitive to small signals. This minimisation is limited by noise, as for small threshold-baseline offsets also noise signals can trigger the pixel. So the tun-

ing procedure is a trade-off between efficiency and noise.

In the current scheme we choose the noise rate as the tuning variable. It allows to tune on the highest acceptable noise rate per pixel, which will give at the same time a good efficiency.

A further possibility is the tuning to a minimal detectable signal. However, this is not implemented for the current prototype as it requests a well known and stable test pulse procedure. This feature is under investigation.

30.2.2 TUNING PROCEDURE

The currently used tuning procedure uses the per-pixel noise rate as tuning parameter, which depends of the threshold-baseline offset and varies for different pixels. The noise rate can be measured for all pixels in parallel by using the digital chip readout and counting the hits without an external stimulation in a fixed time interval.

In a first step, the global tuning bias VPDAC needs to be determined. To do this the global threshold is swept from a noiseless value far from the baseline towards the baseline, approaching it in equidistant steps. Due to the variable noise characteristics of the pixels some will cross the chosen noise rate earlier. These pixels are than muted by setting their TDAC to the maximum and enlarging VPDAC until these pixels are below the accepted noise rate. This sweep is continued until all pixels have been muted once and VPDAC is increased whenever a muted pixel crosses the noise limit.

The stopping threshold of this VPDAC tuning procedure marks the global threshold for which all pixels show noisy behaviour in the untuned state, but can be tuned to a noise rate below the aimed per pixel noise level with the determined VPDAC value. In the following this threshold is referred to as nominal threshold.

In a next step the individual 4-bit TDAC values need to be determined. This is done in the same manner by resetting the individual TDAC values to 0 and sweeping the threshold in equidistant steps to the nominal threshold. Whenever a pixel exceeds the chosen per-pixel noise level its TDAC is increased by one. By design the maximal needed TDAC value can not be larger than the maximum TDAC value.

There are several possibilities to optimise this tuning procedure. One main focus is definitely the time which is needed to perform the proce-




Figure 30.6: The resulting TDAC pixel map for the 0.1 Hz tune from [139].

ture. For this optimisation the time intervals for the noise measurements can be shortened which comes at the cost of a less precise noise determination. This however can be compensated by a safety margin for VPDAC to ensure that all pixels can be tuned below the requested noise rate.

Further this method by definition tunes all pixels below the aimed per-pixel noise rate which may cause an over-tuning, which reduces the efficiency. To counteract this an additional step can be added in the end to measure the chips noise distribution at the nominal threshold and reduce the tuning values for pixels, which have a noise rate below a lower noise rate constraint.

The effect of the basic tuning can be observed in figures 30.3,30.4,30.5 and 30.6.




Figure 30.5: The resulting TDAC distribution for the 0.1 Hz tune from [139].



FIBRE CALIBRATION

This chapter describes the different handles for calibration and performance monitoring of the Fibre Detector. Calibration refers mainly to the readout system, whereas the performance of the fibre and SiPM arrays have to be carefully monitored.

The different readout channels have to be aligned in time, accounting for different signal path lengths and delays inside the readout ASIC. The MuSTiC, the aspired readout solution in the baseline design, determines the time of an event with a leading edge discriminator, see 13.2. The levels of these thresholds have to be tuned to ensure on one hand side a uniform response and on the other hand to account for variations in the gain of the SiPM arrays. Such variations can be caused by temperature fluctuations or SiPM degradation due to radiation damage. This issues may also be addressed with HV compensation which is highly related to threshold levels. More technical details of such threshold scans are provided in section 31.1. Further variations in the signal height caused by fibre degradation or fibre-SiPM coupling inconsistencies have to be carefully monitored with complementary tools.

Since the current design foresees no energy measurement, there is no handle for time-walk calibration.

31.1 MuSTiC Thresholds

The Fibre Detector MuSTiCs are operated in *time mode*, the *energy threshold* is disabled. The discrete values of SiPM response signals¹ can be determined by measurements of event rate at dif-

ferent threshold values. Figure 13.1 shows such a step function where only SiPM dark counts are used. Too low threshold values lead to channel saturation due to the ASIC internal noise². This is followed by several steps, whose plateaus show the dark rate for $\leq n$ firing pixels. Multiple pixels fire because of SiPM cross-talk.

This procedure was used with the available STiC3.1 for the measurements described in section 13.2.1. Due to the limited data rate of this STiC version the channels have to be calibrated in series. Data acquisition of 1 s per threshold provide more than enough statistics and could be reduced. At the moment the calibration of 32 channels requires slightly more than 1 h. For usage in the experiment this has to be improved by calibrating all 32 channels of one ASIC at the same time and possibly by a faster loading of DAC settings.

31.2 Inter Channel Calibration

Once all the channels use the same thresholds, a inter channel calibration regarding time is needed for a proper background suppression. Due to the array-like readout scheme, the scintillating photons of a crossing particle spread out over several columns in the SiPM arrays. This clusters can be exploited to align the read out channels inside one ribbon. If a Fibre Detector hit can be matched to a reconstructed track, the hit position is known.

¹charge = number of fired pixels x gain

²note that STiC is designed for both input polarities, hence zero threshold is around the middle of the range of [0,63].



Taking this position and the time of photon propagation inside the fibres into account, an alignment of the channels up- and downstream is achievable.

Time calibration between ribbons relies on reconstructed tracks which cross the Fibre Detector multiple times (8-hit segments). The possibility to trigger the MuSTiC TDC manually with an external pulse, as described in 13.2.2, offers an additional handle.

Furthermore, a leaky optical fibre could be run along the SiPM arrays. It allows to inject much more light than usually recorded from the ribbons, thus even the *energy threshold* could be used. These much larger signals, resulting in much better time resolution, could potentially lead to faster and better inter-channel and ribbon calibration.

31.3 Monitoring Tools

Different tools for performance monitoring are under discussion. To monitor the MuSTiC performance, in particular the PLL status, one TDC is continuously triggered by a configurable external signal in phase to the reference clock, see section 13.2.2.

Furthermore, a leaky optical fibre along the SiPM arrays allows to inject different sized signals directly into the SiPM arrays as well as into the fibres. These signals provide a handle for debugging and allow to monitor possible degradations in the fibres and SiPMs due to radiation damage or inconsistencies in the fibre-SiPM couplings.

Discussions to add a second time-like threshold to the MuSTiC ASIC are ongoing. Such a second threshold would not provide a timestmap, thus no additional TDC is needed; crossing the threshold would be indicated by a single bit.



TILE CALIBRATION

Patrick: Describe Tile Calibration

Patrick: Can we use cosmics? - Nik

The calibration and monitoring of the Tile detector response is essential for achieving a high time resolution. The main task of the calibration is to synchronise all channels. In particular, the delay arising from the length of the signal lines has to be corrected for. Furthermore, the energy dependence of the timestamps has to be characterised, in order to correct for time-walk effects. Besides the detector calibration, the monitoring of the detector response is required to assure a stable operation. This is necessary since the response is expected to vary over time, due to fluctuations in the ambient temperature and degradation of the SiPM response caused by radiation damage.

For the calibration and monitoring, coincident signals are required. There are three processes which can be used for calibration and monitoring: internal conversion decays, e^+e^- pairs from Bhabha scattering and hit clusters from a single track. Coincident signals from internal conversion

decays and e^+e^- pairs offer the most precise and direct way to calibrate the detector. However, since the event rate is relatively low, this method is not suitable for monitoring purposes and a fast detector calibration.

For the inter-calibration of different tiles, it can be exploited that a single electron track produces a small cluster of hits in the tiles with an average cluster size of about two. The different signals within a cluster are basically coincident, up to the time of flight between the individual tiles of a cluster of about 10 ps. This allows for a relative timing calibration of all cells within a cluster. Since all particle tracks can be used, this provides a fast way to calibrate and monitor the response. As this method is limited to the inter-calibration within one re-curl station, it is complementary to the global calibration with internal conversion decays.

More detailed studies of the calibration and monitoring scheme described above are presented in [140].



DETECTOR ALIGNMENT

In order for the reconstruction algorithms to work optimally, the position and orientation of all active detector elements as well as the stopping target have to be known to good precision. The position of the pixels inside a sensor is given by the tolerances of the manufacturing process, which are much better than the minimal feature size of 180 nm; compared to all other sources of misalignment, this is completely negligible. The task of detector alignment is thus to determine the position and orientation of all active detector parts (HV-MAPS chips, fibres, tiles).

The first step in ensuring a well-aligned detector is the careful assembly of modules and layers using precision tools, followed by a detailed survey. After detector installation, movements of larger detector parts (e.g. the recurl stations with regards to the central detector) can be followed by a system of alignment markers observed by digital cameras inside the magnet. The ultimate alignment precision is however only reached with track-based alignment methods.

33.1 Effects of Misalignment

We have studied the effects of a misaligned pixel detector on reconstruction efficiency and tracking resolution using the full detector simulation. For technical reasons, the sensors are all in their nominal position for the Geant4 simulation, the reconstruction is then however performed assuming different sensor positions.

Random shifts and rotations of the sensors (with shifts and rotations drawn from a normal distribution with given sigma) lead to a worsening

of the reconstruction efficiency; there is however an efficiency plateau if the sensor positions are known to about the pixel size and the rotations to better than 0.3° , see figure 33.1.

Global shifts and rotations of detector parts also affect the efficiency, see figure 33.2 for some of the studied modes.

The momentum resolution is also negatively affected by misalignment, see figures 33.3 and 33.4 as well as figure 33.5. The tails and the core of the resolution distributions are somewhat differently affected by both the misalignment and the selection effects, explaining the differences between figures 33.3 and 33.4. The constraints on the alignment accuracy for achieving optimal momentum resolution are much tighter than for the efficiency - ideally positions and rotations should be known to better than $10\ \mu\text{m}$ and rotations to better than 0.1° . This in turn implies that a preliminary alignment of the detector will have to fulfil the efficiency requirements; track based methods can then be used to further improve until we reach the resolution requirements. The absolute size of the requirements seems well within reach - we regularly align the MUPIX telescope to better than $150\ \mu\text{m}$ by shifting the physical planes, track based alignment then improves this to better than $10\ \mu\text{m}$.

33.2 Position Monitoring System

The positions of the detector stations relative to each other will be monitored by a series of cameras mounted to the cage for inserting the detector




Figure 33.1: Reconstruction efficiency for 4-hit segments (left), 6-hit segments (middle) and 8-hit segments (right) for random sensor shifts and rotations normalized to the efficiency of a perfectly aligned detector.




Figure 33.2: Reconstruction efficiency for 4-hit segments for various global misalignments. Innermost layer: shift along x axis (top left), shift along z axis (top centre), rotation around x axis (top right), rotation around z axis (middle left). Innermost two layers, shift along x axis (middle centre), shift along z axis (middle right). Innermost two layers, rotation around x axis (bottom left), rotation around z axis (bottom center). Torsion of the complete detector around the z axis (bottom right).




Figure 33.3: Root mean square momentum resolution of 4-hit segments (left), 6-hit segments (centre) and 8-hit segments (right) for random sensor shifts and rotations in MeV/c. The slight improvement of the resolution for very large misalignments is due to a selection effect: Only tracks passing through sensors close to their nominal positions can be reconstructed.




Figure 33.4: Sigma of a Gaussian fit to the core of the momentum resolution distribution of 4-hit segments (left), 6-hit segments (centre) and 8-hit segments (right) for random sensor shifts and rotations in MeV/c. The slight improvement of the resolution for very large misalignments is due to a selection effect: Only tracks passing through sensors close to their nominal positions can be reconstructed.




Figure 33.7: Position in r and z of the points of closest approach to the beam line for $1 \cdot 10^8$ stopped muons. The target is clearly visible, as are parts of the support wires. Negative radius implies that the beam line is inside of the track circle, positive is outside.




Figure 33.6: Schematic view of a possible alignment system using three cameras. The detector support cage is shown in black, the blue tube represents the detector stations with the end-rings shown in red. The three cameras and their fields of view are shown in red, green and blue.

into the magnet. They will be complemented by light sources (the detector is usually operated in the dark) and alignment marks on the end-rings of the detector stations. Cameras with a 85° field of view are sufficient to view all end rings in the phase Ib detectors when mounted to the cage at the centre of the magnet. A system of three cameras (e.g. top and $\pm 60^\circ$ from the bottom) also allows to track relative movements of the cameras, as they can see each other. Sub millimetre resolution will require fairly high resolution cameras




Figure 33.5: Root mean square momentum resolution of 4-hit segments for various global misalignments. Innermost layer: shift along x axis (top left), shift along z axis (top centre), rotation around x axis (top right), rotation around z axis (middle left). Innermost two layers, shift along x axis (middle centre), shift along z axis (middle right). Innermost two layers, rotation around x axis (bottom left), rotation around z axis (bottom center). Torsion of the complete detector around the z axis (bottom right).



(2K or 4K) or the use of separate cameras with zoom lenses focused on the station-station transitions. A possible three-camera system is shown in figure 33.6.

33.3 Track-Based Alignment

The fine alignment of the silicon sensors (as well as the fibres and tiles) will be performed using track based methods initially developed in the H1 experiment [141] and since successfully applied to a variety of large and very large tracking systems, e.g. CMS at the LHC [126, 142–148].

The alignment of the complete detector is a large minimization problem, where, for a very large sample of tracks, the residuals from the measured hits to the fitted tracks have to be minimized under variation of both all track and all alignment parameters. If a rough detector alignment is known, corrections will be small and the minimization problem can be linearised.

To this end, tracks reconstructed with the standard reconstruction algorithms described in chapter 28 are re-fitted using the general broken lines (GBL) algorithm [126, 127]. The GBL software can calculate and output the complete covariance matrix between track- and alignment parameters. As the track parameters are not correlated with each other and only relate to the alignment of the small subset of sensors which are hit by the track, the resulting matrix is sparse. There are efficient algorithms for the inversion of such gigantic sparse matrices, one of which is implemented in the *Millepede II* program [149], which we intend to use.

Whilst the sensor alignment is locally well constrained via the overlap of the sensors in azimuthal

direction and the closeness of the double layers, overall deformations such as shifts of the top part with regards to the bottom part, an overall torsion or the position of the recurv stations are only weakly constrained by using tracks from muon decays. These so-called *weak modes* need additional input from tracks which are correlating distant parts of the detector. These tracks are provided by cosmic ray muons. As the cosmic rate is tiny compared to the beam muon rate, it is imperative to have a special trigger stream (potentially using a scintillator coincidence above and below the experiment) to collect enough cosmics for alignment.

We are currently studying the alignment procedure, the amount of statistics needed and potentially problematic weak modes using the simulation; first results are expected in summer 2016.

33.4 Target Alignment

The position of the target needs to be known with very high accuracy as an input for cuts on the vertex distance to the target. As the target is passive, the residual-based method described in the previous section is not applicable. We however do know, that the overwhelming majority of decay positrons originates on the target surface and will thus have a point of closest approach (POCA) to the beam axis inside the target.

This can be used to determine the target position by plotting the distribution of the POCAs in the transverse plane in slices of z for many tracks, which will give an accurate determination of the position of the outer target edge, see figures 33.7 to 33.9. The target thickness has to be determined during manufacture or using photon conversions.




Figure 33.8: Position in x and y of the points of closest approach to the beam line for a 1 mm slice in z at -6 mm (left) and -28 mm (right) for $1 \cdot 10^8$ stopped muons. The target is clearly visible at its nominal position.




Figure 33.9: Position in x and y of the points of closest approach to the beam line for a 1 mm slice in z at -6 mm (left) and -28 mm (right) for 500 000 stopped muons. The simulation was performed with a target shifted by 2 mm in x direction.

PART X



SIMULATED PERFORMANCE



SENSITIVITY STUDY

We study the performance of the detector described in the preceding parts by running the Geant4 simulation and the reconstruction program for both the phase Ia and phase Ib setups. Even under optimistic assumptions, only a handful of signal decays are expected in the data. Nonetheless, we use relatively large signal samples to study the detector performance and deduce a very preliminary and rough event selection.

For the various expected backgrounds, in principle the simulation of several times the expected number of decays in data is required. This is impractical both in terms of processing time and available storage space. We thus try to identify important sources of background from either general considerations (internal conversion) or from simulating a few seconds of run time (accidentals). From these starting points we then generate special samples and — in case of the accidental background — make use of the fact that the background suppression via vertexing, via timing and via event kinematics are at least to a reasonable approximation independent.

34.1 Signal Performance

We study the nominal performance of the detector setup for signal decays using 1 million decays each generated for the phase Ia and phase Ib detector configurations. The decay electrons have a phase space distribution. Efficiencies are determined relative to all muons decaying inside a cylinder with the outer dimensions of the stopping target.

In a first step, all three tracks from the signal decay have to be reconstructed to at least short

(4-hit) tracks; for the efficiency and resolution of the track reconstruction, see chapter 28.

34.1.1 VERTEX FIT

The three tracks from signal decays should intercept in a common point on the surface of the target. We look at all combinations of a track with negative charge and two tracks of positive charge. In order not to fit recurling tracks with themselves, the track tangent vector at the point of closest approach is determined and if cosine of the opening angle between two tracks is more than




Figure 34.1: χ^2 of the vertex fit for signal events in the phase Ia detector configuration.

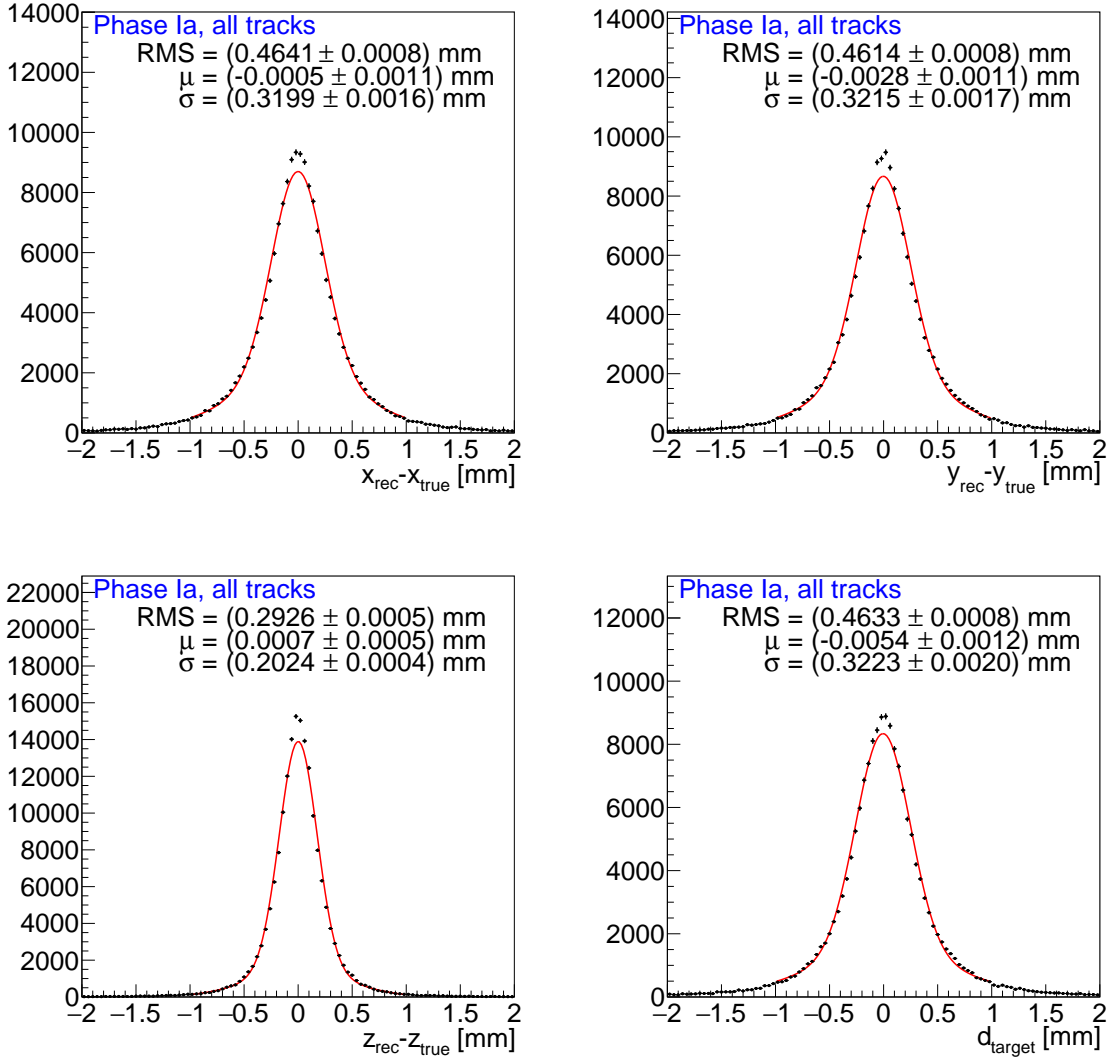


Figure 34.2: Vertex resolution for signal decays in the Phase Ia detector configuration. All tracks with four hits or more are used. The fits are the sum of two Gaussian distributions and the quoted σ is the area-weighted mean. Top left in x , top right in y , bottom left in z and bottom right in the distance to the target; negative target distances denote a reconstructed vertex position inside the target.

0.99 and the momentum difference is less than 1 MeV/c, the combination is not further considered. These selection cuts are subject to an ongoing study and will be further optimized, especially in view of their implementation in the online selection.

The multiple scattering vertex fit as described in section 28.5 is performed; if successful, it gives to handles for background suppression, namely the χ^2 of the fit (see figure 34.1) and the distance

to the target surface; the precision of the vertex reconstruction is illustrated in figure 34.2 for the phase Ia setup and no requirement on track length and in figure 34.3 for the phase Ib setup and the requirement of all three tracks being reconstructed including recurlers. The improvement in the vertex resolution from the improved momentum reconstruction in phase Ib is marginal; this is to be expected, as the uncertainty is dominated by the scattering in the first plane and the position

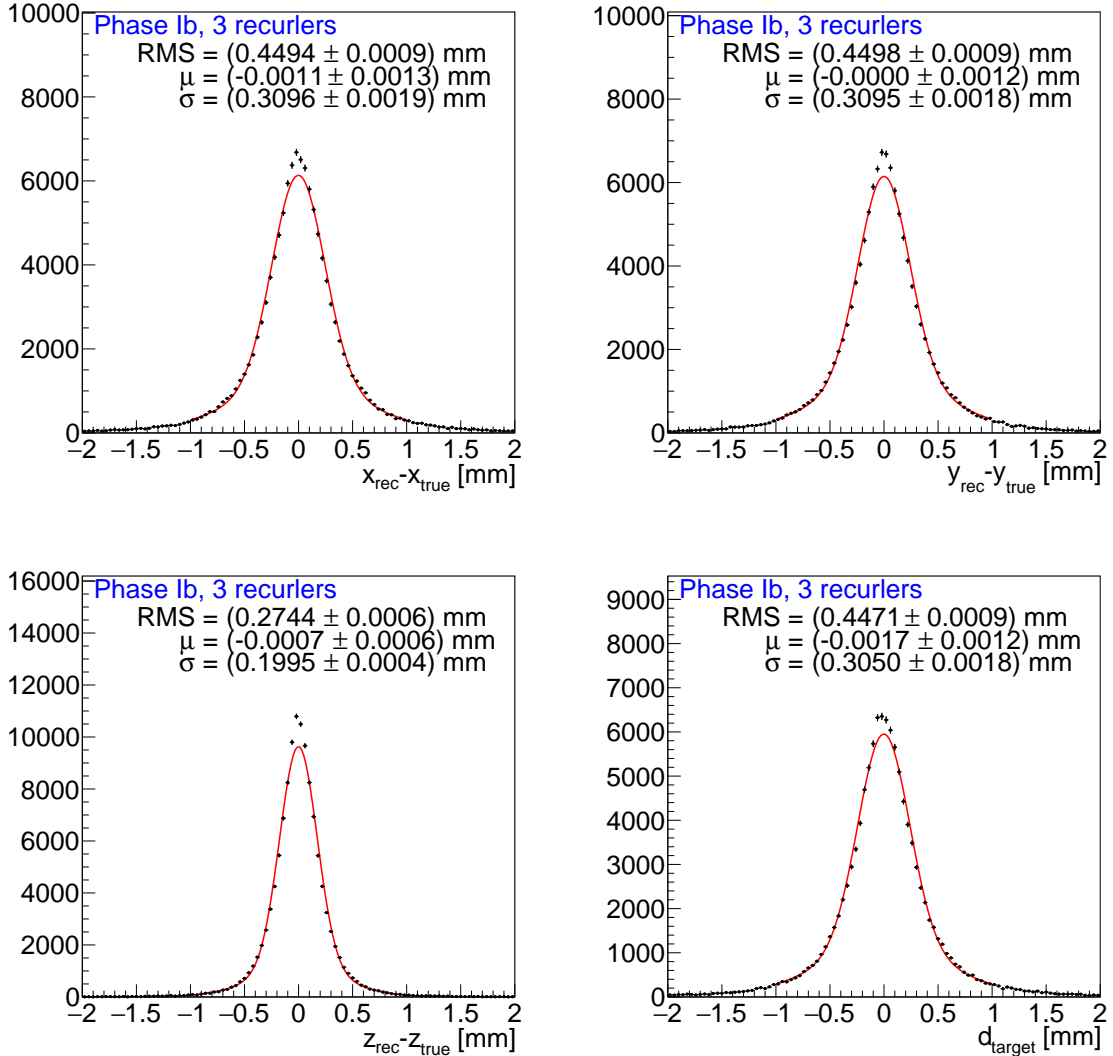


Figure 34.3: Vertex resolution for signal decays in the Phase Ib detector configuration. Three tracks with recurlers are selected. The fits are the sum of two Gaussian distributions and the quoted σ is the area-weighted mean. Top left in x , top right in y , bottom left in z and bottom right in the distance to the target; negative target distances denote a reconstructed vertex position inside the target.

resolution in the first plane and knowledge of the curvature plays an inferior role.

34.1.2 MASS AND MOMENTUM RECONSTRUCTION

For all candidates with a vertex fit $\chi^2 < 15$ and a distance between reconstructed vertex and target of less than 1 mm, the tracks are extrapolated to the vertex and four-vectors are constructed with

an electron mass assumption. From the three four-vectors, the mass of the decaying particle (should correspond to the muon mass) and the momentum of the center-of-mass system (CMS) in the detector frame (should be zero for decays at rest) are determined.

The resolution for the muon momentum is depicted in figures 34.4 and 34.5 for phase Ia and phase Ib respectively. The magnitude of the reconstructed momentum is shown in figure 34.6.




Figure 34.4: Reconstructed decay muon momentum in x , y and z direction (which corresponds to the resolution for p_x , p_y and p_z for muons decaying at rest). The phase Ia detector setup was used and all tracks enter the analysis.




Figure 34.5: Reconstructed decay muon momentum in x , y and z direction (which corresponds to the resolution for p_x , p_y and p_z for muons decaying at rest). The phase Ib detector setup was used and only long tracks enter the analysis.

We define the decay plane from the three momenta \vec{p}_i ,

$$\vec{a} = \frac{(\vec{p}_1 - \vec{p}_2) \times (\vec{p}_3 - \vec{p}_2)}{|(\vec{p}_1 - \vec{p}_2) \times (\vec{p}_3 - \vec{p}_2)|} \quad (34.1)$$

is a vector \vec{a} perpendicular to the decay plane (if the tracks are from a muon decaying to the signal channel at rest).

The SINDRUM experiment based their selection on the projection of the CMS momentum onto this vector, called acoplanar momentum

$$\vec{p}_{\text{acoplanar}} = \vec{p}_{\text{CMS}} \cdot \vec{a}. \quad (34.2)$$

and the coplanar momentum

$$\vec{p}_{\text{coplanar}} = \vec{p}_{\text{CMS}} \times \vec{a}. \quad (34.3)$$

To first order, the resolution for the acoplanar momentum is only dependent on the measurement

(and thus resolution) of the track angle measurement, whereas the coplanar momentum, is dominated by the absolute momentum resolution.

The corresponding distributions are shown in figure 34.7; for the Mu3e setup (as it was for SINDRUM) the resolution in the acoplanar momentum is clearly superior. It however turns out that the acoplanar momentum is not well suited to suppress the dominating accidental background from Bhabha scattering where both the electron and the positron end up in the detector acceptance; in this case the coplanar momentum or the total momentum of the CMS system provide a better coplanar momentum or the total momentum of the CMS system provide a better handle. SINDRUM performed the selection on a resolution weighted combination of the coplanar and acoplanar momenta. No detailed study of the momentum selection has been performed for Mu3e so far,




Figure 34.6: Center of mass system momentum reconstructed for signal events in phase Ia with all tracks considered (left) and phase Ib with three recurlers required (right).

so for the distributions shown in this report, we used the requirement of $p_{CMS} < 4 \text{ MeV}/c$.

Finally, we show the resolution for the reconstructed mass in figures 34.8 and 34.9. As the distributions show, the core of the mass resolution fulfils the criteria set out in chapter ?? and especially for recurling tracks, sizeable, Landau-like tails only appear on the low mass side.

34.1.3 SIGNAL EFFICIENCY

For every reconstruction step, there is a possibility of signal loss; the largest loss is due to the geometrical acceptance of the detector. For phase-space signal decays in the target, 36.5 % have all three electrons traverse the four layers of the central detector in the active region. If recurling tracks are required, the acceptance is further reduced. There are unavoidable inefficiencies in the reconstruction and vertex fit, especially due to the χ^2 cuts, which mostly get rid of tracks with large angle scattering, precluding a reliable and precise reconstruction. The overall efficiency after applying all mentioned cuts is shown in figure 34.10 in dependence of the required number of recurling tracks.

The phase Ia setup offers an overall efficiency of about 24 % if all tracks with four hits are used. In the phase Ib setup, the addition of the recurl stations allows for a more accurate track recon-

struction with more signal events passing the total momentum requirement and thus giving a higher overall efficiency of 25 %. If three recurling tracks are required, the overall efficiency is about 16 %.

The preceding considerations do not yet include inefficiencies due to the timing detectors; more studies are needed to devise appropriate selection criteria.

34.2 Backgrounds

34.2.1 INTERNAL CONVERSION BACKGROUND

We simulate the internal conversion background as described in section Simulation:MuonDecays using the matrix element provided by Signer et al. [133]. The total branching fraction for this decay is $3.4 \cdot 10^{-5}$ [5], so a complete simulation is challenging. We are however mostly interested in the region of phase space where the neutrinos carry little momentum; the branching fraction for this high visible mass region (we used a lower cutoff of $90 \text{ MeV}/c^2$ for the studies presented here), is strongly suppressed and we can generously oversample in the simulation. In addition, we use weighted events in order to better populate the high mass tail. Migrations from lower masses than $90 \text{ MeV}/c^2$ into the signal region are very strongly suppressed if all tracks are used and completely

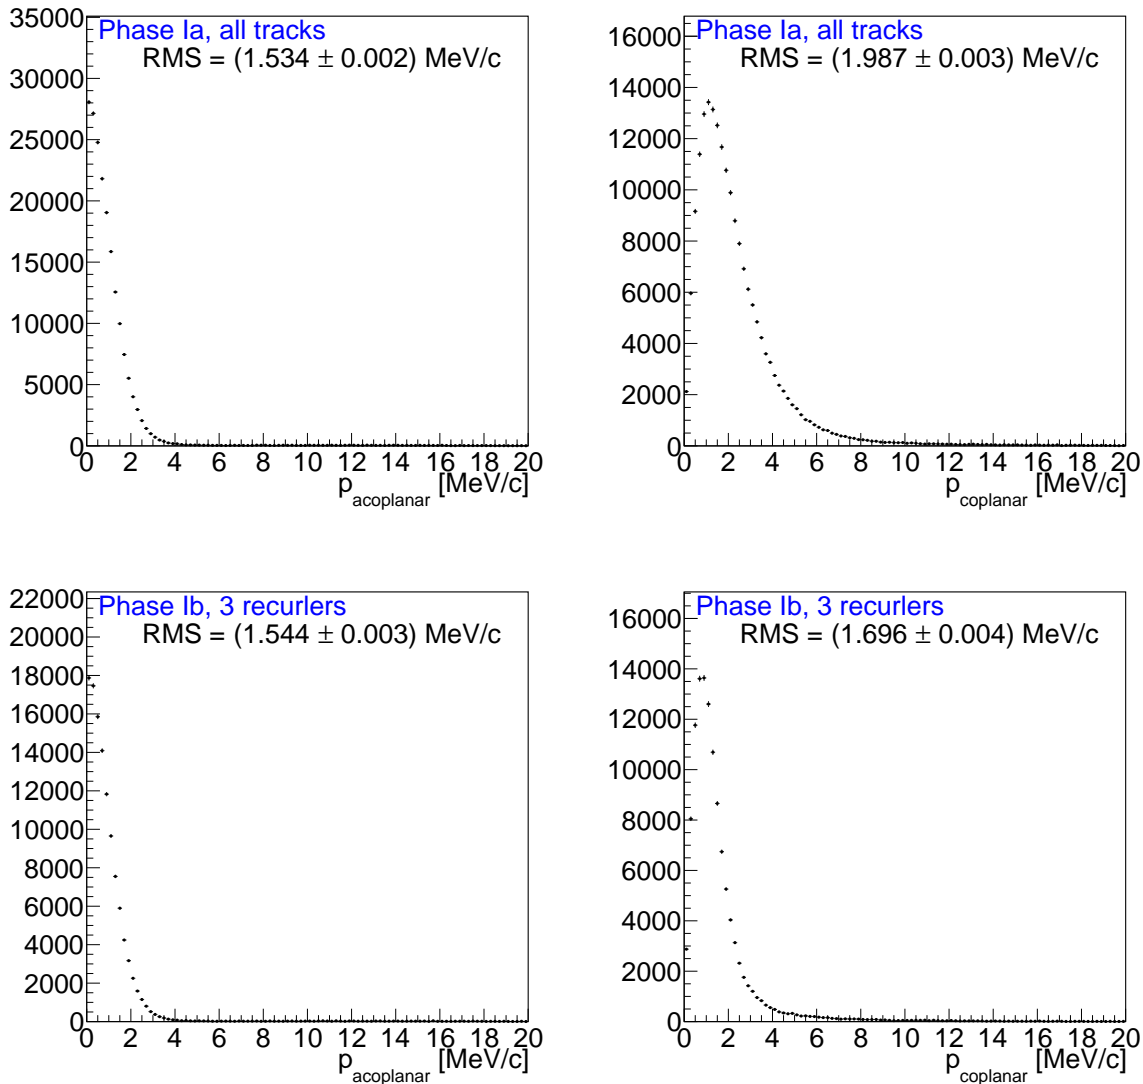


Figure 34.7: Acoplanar momentum (left) and coplanar momentum (right) reconstructed for signal events in phase Ia with all tracks considered (top) and phase Ib with three recurlers required (bottom).

absent if three recurling tracks are required, see figure 34.11.

34.2.2 ACCIDENTAL BACKGROUND

Accidental background arises from the combination of two Michel positrons with an electron. It is thus important to understand and limit electron production in the target region, especially for processes such as Bhabha-scattering, where the electron and positron tracks do intersect in space

and time and only the separation to the second positron remains as a suppression criterion.

Electron Production in the Target

The default target is part of the Geant4 detector simulation as described in chapter 27. This simulation was used to study stopping distributions and electron background produced in the target. Figure 34.12 shows the simulated positions of muon decays in the target region. In can be seen that the target is evenly illuminated by

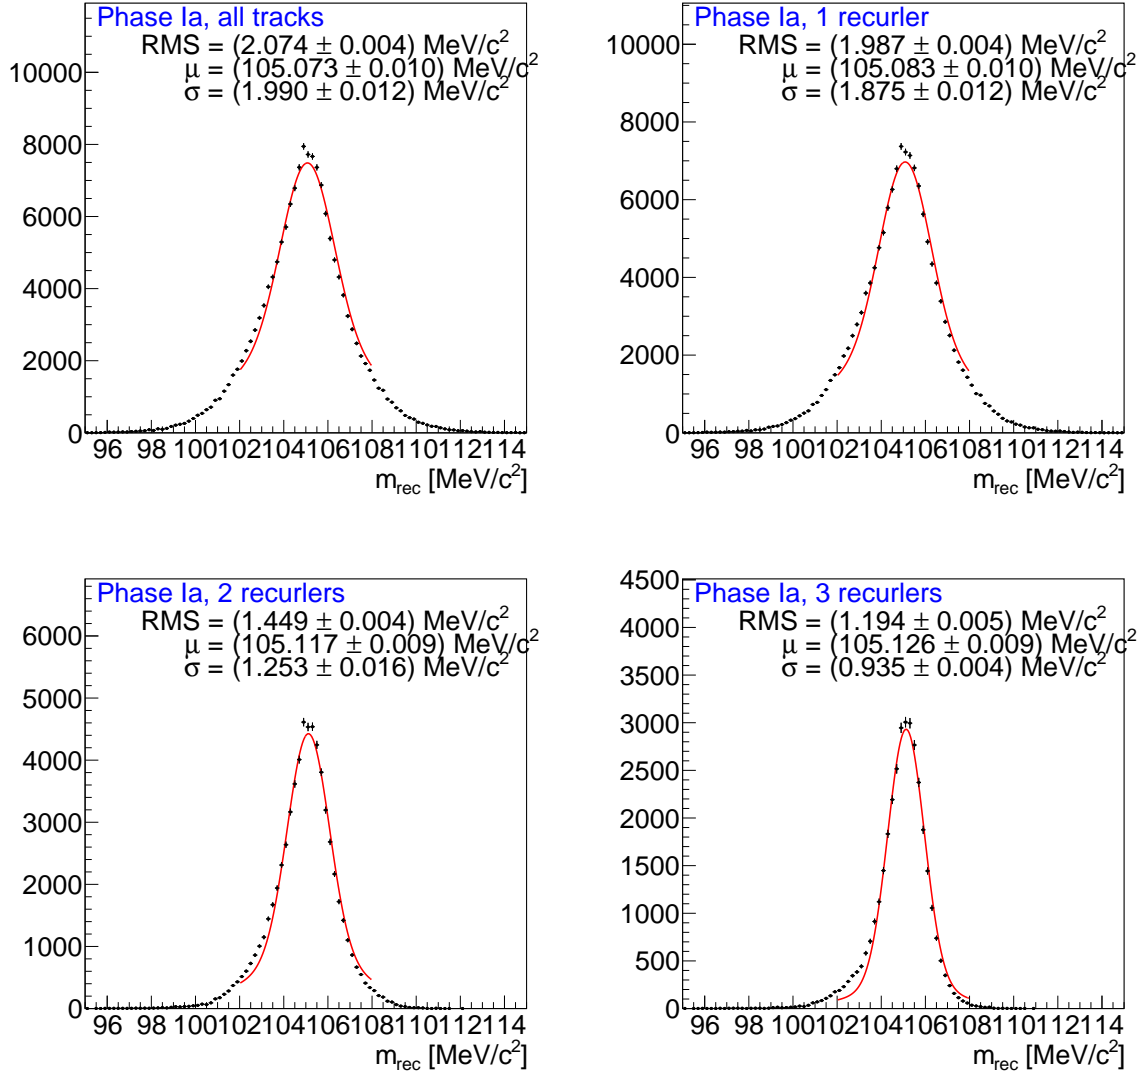


Figure 34.8: Reconstructed muon mass for the phase Ia setup for all tracks (top left), at least one recurler (top right), at least two recurlers (bottom left) and three recurlers (bottom right). The fits are the sum of two Gaussian distributions and the quoted σ is the area-weighted mean.

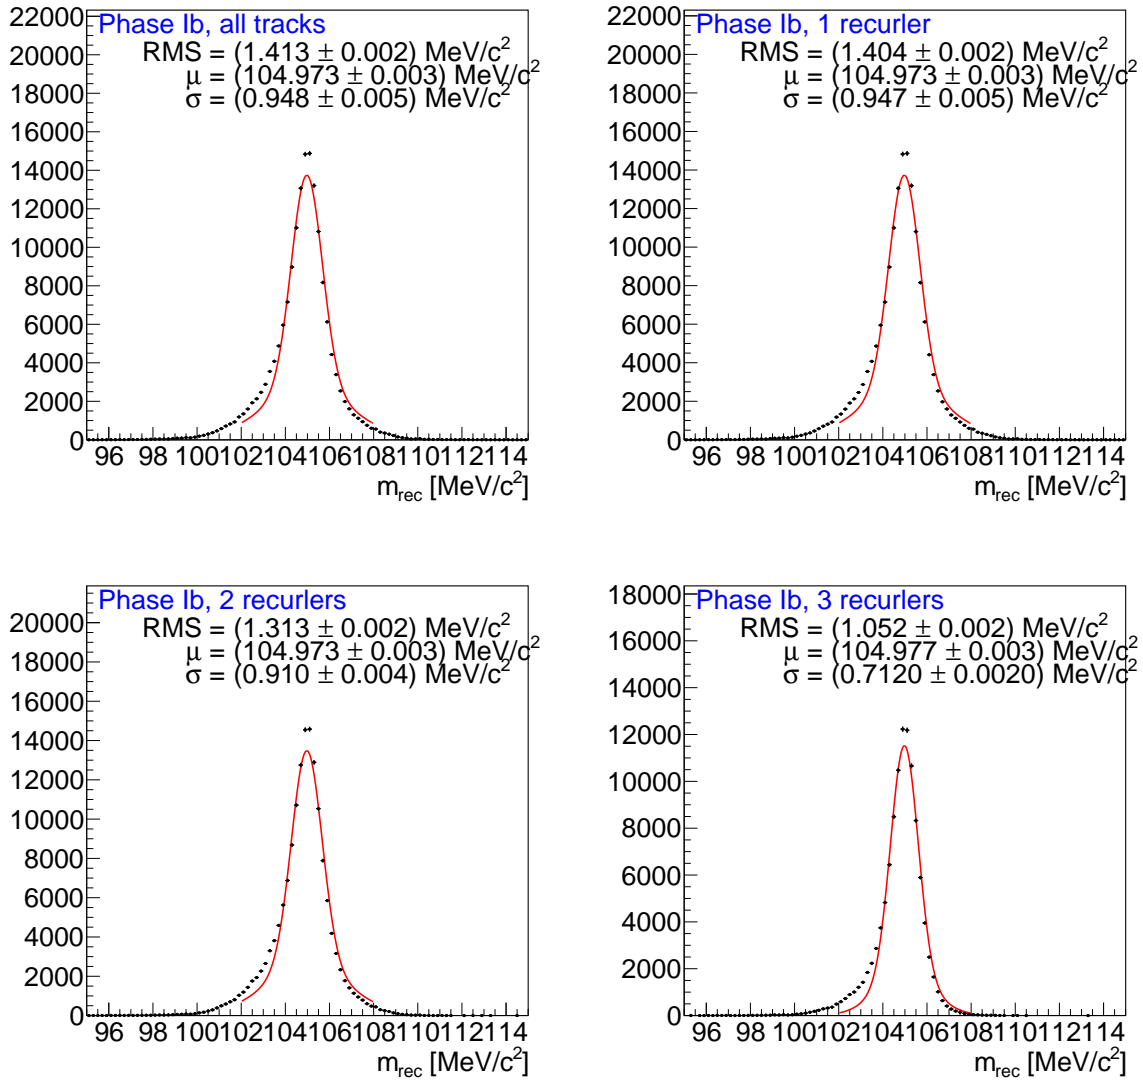


Figure 34.9: Reconstructed muon mass for the phase Ib setup for all tracks (top left), at least one recurler (top right), at least two recurlers (bottom left) and three recurlers (bottom right). The fits are the sum of two Gaussian distributions and the quoted σ is the area-weighted mean.




Figure 34.10: Total efficiency for reconstructing phase-space signal events in the phase Ia (left) and phase Ib (right) detector as a function of the required number of recurling tracks. This includes the geometrical detector acceptance, track and vertex reconstruction and selection inefficiencies.




Figure 34.11: Resolution of the mass reconstruction for internal conversion events with a visible mass above 90 MeV/c² for the phase Ia setup and all tracks (left) and the phase Ib setup and three recurling tracks (right).




Figure 34.12: Longitudinal view (left) and transverse view (right) of the distribution of Michel decays in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.




Figure 34.13: Longitudinal view (left) and transverse view (right) of the locii of Bhabha scattering producing an electron in the detector acceptance in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.

the beam and that muon stops in the first detector layer cannot be completely avoided.

The material of the target is a place where electrons from Bhabha and Compton scattering as well as from photon conversion can be produced and contribute to accidental background as they cannot be suppressed by vertex position requirements. Figures 34.13 to 34.15 show where in the inner detector region these electrons are produced. It can be clearly seen, that Bhabha scattering is the dominating production process for electrons. This needs special attention, as very often both the electron and the positron partaking in the scattering process end up in the detector accep-

tance; the corresponding vertices are shown in figure 34.16. As shown in figure 34.17, almost all the corresponding primary positrons come from muon decays in the target and can thus not be further reduced or shielded.

The total number of electrons produced per Michel decay is shown in table 34.1. As can be easily seen, Bhabha scattering is the most worrying background, especially because typically both the electron and the positron end up in the detector acceptance. The momentum spectrum of the electrons falling very fast, see figure 34.18; this means that many of the electrons end up at the very low edge of the detector and reconstruc-

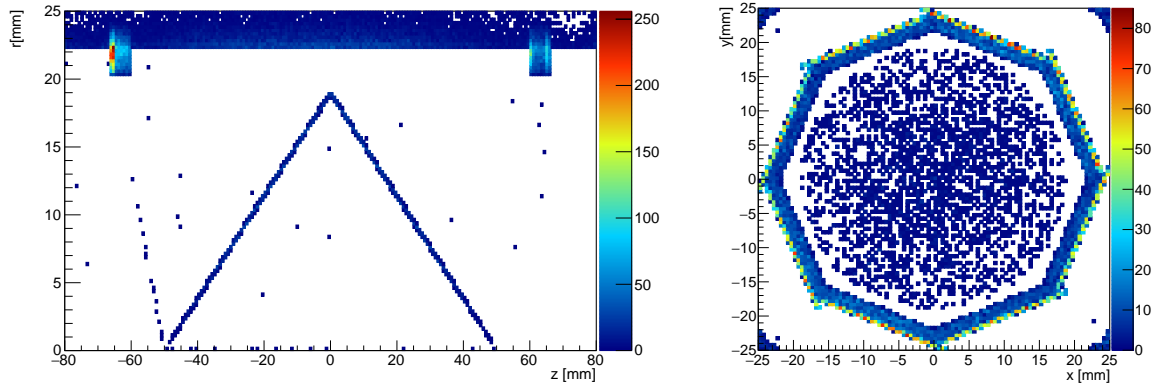


Figure 34.14: Longitudinal view (left) and transverse view (right) of the locii of photon conversion producing an electron in the detector acceptance in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.

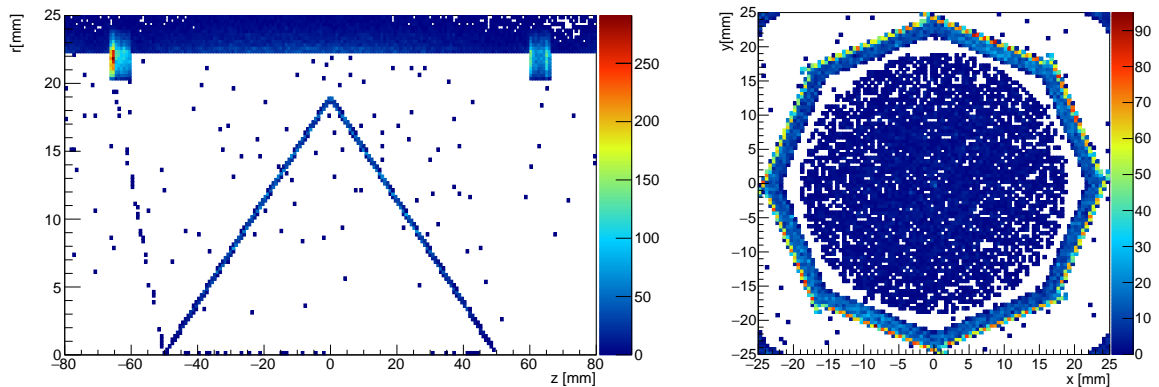


Figure 34.15: Longitudinal view (left) and transverse view (right) of the locii of Compton scattering producing an electron in the detector acceptance in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.

	Electron + 2 Michel	Bhabha + Michel	number of electrons actually reconstructed, see table 34.1.
Phase Ia			
Phase Ib, all tracks	$2.2 \cdot 10^{-4}$		$2.0 \cdot 10^{-2}$
Phase Ib, long tracks	$1.3 \cdot 10^{-4}$		$1.6 \cdot 10^{-2}$

Table 34.2: Vertex suppression factors: How often does an electron (respectively a Bhabha electron-positron pair) get fitted to two (one) Michel positrons in a 50 ns frame.

Vertex suppression

The suppression of accidental background by the vertex fit and associated cuts is highly dependent on the track density. It also depends on the kinematics of the three involved particles. We determine the suppression by running the full simulation for the equivalent of a few seconds of runtime both at phase Ia and phase Ib intensities. We then select all 50 ns reconstruction frames containing at least one electrons. These frames are

tion acceptance, explaining the significantly lower

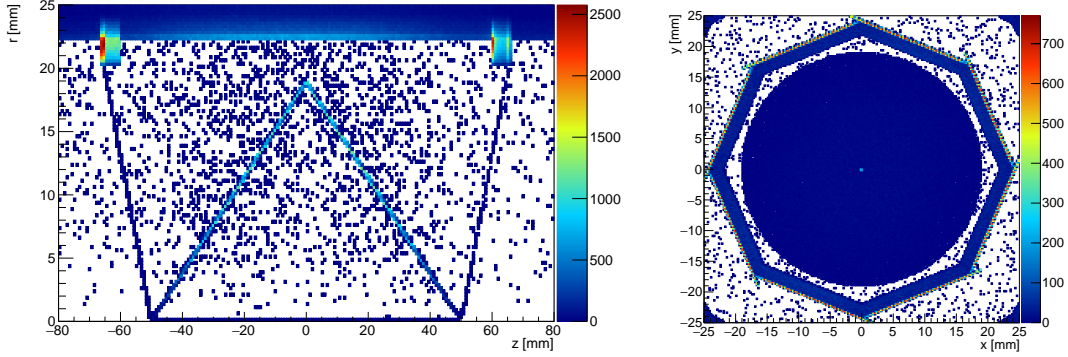


Figure 34.16: Longitudinal view (left) and transverse view (right) of the locii of Bhabha scattering producing an electron and a positron both in the detector acceptance in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.




Figure 34.17: Longitudinal view (left) and transverse view (right) of muon decay vertices leading to a positron then undergoing Bhabha scattering in the target resulting in an electron and a positron both in the detector acceptance in the target region for one second of running at $1 \cdot 10^8$ muon stops per second.

Electron source	Produced in inner detector	Produced in target region	Reconstructed inner detector short tracks	Reconstructed target region short tracks	Reconstructed inner detector long tracks	Reconstructed target region long tracks
Bhabha scattering	$5.0 \cdot 10^{-4}$	$1.2 \cdot 10^{-4}$	$3.0 \cdot 10^{-4}$	$6.7 \cdot 10^{-5}$	$2.3 \cdot 10^{-4}$	$5.0 \cdot 10^{-5}$
both visible	$3.8 \cdot 10^{-4}$	$7.8 \cdot 10^{-5}$	$1.7 \cdot 10^{-4}$	$3.4 \cdot 10^{-5}$	$1.1 \cdot 10^{-5}$	$2.1 \cdot 10^{-5}$
Photon conversion	$1.6 \cdot 10^{-5}$	$1.9 \cdot 10^{-6}$	$9.7 \cdot 10^{-6}$	$1.1 \cdot 10^{-6}$	$7.5 \cdot 10^{-6}$	$8.1 \cdot 10^{-7}$
both visible	$4.5 \cdot 10^{-6}$	$5.2 \cdot 10^{-7}$	$1.7 \cdot 10^{-6}$	$1.8 \cdot 10^{-7}$	$1.1 \cdot 10^{-6}$	$1.1 \cdot 10^{-7}$
Compton scattering	$2.2 \cdot 10^{-5}$	$4.3 \cdot 10^{-6}$	$1.2 \cdot 10^{-5}$	$2.3 \cdot 10^{-6}$	$9.6 \cdot 10^{-6}$	$1.7 \cdot 10^{-6}$
Internal conversion	$3.5 \cdot 10^{-5}$	$2.8 \cdot 10^{-5}$	$1.7 \cdot 10^{-5}$	$1.4 \cdot 10^{-5}$	$1.3 \cdot 10^{-5}$	$1.0 \cdot 10^{-5}$
two visible	$7.3 \cdot 10^{-7}$	$6.2 \cdot 10^{-7}$	$2.1 \cdot 10^{-7}$	$1.8 \cdot 10^{-7}$	$1.3 \cdot 10^{-7}$	$1.1 \cdot 10^{-7}$
Total	$5.8 \cdot 10^{-4}$	$1.5 \cdot 10^{-4}$	$4.5 \cdot 10^{-4}$	$1.1 \cdot 10^{-4}$		

Table 34.1: Electrons with transverse momentum larger than 10 MeV created per Michel decay in the target for the phase Ib detector setup. The inner detector region is a cylinder including the vacuum window and the first pixel layer, the target region is a cylinder just containing the target.




Figure 34.18: Momentum spectrum (left) and transverse momentum spectrum (right) of electrons produced in the target region.

reconstructed and the vertex fit is performed for all electron- positron-positron combinations. The fit is required to have a χ^2 of less than 15 (for three degrees of freedom) and the reconstructed vertex is required to be less than 1 mm from the nominal stopping target surface.

We determine suppression factors for the accidental combination of an electron with two independent Michel positrons and for a Bhabha electron-positron pair combined with an independent Michel positron, see table ?? . Unfortunately, the suppression for electron-positron pairs from photon conversion or internal conversion decays cannot be determined in this manner, as the rate is too low to produce sizeable samples; we currently assume that the vertex suppression is comparable to the Bhabha case and reuse that number.

Timing suppression

Time information of hits in the fibre and tile detectors provide an additional handle on the suppression of accidental backgrounds.

As the vertex suppression the timing suppression strongly depends on the track density thus on the muon stopping rate and the length of the reconstruction frames. If the timestamps used in the pixel detector and the pixel time resolution are smaller than the length of the reconstruction frames this can already be used to suppress accidental background.

In phase Ib the two timing detectors are added to obtain much more precise time information. The precise timing of a track is determined by the number of assignable hits in the fibre detector and the existence of a matched tile hit. If a track

reaches the tile detector in the recurl stations, its timing is dominated by the much more accurate tile detector.

To estimate the achievable time suppression, a large sample of particle tracks is generated and reconstructed as described in chapters 27 and 28. Each uniquely reconstructed track is propagated towards the timing detectors. Taking the detection efficiencies folded with matching efficiencies and time resolutions into account the achievable time resolution of this track is determined.

Based on this set of tracks the time suppression for different superpositions (e.g. positrons from Michel decay with electrons from Bhabha scattering with another positron from a Michel decay or miss-reconstructed track and so on) is calculated which leads to two main weighted time suppression factors. On one hand a mean suppression if two of the three tracks are correlated in time, as for electron positron pairs from Bhabha scattering or the superposition of such a pair from internal conversion with a positron from an ordinary Michel decay.

Figure 34.19 shows the weighted mean survival fraction (suppression factor = $\frac{1}{\text{survivalfraction}}$) depending on the fibre detector time resolution and efficiency. A fully efficient tile detector with a time resolution of 60 ps and reconstruction and pixel timestamps of 50 ns were assumed. Figure 34.21 shows the same for the tile detector where a 96 % efficient fibre detector with a time resolution of 400 ps is assumed. A much higher timing suppression can be achieved if all three tracks are uncorrelated. Figures 34.20 and 34.22 show this three time vertices survival fraction. Since the probability for these accidentals is much lower (see ta-




Figure 34.19: Weighted mean survival fraction of accidental background with two tracks correlated in time versus fibre time resolution (left) and fibre efficiency (right). A fully efficient tile detector with 60 ps time resolution and 50 ns reconstruction frames are assumed.




Figure 34.20: Weighted mean survival fraction of accidental background with three uncorrelated tracks versus fibre time resolution (left) and fibre efficiency (right). A fully efficient tile detector with 60 ps time resolution and 50 ns reconstruction frames are assumed.

ble 34.1) and the vertex suppression also performs much better, this three vertex timing suppression is of less importance.

For the over all sensitivity studies in figure ?? a fully efficient tile detector with a time resolution of 60 ps and a fibre detector with a ribbon efficiency of 96 % and a time resolution of 400 ps in 50 ns reconstruction frames and pixel timestamps is used. This corresponds to the measured values as stated in chapters 12 and 15.

Table 34.3 shows the relative occurrence of Michel/Bhabha pairs in combination with an positron from a Michel decay ordered by the timing detector providing the dominant time informa-

tion. Over 63 % of all reconstructed combinations of such tracks have six or more pixel hits.

Different combinations of dominant timing detector and multiple hits in the Fibre Detector lead to a different possible suppression for this combinations. Table 34.4 summarized this different factors where t stands for Tile Detector, f(2) and f(1) for two respectively one hit in the Fibre Detector and x for no timing information at all.

These assumptions above lead to a timing suppression factor of 91 for the dominant accidental background where two tracks are correlated in time and 1672 for three uncorrelated tracks. Figure 34.5 summarizes this results for all combination of the timing detectors.




Figure 34.21: Weighted mean survival fraction of accidental background with two tracks correlated in time versus tile time resolution (left) and tile efficiency (right). A 96 % efficient fibre detector with 400 ps time resolution and 50 ns reconstruction frames are assumed.




Figure 34.22: Weighted mean survival fraction of accidental background with three uncorrelated tracks versus tile time resolution (left) and tile efficiency (right). A 96 % efficient fibre detector with 400 ps time resolution and 50 ns reconstruction frames are assumed.

dominant timing	all tracks [%]	≥ 6 hits [%]
$e^+/e^-, e^+$	100.0	63.4
t/t, t	6.9	7.0
t/t, f	8.9	8.0
t/f, t	14.5	14.2
t/f, f	9.0	5.3
f/t, t	2.0	0.8
f/t, f	6.1	5.1
f/f, t	9.0	7.6
f/f, f	43.6	52.0

Table 34.3: Fraction of combinations of positron/electron pairs from Bhabha scattering in superposition with a positron from ordinary Michel decay grouped by the dominant timing detector. Tile Detector (t), Fibre Detector (f).




Figure 34.23: Small invariant mass of e^+e^- pairs versus $e^+e^-e^+$ invariant mass for accidental combinations of a Bhabha e^+e^- pair with a Michel positron.

Kinematic suppression

The largest suppression factors for accidental background come from kinematics, i.e. the requirement that the three momenta sum up to zero (enforced by the total momentum selection) and a mass window around the muon mass.

We estimate this suppression factor for the Bhabha e^+e^- plus Michel background by taking electron-positron pair kinematics from the full simulation and combining them randomly with analytically generated Michel positrons; all tracks are weighted with the reconstruction efficiency determined in the full simulation. Note that the Bhabha pairs almost all have the same

bhabha pair/michel	t	f(2)	f(1)	x
t/t	486	129	90	1.5
t/f(2)	231	114	84	1.5
t/f(1)	174	102	82	1.5
t/x	369	88	63	1.0
f(2)/f(2)	169	103	83	1.5
f(2)/f(1)	151	98	80	1.5
f(1)/f(1)	129	95	76	1.5
f(2)/x	159	79	59	1.0
f(1)/x	118	74	56	1.0
x/x	1.0	1.0	1.0	1.0

Table 34.4: Timing suppression for given positron/electron pairs from Bhabha scattering in superposition with a positron from ordinary Michel decay. Tile Detector (t), Fibre Detector (f(n)) with n hits and no timing information x.

	only fibres	only tiles	both
all hits	43	5.3	91
≥ 6 hits	54	5.3	132

Table 34.5: Timing suppression for positron/electron pairs from Bhabha scattering in superposition with a positron from ordinary Michel decay.

Process	0.5 MeV/c ² mass window	1 MeV/c ² mass window
2 Bhabha + Michel	$1.3 \cdot 10^{-6}$	$3.1 \cdot 10^{-6}$
Electron + 2 Michel	$9.1 \cdot 10^{-6}$	$1.8 \cdot 10^{-5}$
IC + Michel	1	1
IC + 2 Michel	1	1

Table 34.6: Kinematic suppression factors; all numbers are determined requiring the total momentum to be < 4 MeV/c

e^+e^- invariant mass of around 7 MeV/c², see figure ?? which comes from the minimum momentum transfer required to kick the electron (initially at rest) into the detector acceptance folded with the strongly forward peaked Bhabha cross section. A requirement on this mass can greatly reduce the Bhabha background, would however also remove a specific part of the signal kinematics.

The kinematic suppression factor for single electrons combined with two random Michel decays we estimate by generating single electrons with an exponentially falling momentum spectrum, where the exponent is fitted in the full simulation. These electrons are then combined with two Michel positrons generated from the analytic distribution. The tracks are weighted with the reconstruction efficiency determined from the full simulation.

Finally, for the combination of electrons or electron positron pairs from internal conversion decays in combination with Michel electrons, we generate internal conversion kinematics using the matrix element and Michel electrons from the analytic distribution and then weight with the reconstruction efficiency.

The kinematic suppression factors obtained thus are listed in table 34.6.

Expected combinatorial background



	Bhabha pair + 1 Michel	Bhabha e^- + 2 Michel	γ conv. pair + 1 Michel	γ conv. + 2 Michel	Compton + 2 Michel	IC pair + 1 Michel	IC + 2 Michel
e^- produced	$5.6 \cdot 10^{-5}$	$8.7 \cdot 10^{-5}$	$1.1 \cdot 10^{-6}$	$1.7 \cdot 10^{-6}$	$3.8 \cdot 10^{-6}$	$2.6 \cdot 10^{-5}$	$2.8 \cdot 10^{-5}$

Table 34.7: Expected combinatorial background for phase Ia running. All numbers are relative to the muon stops in the target.



APPENDIX

all: Bring theses and acknowledgements up to date

A.1 Mu3e theses

Several theses involving the Mu3e project have been completed:

- M. Kiehn, diploma thesis *Track Fitting with Broken Lines for the Mu3e Experiment*, Heidelberg University, 2012.
- M. Zimmermann, bachelor thesis *Cooling with Gaseous Helium for the Mu3e Experiment*, Heidelberg University, 2012.
- H. Augustin, bachelor thesis *Charakterisierung von HV-MAPS*, Heidelberg University, 2012.
- A.-K. Perrevoort, master thesis *Characterisation of High Voltage Monolithic Active Pixel Sensors for the Mu3e Experiment*, Heidelberg University, 2012.
- T. Hartwig, bachelor thesis *Messung der Zeitauflösung eines Szintillator-SiPM-Systems für den Kacheldetektor des Mu3e-Experiments*, Heidelberg University, 2013.
- C. Licciulli, master thesis *Präzise Zeitmessung für das Mu3e-Experiment*, Heidelberg University, 2013.
- A. Damyanova, master thesis *Development of a Scintillating Fibre Tracker/Time-of-Flight Detector with SiPM Readout for the Mu3e Experiment at PSI*, Geneva University, 2013.
- S. Schenk, bachelor thesis *A Vertex Fit for Low Momentum Particles in a Solenoidal Magnetic Field with Multiple Scattering*, Heidelberg University, 2013.
- M. Baumgartner, bachelor thesis *Comparison of Measured and Simulated Light Yield in Plastic Scintillators*, Zürich University, 2013.
- C. Grzesik, bachelor thesis *Fast Optical Readout for the Mu3e Experiment*, Heidelberg University, 2014.
- L. Huxold, bachelor thesis *Cooling of the Mu3e Pixel Detector*, Heidelberg University, 2014.
- R. Philipp, master thesis *Characterisation of High Voltage Monolithic Active Pixel Sensors for the Mu3e Experiment*, Heidelberg University, 2014.
- S. Corrodi, master thesis *Fast Optical Readout of the Mu3e Pixel Detector*, ETH Zürich and Heidelberg University, 2014.
- F. Förster, master thesis *HV-MAPS Readout and Direct Memory Access for the Mu3e Experiment*, Heidelberg University, 2014.
- J. Kröger, bachelor thesis *Data Transmission at High Rates via Kapton Flexprints for the Mu3e Experiment*, Heidelberg University, 2014.



- L. Huth, master thesis *Development of a Tracking Telescope for Low Momentum Particles and High Rates consisting of HV-MAPS*, Heidelberg University, 2014.
- F. Reiß, bachelor thesis *Sensitivity Study of the Phase I Detector for the Mu3e Experiment*, Heidelberg University, 2014.
- H. Augustin, master thesis *Characterization of a novel HV-MAPS Sensor with two Amplification Stages and first examination of thinned MuPix Sensors*, Heidelberg University, 2014.
- A. Herkert, master thesis *Gaseous Helium Cooling of a Thin Silicon Pixel Detector for the Mu3e Experiment*, Heidelberg University, 2015
- P. Eckert, Ph. D. thesis *The Mu3e Tile Detector*, Heidelberg University, 2015.
- J. Hammerich, bachelor thesis *Studies of HV-MAPS analog performance*, Heidelberg University, 2015.
- J. Philipp, bachelor thesis *Effizienzanalyse von HV-MAPS anhand des MuPix Teleskops*, Heidelberg University, 2015.
- Y.W. Ng, master thesis *Finite Element Analysis of the Cooling System for the Mu3e Experiment*, University of Applied Science Jena and Heidelberg University, 2015.
- R.P. Austermühl, bachelor thesis *Analyse von Michelson- Interferometriedaten von Vibrationsmessungen eines dünnen gasgekühlten Pixeldetektors*, Heidelberg University, 2015.
- L. Henkelmann, bachelor thesis *Optical Measurement of Vibration and Deformation of the Mu3e Silicon Pixel Tracker*, Heidelberg University, 2015.
- M. Kiehn, doctoral thesis *Track Reconstruction for the Mu3e Experiment*, Heidelberg University, started 2012.
- A. Damyanova, doctoral thesis *Fibre Tracker for the Mu3e Experiment*, Geneva University, started 2013.
- D. vom Bruch, doctoral thesis *Fast Track reconstruction on Graphics Processors for the Mu3e Experiment*, Heidelberg and Mainz Universities, started 2014.
- Q. H. Huang, doctoral thesis *The Mu3e Data Acquisition System*, Heidelberg and Mainz Universities, started 2014.
- A. K. Perrevoort, doctoral thesis *Sensitivity Studies and Front-End Firmware for the Mu3e Experiment*, Heidelberg University, started 2014.
- S. Corrodi, doctoral thesis *Readout of the Mu3e Fibre Detector*, ETH Zürich, started 2014.
- L. Huth, doctoral thesis *The Mu3e Pixel Detector*, Heidelberg University, started 2015.
- S. Dittmeier, doctoral thesis *Readout of the Mu3e Pixel Detector*, Heidelberg University, started 2014.
- H. Augustin, doctoral thesis *Sensors for the Mu3e Pixel detector*. Heidelberg University, started 2015.
- U. Hartenstein, doctoral thesis *Calibration and Alignment of the Mu3e Pixel Detector*, Mainz University, started 2015.
- C. Grzesik, master thesis *Online Track Reconstruction using Graphics Processors for the MUPIX Telescope*, started 2015.

These theses are available in portable document format (PDF) from <http://www.psi.ch/mu3e/documents>.

Several more theses are ongoing (thus titles are preliminary):

- R. Gredig, doctoral thesis *Fibre Tracker for the Mu3e Experiment*, Zürich University, started 2012.

A.2 Acknowledgements

all: Add Ph.D. funding sources like IMPRS, HGSFP etc.

N. Berger, S. Shrestha, A. Kozlinskiy, D. vom Bruch, Q. H. Huang and U. Hartenstein would like to thank the *Deutsche Forschungsgemeinschaft* (DFG) for funding their work on the Mu3e experiment through the Emmy Noether



program and thus supporting the experiment at a crucial early stage.

A. Kozlinskiy and N. Berger are grateful for support from the PRISMA cluster of excellence at Johannes-Gutenberg University Mainz.

S. Dittmeier, L. Huth and M. Kiehn acknowledge support by the International Max Planck Research School for Precision Tests of Fundamental Symmetries.

A.-K. Perrevoort acknowledges support by the Research Training Group on Particle Physics beyond the Standard Model.

H. Augustin and A. Herkert acknowledge support by the Research Training Group High Resolution and High Rate Detectors in Nuclear and Particle Physics.



BIBLIOGRAPHY

- [1] A. Blondel, S. Bravar, M. Pohl, S. Bachmann, N. Berger, A. Schi \ddot{u} ning, D. Wiedner, P. Fischer, I. Peric, M. Hildebrandt, P.-R. Kettle, A. Papa, S. Ritt, G. Dissertori, Ch. Grab, R. Wallny, P. Robmann and U. Straumann, Letter of intent for an experiment to search for the decay $\mu \rightarrow eee$, 2012.
- [2] A. Blondel et al., “*Research Proposal for an Experiment to Search for the Decay $\mu \rightarrow eee$* ”, ArXiv e-prints, January 2013, (arXiv:1301.6113 [physics.ins-det]).
- [3] Y. Kuno and Y. Okada, “*Muon decay and physics beyond the standard model*”, Rev. Mod. Phys., **73** 151–202, 2001, (arXiv:hep-ph/9909265).
- [4] R. M. Djilkibaev and R. V. Konoplich, “*Rare Muon Decay $\mu^+ \rightarrow e^+e^-e^+\nu_e\bar{\nu}_\mu$* ”, Phys.Rev., **D79** 073004, 2009, (arXiv:0812.1355 [hep-ph]).
- [5] W. Bertl et al., [SINDRUM Collaboration], “*Search for the decay $\mu^+ \rightarrow e^+e^+e^-$* ”, Nucl. P, **B 260**(1) 1 – 31, 1985.
- [6] R. R. Crittenden, W. D. Walker and J. Ballam, “*Radiative Decay Modes of the Muon*”, Phys. Rev., **121** 1823–1832, Mar 1961.
- [7] S. Egli et al., [SINDRUM], “*Measurement of the Decay $\pi^+ \rightarrow e^+\nu_e e^+e^-$ and Search for a Light Higgs Boson*”, Phys. Lett., **B222** 533, 1989.
- [8] G. Bressi, G. Carugno, S. Cerdonio, E. Conti, A.T. Meneguzzo and D. Zanello, “*New measurement of the $\pi \rightarrow \mu\nu\gamma$ decay*”, Nuclear Physics B, **513**(3) 555 – 572, 1998.
- [9] A.E. Pifer, T. Bowen and K.R. Kendall, “*A High Stopping Density μ^+ Beam*”, Nucl.Instrum.Meth., **135** 39–46, 1976.
- [10] A.M. Baldini, F. Cei, C. Cerri, S. Dussoni, L. Galli et al., “*MEG Upgrade Proposal*”, ArXiv e-prints, January 2013, (arXiv:1301.7225 [physics.ins-det]).
- [11] http://aea.web.psi.ch/Urs_Rohrer/MyWeb/trancomp.htm, PSI Graphic Transport Framework by U. Rohrer, based on a CERN-SLAC-FERMILAB version by K. L. Brown *et al.*
- [12] http://aea.web.psi.ch/Urs_Rohrer/MyWeb/turtcomp.htm, PSI Graphic Turtle Framework by U. Rohrer based on a CERN-SLAC-FERMILAB version by K. L. Brown *et al.*
- [13] T. Roberts, G4Beamline, <http://g4beamline.muonsinc.com>.
- [14] U. Bellgardt et al., [SINDRUM Collaboration], “*Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$* ”, Nucl.Phys., **B299** 1, 1988.
- [15] W. Bertl, SINDRUM I, Presentation at PSI, 2008.
- [16] I. Perić, “*A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology*”, Nucl.Instrum.Meth., **A582** 876, 2007.
- [17] I. Perić and C. Takacs, “*Large monolithic particle pixel-detector in high-voltage CMOS technology*”, Nucl. Instrum. Meth., **A624**(2) 504, 2010.
- [18] I. Perić, C. Kreidl and P. Fischer, “*Particle pixel detectors in high-voltage CMOS technology - New achievements*”, Nucl. Instr. Meth., **A 650**(1) 158, 2011.
- [19] I. Perić, Private communication, 2010.
- [20] H. Augustin, *Characterization of a novel HV-MAPS Sensor with two Amplification Stages and first examination of thinned MuPix Sensors*, Master thesis, Heidelberg University, 2014.
- [21] IBM Corporation, *IBM CMOS 7 HV*, 2012, TGD03019USEN.
- [22] L. Huth, *Development of a Tracking Telescope for Low Momentum Particles and High Rates consisting of HV-MAPS*, Master thesis, Heidelberg University, 2014.



- [23] A.-K. Perrevoort, *Characterisation of High Voltage Monolithic Active Pixel Sensors for the Mu3e Experiment*, Master thesis, Heidelberg University, 2012.
- [24] I. Perić et al., “*Overview of HVCMOS pixel sensors*”, JINST, **10**(05) C05021, 2015.
- [25] I. Perić et al., “*High-voltage pixel sensors for ATLAS upgrade*”, Nucl. Instrum. Meth., **A765** 172–176, 2014.
- [26] K. Akiba, M. Artuso, R. Badman, A. Borgia, R. Bates et al., “*Charged Particle Tracking with the Timepix ASIC*”, Nucl.Instrum.Meth., **A661** 31–49, 2012, (arXiv:1103.2739 [physics.ins-det]).
- [27] R. Philipp, *Characterisation of High Voltage Monolithic Active Pixel Sensors for the Mu3e Experiment*, Master thesis, Heidelberg University, 2014.
- [28] M. Kiehn, *Pixel Sensor Evaluation and Track Fitting for the Mu3e Experiment*, PhD thesis, Heidelberg University, In preparation.
- [29] J. Repenning, *Charakterisierung von HV-MAPS*, Bachelor thesis, Heidelberg University, in preparation.
- [30] J. Philipp, *Effizienzanalyse von HV-MAPS anhand des MuPix-Teleskops*, Bachelor thesis, Heidelberg University, 2015.
- [31] J. Kroeger, *Data Transmission at High Rates via Kapton Flexprints for the Mu3e Experiment*, Bachelor thesis, Heidelberg University, 2014.
- [32] O. Harper, *Flexprint studies for the Mu3e experiment*, Master thesis, Université Paris-Sud, Heidelberg University, 2015.
- [33] G. Bolla et al., “*Wire-bond failures induced by resonant vibrations in the CDF silicon detector*”, Nucl. Instrum. Meth., **A518** 277–280, 2004.
- [34] Joseph M. Izen, Matthew Kurth and Rusty Boyd, “*Polyurethane spray coating of aluminum wire bonds to prevent corrosion and suppress resonant oscillations*”, 2015, (arXiv:1511.05436 [physics.ins-det]).
- [35] A. X. Widmer and P. A. Franaszek, “*A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code*”, IBM Journal of Research and Development, **27** 440, 1983.
- [36] P.A. Franaszek and A.X. Widmer, Byte oriented DC balanced (0,4) 8B/10B partitioned block transmission code, December 4 1984, US Patent 4,486,739.
- [37] J. Beringer et al., [Particle Data Group], “*Review of Particle Physics (RPP)*”, Phys.Rev., **D86** 010001, 2012.
- [38] Kuraray, *SCINTILLATION MATERIALS*, 2014.
- [39] Saint-Gobain Crystals, *Scintillation Products*, 2014.
- [40] LHCb Collaboration, LHCb Tracker Upgrade Technical Design Report, Technical Report CERN-LHCC-2014-001. LHCb-TDR-015, CERN, Feb 2014.
- [41] A. Kuonen, O. Girard and G. Haefeli, Characterisation of the Hamamatsu MPPC multi-channel array for LHCb SciFi Tracker v.9.2014, LHCb internal note LPHE Note 2015-01, April 2015.
- [42] S. Ritt, “*The DRS Chip: Cheap Waveform Digitizing in the GHz Range*”, Nucl.Instrum.Meth., **A518** 470–471, 2004.
- [43] K. Sedlak, A. Stoykov and R. Scheuermann, “*A GEANT4 study on the time resolution of a fast plastic scintillator read out by a G-APD*”, Nucl. Instr. Meth A., **696** 40 – 45, 2012.
- [44] A. Papa, F. Barchetti, F. Gray, E. Ripiccini and G. Rutar, “*A multi-purposed detector with silicon photomultiplier readout of scintillating fibers*”, Nucl. Instrum. Methods Phys. Res., **797** 130, 2015.
- [45] CERN Program Library Long Writeup W5013, *GEANT – Detector Description and Simulation Tool*, 1993.
- [46] A. Papa, G. Cavoto, E. Ripiccini and M. De Gerone, “*A simulation tool for scintillating fibers coupled to SiPM for MIP and heavy ionizing particle identification*”, Journal of Instrumentation, **9**(05) C05066, 2014.
- [47] C. Casella, S. Corrodi, J. Fischer, A. Howard, W. Lustermaun, and I. Mikio, SPTR STiC paper, in preparation/personal contact.
- [48] T. Harion, K. Briggel, H. Chen, P. Fischer, A. Gil, V. Kiworra, M. Ritzert, H.-C. Schultz-Coulon, W. Shen and V. Stankova, “*STiC: a mixed mode silicon photomultiplier readout ASIC for time-of-flight applications*”, Journal of Instrumentation, **9**(02) C02003, 2014.
- [49] C. Casella, S. Corrodi, J. Fischer, A. Howard, W. Lustermaun, and I. Mikio, High Rate Performance Comparison between STiC and TOF-PET ASIC, in preparation.
- [50] T. Iwamoto, [MEG Collaboration], “*Performance of the MEG detector to search for $\mu^+ \rightarrow e^+\gamma$ decays at PSI*”, PoS, **ICHEP2010** 489, 2010.
- [51] I. Sacco, P. Fischer and M. Ritzert, “*PETA4: a multi-channel TDC/ADC ASIC for SiPM readout*”, JINST, **8** C12013, 2013.



- [52] C. Piemonte et al., “Performance of FBK SiPMs coupled to PETA3 read-out ASIC for PET application”, *Nucl. Instrum. Meth.*, **A 718** 345, 2013.
- [53] K. Briggel, M. Dorn, R. Hagdorn, T. Harion, H. C. Schultz-Coulon and W. Shen, “KLauS: an ASIC for silicon photomultiplier readout and its application in a setup for production testing of scintillating tiles”, *JINST*, **9** C02013, 2014.
- [54] M. Rolo, R. Bugalho, F. Gon alves, A. Rivetti, G. Mazza, J. Silva, R. Silva and J. Varela, “A 64-channel ASIC for TOFPET applications”, *IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, pages 1460–1464, 2012.
- [55] Ricardo, “TOFPET Poster @ Agus LIP lab”, “”, 2014.
- [56] S. Ahmad, J. Fleury, C. de la Taille, N. Seguin-Moreau, F. Dulucq, G. Martin-Chassard, S. Callier, D. Thienpont and L. Raux, “Triroc: A Multi-Channel SiPM Read-Out ASIC for PET/PET-ToF Application”, *Nuclear Science, IEEE Transactions on*, **62**(3) 664–668, June 2015.
- [57] F. Corsi, M. Foresta, C. Marzocca, G. Matarrese and A. Del Guerra, “BASIC: an 8-channel Front-end ASIC for Silicon Photomultiplier Detectors”, *IEEE Nuclear Science Symposium Conference Record*, pages 1082–1087, 2009.
- [58] J. Castilla, J.-M. Celaand A Comerma, J-M Fernandez-Varea, L Freixas, L Garrido, D Gascon, J Marin, G Martinez amd J-M Perez amd P Rato-Mendes and I Sarasola, “Evaluation of the FlexToT ASIC on the Readout of SiPM Matrices and Scintillators for PET”, *IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2013.
- [59] J. Fleury, S. Callier, C. de La Taille amd N. Seguin, D. Thienpont, F. Dulucq, S. Ahmad and G. Martin, “Petiroc, a new front-end ASIC for time of flight application”, *IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, pages 1–5, 2013.
- [60] J. Fleury, S. Callier, C. de La Taille amd N. Seguin, D. Thienpont, F. Dulucq, S. Ahmad and G. Martin, “Petiroc and Citiroc: front-end ASICs for SiPM read-out and ToF applications”, *JINST*, **9**(01) C01049, 2014.
- [61] F. Anghinolfi, P. Jarron, F. Krummenacherand E. Usenko and M.C.S Williams, “NINO: AnUltrafast Low-Power Frond-End Amplifier Discriminator for the Time-of-Flight Detector in the ALICE Experiment”, *IEEE transactions on nuclear science*, **51**, 2004.
- [62] E. Auffray, B. Frisch, F. Geraci, A. Ghezzi, S. Gundacker, H. Hillemanns, P. Jarron, T. Meyer, M. Paganoni, K. Pauwels, M. Pizzichemi and P. Lecoq, “A Comprehensive and Systematic Study of Coincidence Time Resolution and Light Yield Using Scintillators of Different Size, Wrapping and Doping”, *IEEE Nuclear Science Symposium Conference Record*, 2011.
- [63] M. Bouchel, S. Callier, F. Dulucq, J. Fleury, J. J. Jaeger, C. de La Taille, G. Martin-Chassard and L. Raux, “SPIROC (SiPM Integrated Read-Out Chip): dedicated very front-end electronics for an ILC prototype hadronic calorimeter with SiPM read-out”, *Journal of Instrumentation*, **6**(01) C01098, 2011.
- [64] P. Eckert, *The Mu3e Tile Detector*, Phd thesis, Heidelberg University, 2015.
- [65] P. Eckert, R. Stamen and H.C. Schultz-Coulon, “Study of the response and photon-counting resolution of silicon photomultipliers using a generic simulation framework”, *JINST*, **7** P08011, 2012, (arXiv:1206.4154 [physics.ins-det]).
- [66] T. Harion, *The STiC ASIC: High Precision Timing with Silicon Photomultipliers*, PhD thesis, Heidelberg University, 2015.
- [67] S. Majewski, M. Bowen, C. Zorn, K. Johnson, V. Hagopian, J. Thomaston and H. Wahl, “Radiation damage studies in plastic scintillators with a 2.5-MeV electron beam”, *Nucl. Instr. Meth. A.*, **281**(3) 500–507, 1989.
- [68] EndoTOFPET-US Proposal, “Novel Multimodal Endoscopic Probes for Simultaneous PET/Ultrasound Imaging for Image-Guided Interventions”, *European Union 7th*, **186** 2007–2013.
- [69] W. Shen, K. Briggel, H.S. Chen, P. Fischer, A. Gil, T. Harion, M. Ritzert and H.-C. Schultz-Coulon, STiC-A Mixed Mode Chip for SiPM ToF Applications, In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE*, pages 877–881. IEEE, 2012.
- [70] W. Shen, K. Briggel, H.S. Chen, P. Fischer, A. Gil, T. Harion, V. Kiworra, M. Ritzert, H.-C. Schultz-Coulon and V. Stankova, STiC2-characterization Results of a SiPM Readout ASIC for Time-of-Flight Applications, In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE*, pages 1–5. IEEE, 2013.
- [71] S. Ritt, R. Dinapoli and U. Hartmann, “Application of the DRS chip for fast waveform digitizing”, *Nucl.Instrum.Meth.*, **A623** 486, 2010, (arXiv:1011.0226 [physics.ins-det]).



- [72] A. Herkert, *Gaseous Helium Cooling of a Thin Silicon Pixel Detector for the Mu3e Experiment*, Master thesis, Heidelberg University, 2015.
- [73] Y.W. Ng, *Finite Element Analysis of the Cooling System for the Mu3e Experiment*, Master thesis, University of Applied Science Jena, 2015.
- [74] R. P. Austerer, *Analyse von Michelson-Interferometriedaten von Vibrationsmessungen eines dünnen gasgekühlten Pixeldetektors*, Bachelor thesis, Heidelberg University, 2015.
- [75] L. Henkelmann, *Optical Measurements of Vibration and Deformation of the Mu3e Silicon Pixel Tracker*, Bachelor thesis, Heidelberg University, 2015.
- [76] K. Weir, W.J.O. Boyle, B.T. Meggitt, A.W. Palmer and K.T.V. Grattan, “A novel adaptation of the Michelson interferometer for the measurement of vibration”, *Journal of Lightwave Technology*, **10** 700–703, 1992.
- [77] R. Schmidt S. Ritt, MSCB (MIDAS Slow Control Bus), 2001, <http://midas.psi.ch/mscb>.
- [78] EPICS (Experimental Physics and Industrial Control System), <http://www.aps.anl.gov/epics>.
- [79] Avago, MiniPod Datasheet, Technical report, 2013.
- [80] Ken Boyette, OpenCores 8bit/10bit encoder and decoder http://opencores.org/project,8b10b_encdec.
- [81] Altera, Stratix GX Transceiver User Guide, Technical report, 2015.
- [82] S. Corrodi, *Fast Optical Readout of the Mu3e Pixel Detector*, Master thesis, ETH Zürich and Heidelberg University, 2014.
- [83] C. Grzesik, *Fast Optical Readout for the Mu3e Experiment*, Bachelor thesis, Heidelberg University, 2014.
- [84] Altera, Nios ii classic processor reference guide, Technical report, 2015.
- [85] P. Durante, N. Neufeld, R. Schwemmer, U. Marconi, G. Balbi and I. Lax, 100Gbps PCI-express readout for the LHCb upgrade, In *Proceedings, 19th Real Time Conference (RT2014)*, 2014.
- [86] P. Durante, N. Neufeld, R. Schwemmer, G. Balbi and U. Marconi, “100 Gbps PCI-Express readout for the LHCb upgrade”, *JINST*, **10**(04) C04018, 2015.
- [87] J.P. Cachemiche, PCIe40 status report, In *LHCb Electronics upgrade WG*, 2015.
- [88] G. Haefeli, A. Bay, A. Gong, H. Gong, M. Muecke et al., “The LHCb DAQ interface board TELL1”, *Nucl.Instrum.Meth.*, **A560** 494–502, 2006.
- [89] W. Kühn et al., [PANDA Collaboration], “FPGA based compute nodes for high level triggering in PANDA”, *J.Phys.Conf.Ser.*, **119** 022027, 2008.
- [90] Altera Corporation, *DSP Development Kit, Stratix V Edition - Reference Manual*, 1.0 edition, 2012.
- [91] Altera Corporation, *DSP Development Kit, Stratix V Edition - User Guide*, 1.1 edition, 2013.
- [92] NVIDIA Corporation, *NVIDIA GeForce GTX 680*, 2012, Whitepaper.
- [93] AMD Corporation, *AMD Graphics Cores Next (GCN) Architecture*, 2012, Whitepaper.
- [94] PCI Express Base Specification, 2014, Revision 3.1.
- [95] K. Olchanski S. Ritt, P. Amaudruz, Maximum Integration Data Acquisition System, 2001, <http://midas.psi.ch>.
- [96] K. Abe et al., [T2K Collaboration], “The T2K Experiment”, *Nucl.Instrum.Meth.*, **A659** 106–135, 2011, (arXiv:1106.1238 [physics.ins-det]).
- [97] L.M. Brarkov et al., “Search for $\mu^+ \rightarrow e^+ \gamma$ down to 10^{-14} branching ratio”, Research Proposal to PSI, 1999.
- [98] Microsoft skype - www.skype.com.
- [99] Atlassian bitbucket - www.bitbucket.org.
- [100] Atlassian hipchat - www.hipchat.com.
- [101] git source code management www.git-scm.com.
- [102] Linux www.linux.com.
- [103] Linux kernel www.kernel.org.
- [104] Apple mac os x www.apple.com/osx/.
- [105] Opensuse www.opensuse.org.
- [106] B. Stroustrup, *The C++ Programming Language*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 3rd edition, 2000.
- [107] B. Kernighan and D. Ritchie, *The C Programming Language*. Prentice-Hall, 1978.
- [108] Standard for Programming Language C++, 2014.
- [109] Free Software Foundation, Gnu compiler collection - gcc.gnu.org.
- [110] Intel compilers - software.intel.com/en-us/intel-compilers.
- [111] Tom R. Halfhill, “Parallel Processing With CUDA”, *Microprocessor Report*, **01/28/08** 1–6, 2008.






- [112] A. Munshi, The OpenCL specification, version 1.2, Technical report, The Khronos Group, 2011.
- [113] B. Gaster, The OpenCL C++ Wrapper API, Technical report, Khronos OpenCL Working Group, 2010.
- [114] CMake - www.cmake.org.
- [115] S. Agostinelli et al., “*Geant4—a simulation toolkit*”, Nucl. Instr. Meth., **A 506**(3) 250 – 303, 2003.
- [116] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P.A. Dubois et al., “*Geant4 developments and applications*”, IEEE Trans. Nucl. Sci., **53** 270, 2006.
- [117] J. Allison, J. Apostolakis, A. Bagulya, C. Champion, S. Elles et al., “*Geant4 electromagnetic physics for high statistic simulation of LHC experiments*”, J.Phys.Conf.Ser., **396** 022013, 2012.
- [118] V.N. Ivanchenko, O. Kadri, M. Maire and L. Urban, “*Geant4 models for simulation of multiple scattering*”, J.Phys.Conf.Ser., **219** 032045, 2010.
- [119] L. Urbán, A model for multiple scattering in GEANT4. oai:cds.cern.ch:1004190, Technical Report CERN-OPEN-2006-077, CERN, Geneva, Dec 2006.
- [120] N. Berger et al., “*Multiple Coulomb Scattering in Thin Silicon*”, JINST, **9** P07007, 2014, ([arXiv:1405.2759](https://arxiv.org/abs/1405.2759) [[physics.ins-det](https://arxiv.org/abs/1405.2759)]).
- [121] R. Brun and F. Rademakers, “*ROOT: An object oriented data analysis framework*”, Nucl. Instrum. Meth., **A389** 81–86, 1997.
- [122] J. List and B. List, Self-filling histograms: An object-oriented analysis framework, In *Computing in high energy physics and nuclear physics. Proceedings, Conference, CHEP’04, Interlaken, Switzerland, September 27-October 1, 2004*, pages 373–376, 2005.
- [123] Nokia, Qt - Cross-platform application and UI framework, <http://qt.nokia.com/>.
- [124] M. Segal and K. Akeley, The OpenGL graphics system: A specification (version 3.0), Technical report, The Khronos Group, 2008.
- [125] V. Karimäki, “*Effective circle fitting for particle trajectories*”, Nucl. Instr. and Meth., **A305** 187, 1991.
- [126] V. Blobel, C. Kleinwort and F. Meier, “*Fast alignment of a complex tracking detector using advanced track models*”, Comput.Phys.Commun., **182** 1760–1763, 2011, ([arXiv:1103.3909](https://arxiv.org/abs/1103.3909) [[physics.ins-det](https://arxiv.org/abs/1103.3909)]).
- [127] C. Kleinwort, “*General Broken Lines as advanced track fitting method*”, Nucl.Instrum.Meth., **A673** 107–110, 2012, ([arXiv:1201.4320](https://arxiv.org/abs/1201.4320) [[physics.ins-det](https://arxiv.org/abs/1201.4320)]).
- [128] Student (W.S. Gosset), “*The Probable Error of a Mean*”, Biometrika, **6**(1) 1–25, 1908.
- [129] F. Scheck, “*Muon Physics*”, Phys.Rept., **44** 187, 1978.
- [130] W.E. Fischer and F. Scheck, “*Electron Polarization in Polarized Muon Decay: Radiative Corrections*”, Nucl.Phys., **B83** 25, 1974.
- [131] P. Depommier and A. Vacheret, Radiative muon decay, Technical report, TWIST Technote No 55, 2001.
- [132] C. Fronsdal and H. Uberall, “*mu-Meson Decay with Inner Bremsstrahlung*”, Phys.Rev., **113** 654–657, 1959.
- [133] A. Signer et al., Private communication, 2015.
- [134] R. Kleiss, W.J. Stirling and S.D. Ellis, “*A new Monte Carlo treatment of multiparticle phase space at high energies*”, Comp. Phys. Commun., **40** 359 – 373, 1986.
- [135] P. Biass and T. Hebbeker, “*Parametrization of the Cosmic Muon Flux for the Generator CMSCGEN*”, 2009, ([arXiv:0907.5514](https://arxiv.org/abs/0907.5514) [[astro-ph.IM](https://arxiv.org/abs/0907.5514)]).
- [136] A. Schöning, A Fast 3-Dimensional Track Fit with Multiple Scattering, Talk given at the Mu3e Collaboration meeting, Internal note Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, October 2012.
- [137] M. Kiehn, *Track Fitting with Broken Lines for the MU3E Experiment*, Diploma thesis, Heidelberg University, 2012.
- [138] V. L. Highland, “*Some Practical Remarks on Multiple Scattering*”, Nucl.Instrum.Meth., **129** 497, 1975.
- [139] J. Hammerich, *Studies of HV-MAPS Analog Performance*, Bachelor thesis, Heidelberg University, 2015.
- [140] C. Graf, *A Calibration Scheme for the Mu3e Tile Detector*, Master’s thesis, University of Heidelberg, 2015.
- [141] C. Kleinwort, H1 alignment experience, In *Proceedings, first LHC Detector Alignment Workshop, CERN, Geneva, Switzerland, 4-6 September 2006*, pages 41–50, 2006.
- [142] T. Lampén, [CMS], “*Alignment of the CMS silicon tracker*”, J. Phys. Conf. Ser., **523** 012024, 2014.
- [143] S. Chatrchyan et al., [CMS], “*Alignment of the CMS tracker with LHC and cosmic ray data*”, JINST, **9** P06009, 2014, ([arXiv:1403.2286](https://arxiv.org/abs/1403.2286) [[physics.ins-det](https://arxiv.org/abs/1403.2286)]).

- [144] N. Bartosik, [CMS], “*Simultaneous alignment and Lorentz angle calibration in the CMS silicon tracker using Millepede II*”, PoS, **EPS-HEP2013** 074, 2013.
- [145] G. Flucke, [CMS], “*Alignment of the CMS silicon tracker*”, J. Phys. Conf. Ser., **368** 012036, 2012.
- [146] J. Behr, [CMS], “*Alignment procedures for the CMS silicon tracker*”, J. Phys. Conf. Ser., **396** 022005, 2012.
- [147] J. Draeger, *Track based alignment of the CMS silicon tracker and its implication on physics performance*, PhD thesis, Hamburg U., 2011.
- [148] M. Weber, [CMS], “*Calibration, alignment and tracking performance of the CMS silicon strip tracker*”, Nucl. Instrum. Meth., **A628** 59–63, 2011.
- [149] V. Blobel, “*Software alignment for tracking detectors*”, Nucl. Instrum. Meth., **A566** 5–13, 2006.

H G F E D C B A

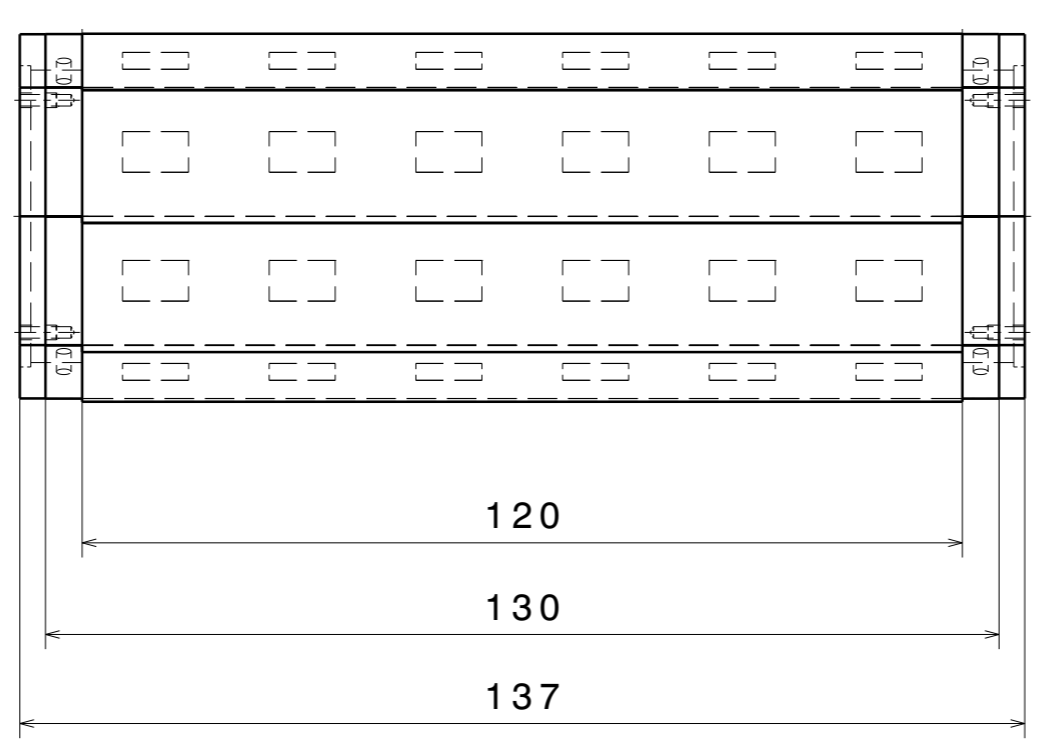
4



4


3

3



2

2



1

1

Oberflächenangaben DIN ISO 1302
 Allgemeintoleranzen DIN ISO 2768-mK

Erstellt von: Kevin Stumpf 25.11.2015		XXX	Material XXX
Stand (Datum): 26.11.2015			
Größe A3		Φ Physikalisches Institut Im Neuenheimer Feld 227a, 69120 Heidelberg	Menge (Stück) XXX
Maßstab 1:1	XXX		Zeichnungsnummer HE_MU3E_RING1_904100_D_8_detector_1_half
Eigentum der Universität Heidelberg. Keine Vervielfältigung oder Weitergabe ohne Zustimmung.			

H G B A

H G F E D C B A

4

3

2


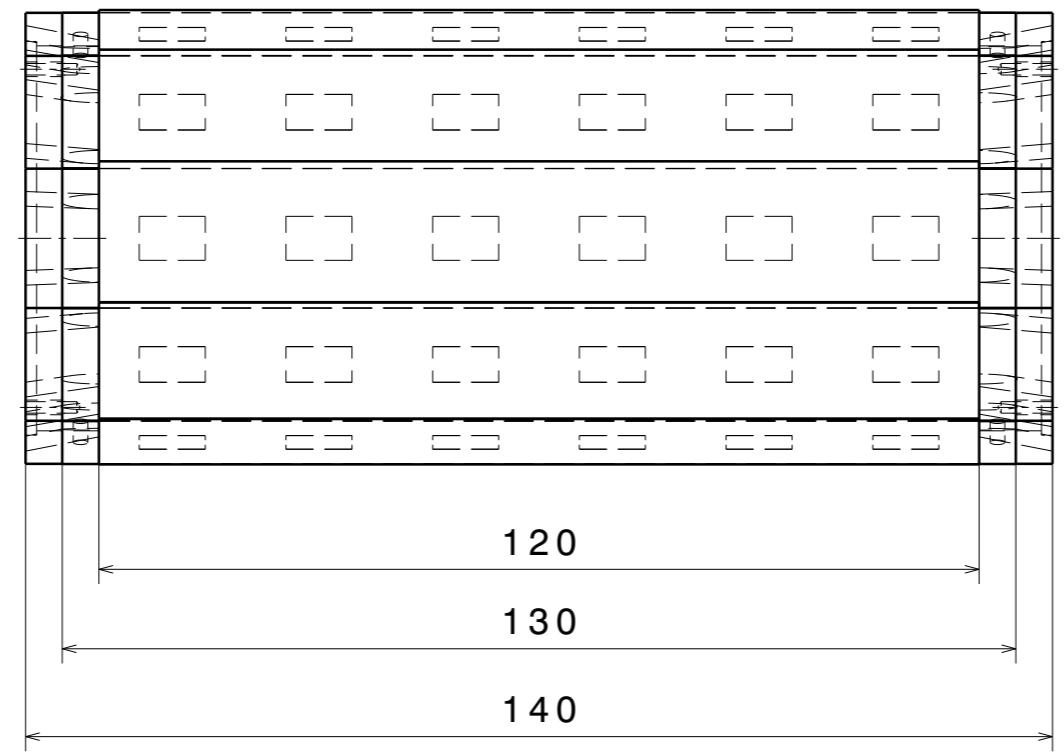



1

4

3

2

1



Oberflächenangaben DIN ISO 1302
Allgemeintoleranzen DIN ISO 2768-mK




Erstellt von: Kevin Stumpf		<h1>Detector_2_half</h1>		Material XXX
Stand (Datum): 26.11.2015				
Größe A3		Φ Physikalisches Institut Im Neuenheimer Feld 227a, 69120 Heidelberg		Menge (Stück) XXX
Maßstab 1:1				
Zeichnungsnummer HE_MU3E_RING2_904200_D_8		Blatt 1/1		

Eigentum der Universität Heidelberg. Keine Vervielfältigung oder Weitergabe ohne Zustimmung.




H G B A

4

4



Schnittansicht A-A
Maßstab: 1:2



Oberflächenangaben DIN ISO 1302
Allgemeintoleranzen DIN ISO 2768-mK

Erstellt von: Kevin Stumpf 30.11.2015		<h1>Detector_3</h1>		Material XXX
Stand (Datum): 01.12.2015				
Größe A3		Φ Physikalisches Institut Im Neuenheimer Feld 227a, 69120 Heidelberg		Menge (Stück) XXX
Maßstab 1:2	XXX			Zeichnungsnummer HE_MU3E_RING3_904300_D_5

Eigentum der Universität Heidelberg. Keine Vervielfältigung oder Weitergabe ohne Zustimmung.

1

1

2

2




3

3



H G F E D C B A

4

4




Schnittansicht A-A
Maßstab: 1:2



2

2



1

1

Oberflächenangaben DIN ISO 1302
Allgemeintoleranzen DIN ISO 2768-mK

Erstellt von: Kevin Stumpf		<h1>Detector_4</h1>		Material
Stand (Datum): 30.11.2015				XXX
01.12.2015		Φ Physikalisches Institut Im Neuenheimer Feld 227a, 69120 Heidelberg		Menge (Stück)
Größe A3				XXX
Maßstab 1:2	XXX	Zeichnungsnummer HE_MU3E_RING4_904400_D_4	Blatt 1/1	


Eigentum der Universität Heidelberg. Keine Vervielfältigung oder Weitergabe ohne Zustimmung.

H G F E D C B A



Mu3e Phase Ia Pixel Optical Readout Links

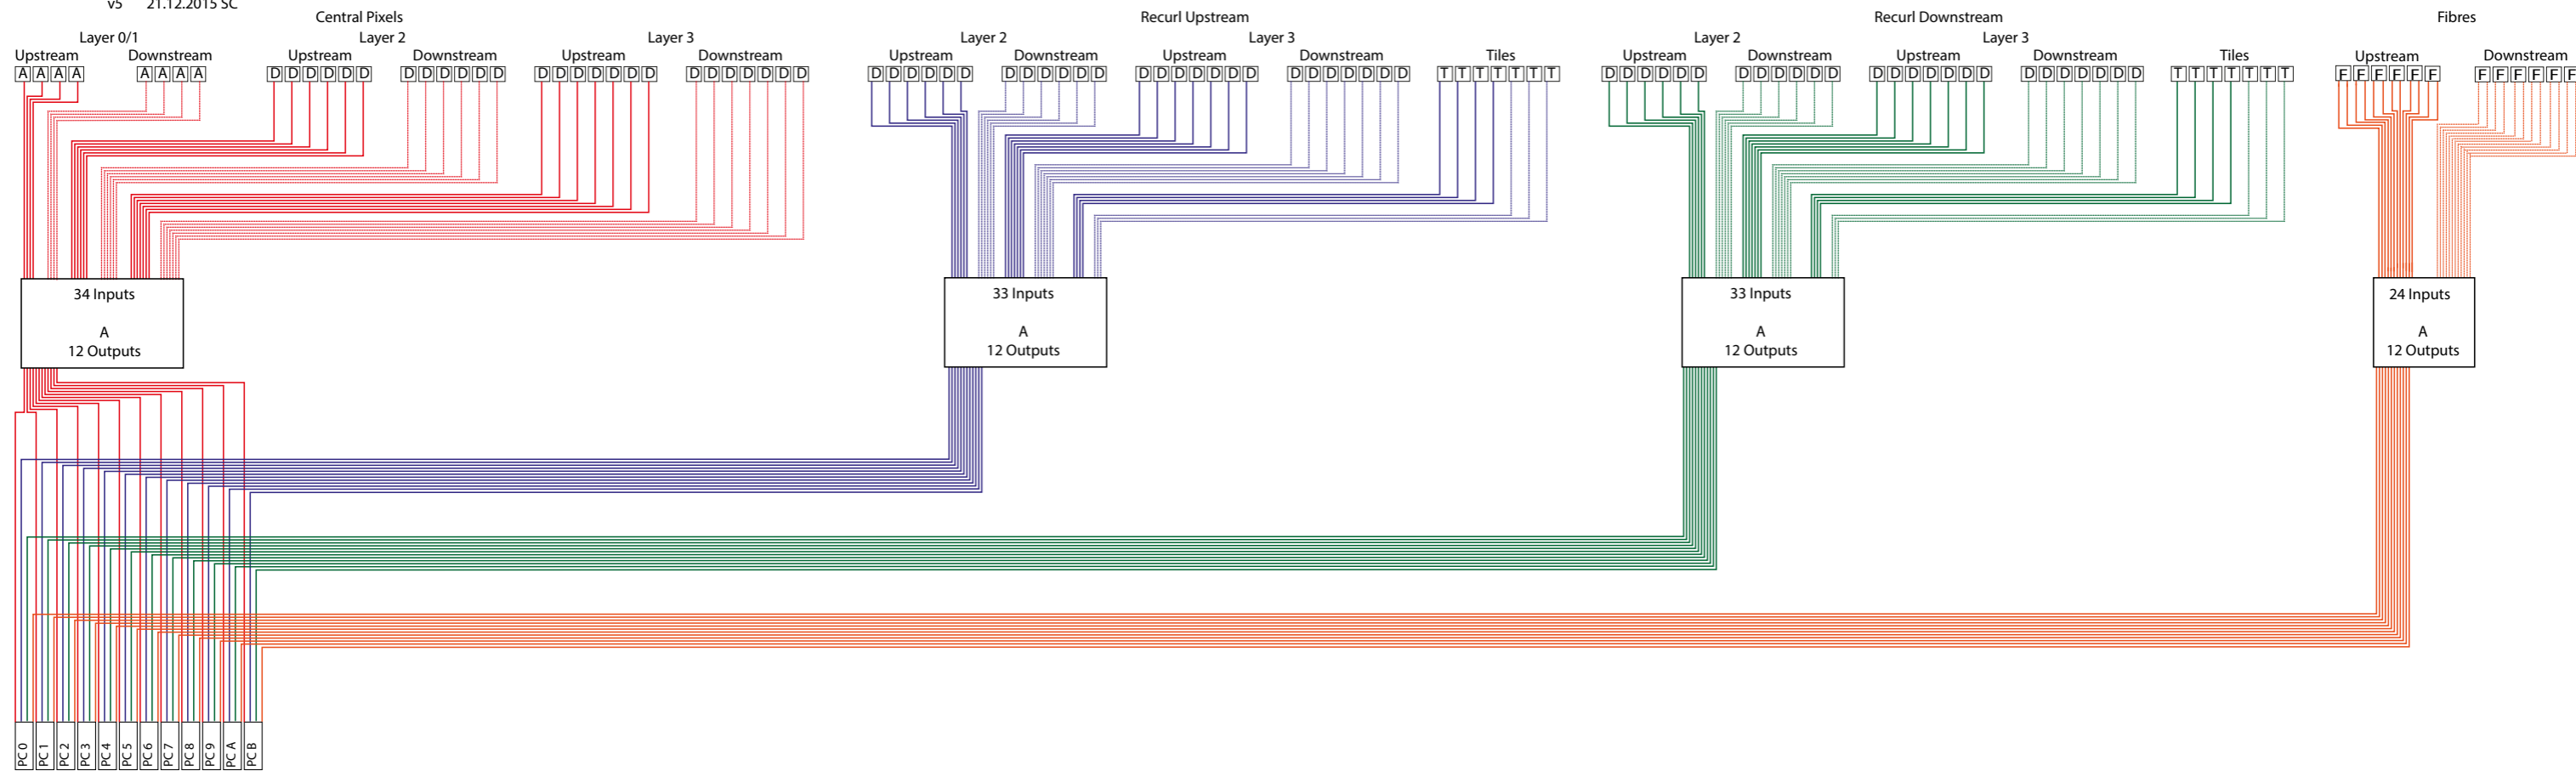
- v1 2.10.2013 NB
- v2 21.10.2013 NB
- v3 27.4.2015 NB
- v4 8.10.2015 NB





Mu3e Phase Ib Optical Readout Links

- v1 2.10.2013 NB
- v2 21.10.2013 NB
- v3 28.4.2015 NB
- v4 8.10.2015 NB
- v5 21.12.2015 SC






Mu3e Phase I Pixel Readout Scheme


- v1 2.10.2013 NB
- v2 21.10.2013 NB
- v3 6.11.2014 NB
- v4 26.11.2014 NB
- v5 31.3.2015 NB

All addresses in hexadecimal notation

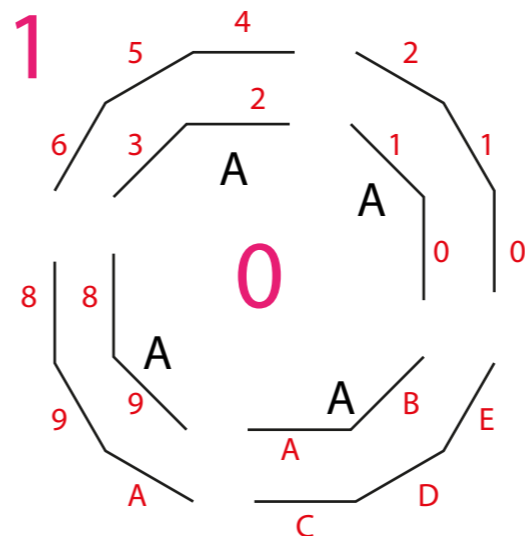


Vertex Layers

Upstream side





Downstream side



FPGA Types:


- A: 5x3 Chips with 1 Link: 15 Links 8 FPGAs
- D: 4x9 Chips with 1 Link: 36 Links 26 FPGAs

Address Scheme:




Outer Layers


Upstream side



Downstream side



Upstream



Downstream



Mu3e Phase I Common Readout Firmware

v1 1.5.2015 NB

