



BERGISCHE
UNIVERSITÄT
WUPPERTAL



JÜLICH
Forschungszentrum

Design and Characterisation of an MRI Compatible Human Brain PET Insert by Means of Simulation and Experimental Studies

Dissertation

zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

der Fakultät für Mathematik und Naturwissenschaften
der Bergischen Universität Wuppertal

vorgelegt von

Mirjam Lenz
aus Wuppertal

2020

Gutachter

Prof. Dr. Uwe Pietrzyk

Institut für Neurowissenschaften und Medizin (INM-4)
Forschungszentrum Jülich GmbH
Fakultät für Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Prof. Dr. Christian Zeitnitz

Fakultät für Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Prof. Dr. Christian Morel

Centre de Physique des Particules de Marseille (CPPM)
Aix-Marseille Université
Institut National de Physique Nucléaire
et de Physique des Particules (IN2P3)
Centre National de la Recherche Scientifique (CNRS)

Mitglieder der Prüfungskommission

Dr. Markus Axer

Institut für Neurowissenschaften und Medizin (INM-1)
Forschungszentrum Jülich GmbH

Prof. Dr. Karl-Heinz Kampert

Fakultät für Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Prof. Dr. Uwe Pietrzyk

Institut für Neurowissenschaften und Medizin (INM-4)
Forschungszentrum Jülich GmbH
Fakultät für Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Prof. Dr. Christian Zeitnitz

Fakultät für Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Tag der mündlichen Prüfung

14. Juli 2020

*The true delight is in the finding out
rather than in the knowing.*

Isaac Asimov

The PhD thesis can be quoted as follows:

urn:nbn:de:hbz:468-20210428-104546-2

[<http://nbn-resolving.de/urn/resolver.pl?urn=urn%3Anbn%3Ade%3A468-20210428-104546-2>]

DOI: 10.25926/d4s5-bc37

[<https://doi.org/10.25926/d4s5-bc37>]

Acknowledgements

First of all, I would like to thank Uwe Pietrzyk, who suggested and supervised this interesting PhD topic. Despite his many activities at the research centre, he always took the time for discussions and provided lots of helpful ideas. I highly appreciate his encouragement and continuous support, which remained absolutely unchanged after his retirement in 2018.

The supervision was taken over by Christoph Lerche, who had co-supervised me from the beginning. I am extremely grateful that he shared his broad expertise on silicon photomultipliers, positioning algorithms and PET instrumentation in general. His door was always open, and he provided constructive advice during our frequent discussions on the most recent findings.

I would also like to thank N. Jon Shah for giving me the opportunity to conduct my PhD project at the INM-4, and especially for supporting my research activities abroad. I gratefully acknowledge that my research was supported by the Initiative and Networking Fund of the Helmholtz Association (grant HVF-0051), and that my research stay in France was supported by the Erasmus Plus programme at the university of Wuppertal.

A special thanks goes to former and current members of the PET group, for countless vivid discussions and the pleasant working atmosphere. Particularly supportive and helpful to me were my office colleagues Liliana Lourenco Caldeira and Hancong Xu. I am thankful for the critical reflections on my findings, our memorable coffee breaks, and that I could always count on their support- the PhD time wouldn't have been the same without them.

I'd also like to extend my gratitude to Matthias Streun, who made his test stand for digital silicon photomultipliers available to me, guided me during the first measurements and was always open for discussions. Holger Nöldgen and Roger Heil provided technical support in the lab, which I thankfully acknowledge.

My sincere thanks go to Rajalekshmi Deepu and Wouter Klijn for supporting my GATE activities on the supercomputer JURECA. I also had the pleasure to work with Stepan Nassyr, who patiently answered my questions about computing. Without their advice and the computation time granted by the John von Neumann Institute for Computing, my simulation studies would not have been feasible. Many thanks also go to Jürgen Scheins, who reconstructed my simulated data sets with PRESTO.

To Francis Loignon-Houle I am indebted, because he made me aware of reflectivity quenching of 3M ESR vikuiti, and therefore provided an important hint for matching simulated and measured data. I would also like to thank Christian Morel and Mathieu Dupont for the warm welcome at Centre de Physique des Particules de Marseille (CPPM). Mathieu supervised my three months stay at CPPM and supported me during the first steps of the GATE digitiser implementation. I am thankful for the insights gained in this time and the exciting time abroad.

A big *thank you* goes to Cláudia Régio Brambilla, Hancong Xu, Matthias Streun, Benjamin Lenz and especially Christoph Lerche, who took the time to proofread my thesis and who gave me valuable feedback. All flaws in this thesis are my own responsibility.

I am extremely grateful to my family, especially my parents and my brothers, for their unwavering support and encouragement in tough times. Without the constant support of my fiancé Michael, the completion of my dissertation would not have been possible. I thank him for being with me in all ups and downs, and for his understanding whenever I needed "*just five more minutes...*" at the desk. Thank you for carrying me to the finish line and reminding me of the important things in life.

Abstract

Positron emission tomography (PET) is a widely used in-vivo imaging technique to visualise metabolism, allowing for a broad spectrum of applications in oncology, cardiology and neuroscience. At present, an MRI compatible human brain PET scanner for applications in neuroscience is being constructed in the scope of a Helmholtz Validation Fund project.

In this thesis, a detector for this novel PET device was designed. The detector concept combined three scintillator layers with a lightguide and digital silicon photomultipliers (dSiPMs). Monte Carlo simulations were used to optimise the dimensions of the scintillator arrays, so that the new scanner design yielded the maximum possible sensitivity. The benefit from the additional depth information, which can be acquired with three scintillator layers, was evaluated and proven to be higher compared to a less expensive two layer geometry. Since a more homogeneous spatial resolution was achieved in the whole field of view, this finding had a high relevance for the envisaged neuroscientific applications.

In order to accurately acquire the depth information, new strategies for decoding the flood map during the calibration of a detector module were developed. This required realistic simulation data with ground truth information, so that the simulation toolkit GATE was extended to model the electronic readout of the dSiPMs. To overcome extended simulation times and to provide simulations on a statistically sound basis, the GATE studies were executed on the supercomputer JURECA. The simulated data were matched to measured data from test detectors. This allowed the determination of an optimum thickness of a lightguide between the scintillators and the dSiPMs. Moreover, the number of correctly identified scintillation events was evaluated by means of different event positioning approaches and different clustering methods during the calibration step. The highest amount of correctly identified events in a single detector block was achieved with model-based clustering and Maximum Likelihood positioning (61.5%). By simulating the whole propagation and detection of scintillation photons including ground truth information, this study provides the opportunity to improve the positioning approaches and to enhance this number in future.

The gained insights were further applied to select a surface finish of the scintillators. Measurements with crystal samples of the final detector dimensions showed that rough lateral crystal surfaces yielded the best signal separation in the calibration flood map.

The experimental and simulation studies presented in this thesis had a major influence on the final detector design of the novel brain PET. The detailed simulations including the propagation and detection of scintillation photons were in good agreement with measured data, and could be a promising approach for future detector design studies.

Contents

1	Introduction	1
2	Background	5
2.1	Introduction to PET	5
2.1.1	Historical Overview	5
2.1.2	Positron Emission	6
2.1.3	Detection of Annihilation Photons	7
2.1.4	Data Corrections	8
2.1.5	The Parallax Error	10
2.1.6	Other Factors Limiting Spatial Resolution	10
2.2	PET Detectors	11
2.2.1	General Detector Concept	11
2.2.2	Choosing Components for a PET Design	16
2.2.3	Crystal Geometries	17
2.2.4	Photodetectors	21
2.2.5	Philips Digital Photon Counting	23
2.2.6	Positioning Algorithms	28
2.3	Hybrid Imaging	30
2.3.1	MRI compatibility of the PET insert	31
2.4	Monte Carlo Simulations	32
2.4.1	GATE: Geant4 Application for Tomographic Emission	33
3	Motivation and Thesis Outline	35
3.1	Requirements for a BrainPET Scanner	35
3.2	Detector Concept in this Work	39
3.3	Objective of this Thesis	42
4	Optimisation of Detector Block Configuration	43
4.1	Simulation of Sensitivity in GATE	44
4.2	“Is it really worth it?”- A cost-benefit assessment	48
4.3	Impact of Capability to Detect DOI Information	50
5	Building Test Detectors	57
5.1	Assembly of Prototypes	57
5.2	Acquisition of Flood Maps	59

6	Simulation of Scintillation Photons in a Detector Block	67
6.1	Simulation Setting	67
6.2	Acquisition of Floodmaps	71
6.3	Simulations Including Reflectivity Quenching	75
6.4	GATE Digitiser Dedicated to DSIPMs	78
6.4.1	Structure of the Digitiser	78
6.4.2	Validation of the Digitiser	81
6.4.3	Results: Application to Previous Simulations	87
6.5	Additional Information from the Simulation	91
7	Study on Lightguide Thickness	95
7.1	Preliminary Study	95
7.2	Simulations with Different Lightguides	97
7.3	Comparison to Measurements with Different Lightguides	101
8	Evaluation of Reference Response for Positioning Algorithms	107
8.1	Maximum Likelihood Event Positioning: Workflow	107
8.2	Simulations: Using the Ground Truth Information	109
8.2.1	Simulation Setup	109
8.2.2	Distinguishing Signals from Different Crystal Layers	110
8.2.3	Application to Overlapping Spots	115
8.3	Clustering Methods	119
8.4	Influence of Response Functions on Positioning Results	123
8.5	Measurements with a Detector Block of Final Size	128
8.5.1	Flood Maps with Ground and Polished Surface Finish	128
8.5.2	Energy Resolution of Both Crystal Arrays	130
8.5.3	Spot Separability in the Final Detector Block	133
9	Conclusions and Outlook	141
A	Additional Material	147
A.1	Schematic Drawing of Scanner Cassettes	147
A.2	DPC Data Processing	148
A.3	Validation of the GATE Digitiser	149
A.4	Simulations Concerning the Reflectivity of 3M ESR Vikuiti	156
	List of Figures	157
	List of Tables	167
	Bibliography	169

List of Abbreviations

APD	Avalanche photodiode
BGO	Bismuth germanate $\text{Bi}_4\text{Ge}_3\text{O}_{12}$
CoG	Centre of gravity
CT	Computed tomography
DCR	Dark count rate
DOI	Depth of interaction
DPC	Digital photon counting
ESR	Enhanced specular reflector
FDG	Fluorodeoxyglucose
FOV	Field of view
FPGA	Field programmable gate array
FWHM	Full width at half maximum
G-APD	Geiger-mode avalanche photodiode
GATE	Geant4 application for tomographic emission
HRRT	High-Resolution research tomograph
LOR	Line of response
LSO	Lutetium Oxyorthosilicate Lu_2SiO_5
LYSO	Lutetium-Yttrium Oxyorthosilicate $\text{Lu}_{1.8}\text{Y}_{0.2}\text{SiO}_5$
MC	Monte Carlo
MLEM	Maximum likelihood expectation maximisation
mM	Millimolar
MPPC	Multi-pixel photon counters
MRI	Magnetic resonance imaging
MRS	Magnetic resonance spectroscopy
μM	Micromolar
NEMA	National Electrical Manufacturers Association
nM	Nanomolar
OSEM	Ordered subset expectation maximisation
PDE	Photon detection efficiency
PDF	Probability density function
PET	Positron emission tomography
pM	Picomolar
PMT	Photomultiplier tube
PRESTO	PET reconstruction software toolkit
PSAPD	Position sensitive avalanche photodiode
PVR	Peak valley ratio

QE	Quantum efficiency
RF	Radio frequency
RTL	Row-trigger-lines
SiPM	Silicon photomultiplier
SNR	Signal to noise ratio
SPAD	Single photon avalanche devices
SPECT	Single-photon emission computed tomography
TDC	Time-to-digital converter
UHF	Ultra-high field

Chapter 1

Introduction

Positron Emission Tomography (PET) is a widely used *in vivo* imaging modality, which is capable of imaging metabolic processes. By labelling a metabolically active substance with a radionuclide, the distribution of the labelled substance can be reconstructed from the detected radiation, which originates from the radioactive decay [1]. Since there are various radiopharmaceuticals available, which label e.g. glucose, amino acids or specific receptor ligands [2], PET allows for a broad spectrum of applications in the fields of oncology, cardiology and neurology.

In contrast to other imaging modalities such as Magnetic Resonance Imaging (MRI) or Computed Tomography (CT), which deliver detailed morphological information on the subject, PET investigates physiological processes inside the human body with a high diagnostic accuracy. This is feasible due to the high sensitivity (nanomolar (nM) to picomolar (pM)), which is a magnitude lower for MRI (μM to nM) and two magnitudes lower for CT (mM to μM) [3]. The acquired metabolic information can be helpful for early diagnosis, treatment planning and determination of a disease state, even if there is no or not yet a visible structural alteration [4]. However, the low spatial resolution of PET images makes a combination with MRI or CT desirable to gain high-resolution structural information and functional information simultaneously [2]. Although the combination of both modalities in hybrid devices like PET-CT or PET-MRI is technically challenging, it also provides the opportunity to develop new imaging protocols for functional or dynamic studies, and to reduce the total image acquisition time [5].

Whole-body PET scanners usually provide a spatial resolution between 3.5 and 5 mm [6], [7], a resolution which is limited by different factors. The dimensions of the scintillation crystals limit the achievable resolution to half of the crystal pitch, but also the range of the emitted positrons and the reconstruction technique contribute to the degradation of spatial resolution [8]. Moreover, the homogeneity of the spatial resolution inside the field of view is limited. Depending on the incidence angle of the annihilation photons, the photons' penetration depth plays a major role in the correct identification of a scintillation event. This effect is referred to as the parallax error or radial elongation, since it increases with radial distance from the scanner axis, and leads to a degradation of spatial resolution at radial offsets [9].

Considering the depth of interaction (DOI) during the image reconstruction can reduce this error and therefore make the spatial resolution more homogeneous. A drawback of most PET scanners is the low geometric sensitivity, which is usually $<5\%$ [10]. However, a high number of detected decay events is crucial for the signal to noise ratio (SNR) in a PET measurement, which is why one aims at maximising the detector sensitivity. Since this aspect is mainly influenced by the thickness of the scintillator crystals and the geometric coverage of the whole scanner, there exist organ-specific devices for dedicated imaging of the brain, breast, heart and prostate [11]. These vary in the shape and configuration of the detector modules, so that the respective organ can be optimally covered.

In this thesis, the focus will be on detector design for a dedicated brain PET scanner for use in neuroscience. Typical applications are the investigation of neuroreceptors' affinity and availability, the progression of Parkinson's or Alzheimer's disease, or identification and staging of brain tumours [2], [4]. To detect signals within the delicate structures of the brain, a higher spatial resolution than those of whole-body scanners is required [12]. Moreover, the spatial resolution is desired to be homogeneous in the field of view, so that signals from the outer cortex of the human brain can also be resolved. A combination with ultra high field (UHF) MRI would benefit from the highly resolved anatomical information in the sub-millimetre range, and MR spectroscopy that could be acquired simultaneously [13].

Two dedicated brain PET scanners have been developed so far, the High Resolution Research Tomograph (HRRT) [14] and a brain PET insert for 3T MRI. The latter has been installed at four sites, at the research centre Jülich, in Tübingen, Boston and Atlanta [15]. However, both, the HRRT and the Siemens BrainPET, were developed as research tomographs, so that there is no commercially available brain PET. This gap is planned to be bridged by a new MRI compatible brain PET insert, that is currently being built within the scope of a Helmholtz Validation Fund at the research centre Jülich. This scanner is envisaged to provide a high sensitivity ($>10\%$), a high spatial resolution ($<2\text{ mm}$), and to be compatible with UHF MRI at 7T.

The design and characterisation of a new detector module for this brain PET scanner will be the aim of this thesis. Since several new technologies will be combined to meet the demands for a high-resolution scanner, the characterisation studies will be conducted by means of test detectors and Monte Carlo simulations using the toolkit GATE [16]. These studies include finding optimum dimensions of the detector components, such as the scintillator arrays and a lightguide. Moreover, the new detectors are envisaged to provide depth of interaction information, so that strategies for processing of the raw data must be developed. The potential improvement compared to the existing Siemens BrainPET will also be investigated.

This thesis is structured as follows: Chapter 2 introduces the basic working principles of PET, while the main focus is set on detector technology. The single components of a detector like scintillator crystals, lightguides and digital silicon photomultipliers and their electronic readout are described in detail. Moreover, Monte Carlo simulations of PET scanners using GATE are outlined.

Currently used PET scanners and their detectors are reviewed in chapter 3. The detector concept for the new brain PET scanner is presented in contrast, and the specific goals of this thesis are explained in more detail.

In chapter 4, the configuration of the single detector modules is optimised regarding sensitivity of the whole scanner. For this study, different module sizes are simulated in GATE. The results are compared to the currently used brain PET by studying the impact of the new detector concept on reconstructed images.

First measurements with test detectors are presented in chapter 5. Closely related are the simulation studies in chapter 6, which tackle the creation, propagation and detection of scintillation photons in a single detector module, in order to gain a deeper understanding of the measured data. Further simulation studies and corresponding measurements are used to find an optimum lightguide thickness for the detector modules. These are described in chapter 7.

Chapter 8 deals with the identification of scintillation events inside a detector module. The insights gained in the previous simulation studies are applied to test different segmentation approaches during the calibration of a single detector block. Moreover, two methods are compared, which determine the crystal where the scintillation had taken place. These studies are based on simulated data, which provide a ground truth information. The chapter closes with measurements of scintillator arrays, which have the dimensions determined in chapter 4. These are used to decide on the surface treatment of the crystals.

Chapter 9 concludes the presented work and gives an outlook to future work.

Chapter 2

Background

2.1 Introduction to PET

2.1.1 Historical Overview

Since its first applications in the 1970s [1], Positron Emission Tomography (PET) has evolved to a widely used *in vivo* technique for functional imaging. A radiopharmaceutical, which contains a metabolic active substance labelled with a radionuclide, is administered to the subject (most commonly by intravenous injection). Since the substance takes part in metabolic processes, along with the radioactive component, the measured distribution of decay products allows us to reconstruct the distribution of the labelled substance in the living organism. For PET, positron emitters like ^{18}F or ^{15}O are used: the emitted positrons annihilate with electrons in the surrounding tissue, giving rise to two nearly collinear annihilation photons (emission angle of $(180 \pm 0.5)^\circ$ [17]). Due to their energy of 511 keV, the photons can still be detected outside the body. Dedicated reconstruction algorithms use this information to calculate the distribution of the labelled substance.

Before a first tomograph for application on humans could be built, several milestones had to be achieved, starting with George Hevesy's idea of "radioactive indicators" in 1923 [18]. He stated that an atom of a compound could be substituted by a radioactive isotope of the same element, while the compound would not alter its chemical characteristics. The measurement of decay products of the radioactive component (i.e. annihilation photons) can yield information on its distribution, even in a living organism. Hevesy was rewarded with the Nobel Prize for Chemistry in 1943 for the development of the tracer principle.

Further milestones in the development of PET devices are the postulation of the positron (Paul Dirac, 1929) and insights on the potential use of the isotopes ^{11}C , ^{13}N , ^{15}O and ^{18}F as radioligands for biomarkers, gained in the late 1920s to late 1940s [1], [19]. Artificial production of these isotopes was made possible by the invention of cyclotrons (Edward Lawrence, 1930s), which were used in clinical environments from the mid-1950s on [1]. Scintillation counters came up around 1948, when Robert Hofstadter discovered the energy conversion of high-energy gamma rays into scintillation photons in a thallium doped sodium iodide crystal (NaI:Tl) [20]. This discovery

enabled the detection of annihilation radiation, so that first applications on the human brain followed quickly. However, it was only in the 1970s, when dedicated PET scanners were first presented [1].

Nowadays, PET provides a broad spectrum of clinical and research applications: it is used for diagnosis and therapy monitoring in oncology, cardiology and neurology. In oncology, a PET scan can be helpful to locate and stage a tumour, and to plan or evaluate a radiotherapy [21]. In neurology, the affinity and availability of neuroreceptors is often visualised, e.g. to localise epileptic foci or to monitor the progression of Parkinson's disease [4]. Moreover, PET is used to understand the biology of diseases and develop new drugs or review existing ones. In early stages of drug development, small-animal PET is used in longitudinal studies, and in studies with genetically altered animals (e.g. to obtain deeper knowledge on drug mechanisms) [22]. Some typical clinical applications and the respectively used tracers are listed in table 2.1.

TABLE 2.1: Overview of common PET tracers and their clinical applications. Republished and adapted with permission of Eureka Science (FZC), from "Clinical Applications of Positron Emission Tomography (PET) Imaging in Medicine: Oncology, Brain Diseases and Cardiology", S. Kitson et al, *Current Radiopharmaceuticals*, vol. 2, no. 4, 2009 [4]; permission conveyed through Copyright Clearance Center, Inc.

PET Radiotracer	Isotope Half-Life	Physiological Process or Function	Clinical Application
$[^{11}\text{C}]$ -Flumazenil	20.3 min	Benzodiazepine receptor antagonist	Epilepsy
$[^{11}\text{C}]$ -Raclopride	20.3 min	D ₂ receptor antagonist	Movement disorders
$[^{15}\text{O}]$ -Carbon dioxide	2.1 min	Blood perfusion	Brain activation studies
$[^{18}\text{F}]$ -Fluorodeoxy-glucose (FDG)	109.8 min	Glucose metabolism	Oncology, cardiology, neuropsychiatry

The following paragraphs will describe the basic principles of PET, the workflow and the components of PET detectors. General information on PET is taken from reference [23], if not denoted otherwise.

2.1.2 Positron Emission

As the name Positron Emission Tomography already indicates, the tomographic scan is based on positron emission. Positron emission, or more general β^+ decay, can be described by the following formula:



Under emission of a positron e^+ and an electron neutrino ν_e , a nucleus X with mass number A and atomic number Z is converted to one with $Z - 1$. This conversion process can be understood as the decay of a proton in the nucleus to a neutron. The emitted positrons can directly annihilate with their antiparticles, electrons, and thus give rise to two annihilation photons ("in-flight annihilation")



or positrons and electrons form Positronium (lifetime $\sim 10^{-10}$ s) prior to annihilation. The annihilation photons propagate in opposing directions. Equation 2.2 assumes that both the electron and the positron have negligibly small kinetic energies compared to their rest energy of 511 keV. This is usually the case, since the positron can undergo inelastic or elastic scattering before annihilation, and thus ionise or excite atoms by transferring kinetic energy or lose kinetic energy due to bremsstrahlung. In compliance with the law of energy and momentum conservation, the remaining kinetic energies of positron and electron are transferred to the annihilation photons. The two photons can thus have an acollinearity of 0.5° . The full width at half maximum (FWHM) distance which a positron can propagate before annihilation is in the range of less than a millimetre for common PET tracers (0.102 mm FWHM for ^{18}F up to 0.501 mm for ^{15}O in Water) [23].

In PET, positron emission takes place in the radioactive component of the tracer, which is produced in a cyclotron and then administered intravenously to the patient in the form of a radiopharmaceutical. Annihilation occurs in the surrounding tissue, so that two annihilation photons emerge with a relative angle of $(180 \pm 0.5)^\circ$ [17].

2.1.3 Detection of Annihilation Photons

The detection of annihilation photons is an essential step of PET. The general idea is to detect annihilation photon pairs in a detector ring, so that the annihilation position can be estimated by means of a reconstruction algorithm. The distribution of annihilations nearly corresponds to the distribution of the tracer, and thus the labelled substance (such as the glucose analogue in FDG). In this way, PET allows for investigation of metabolic processes in the human body.

For this reason, photon detection plays an important role in PET. Detectors are required to stop the photons, and to resolve the interaction positions and times of the photons inside the detector. A dedicated coincidence sorting unit selects photons which arrive within a certain time window (typically a few nanoseconds), to find coincident events, i.e. photons which are assumed to originate from the same annihilation event. Moreover, an energy threshold can be applied to filter out scattered events (e.g. from 480 to 630 keV). This is important for the determination of the so-called Line of Response (LOR), which is the connecting line between the interaction positions of two annihilation photons in the respective detectors.

In general, four types of coincidence events are possible, as shown in figure 2.1:

1. True coincidence: the photons hit the detector directly, without any interactions, i.e. change in energy or propagation direction.
2. Scattered coincidence: at least one of the photons undergoes Compton scattering and changes direction and energy.
3. Random coincidence: two photons reach the detector within the detection time window, but originate from different annihilation events.
4. Multiple coincidence: more than two photons reach the detector within the detection time window. Since the correct LOR cannot be determined, these events are discarded.

Except for true coincidences, all the others lead to wrong LORs and should thus be filtered out, despite the fact that this filtering leads to a decrease in sensitivity.

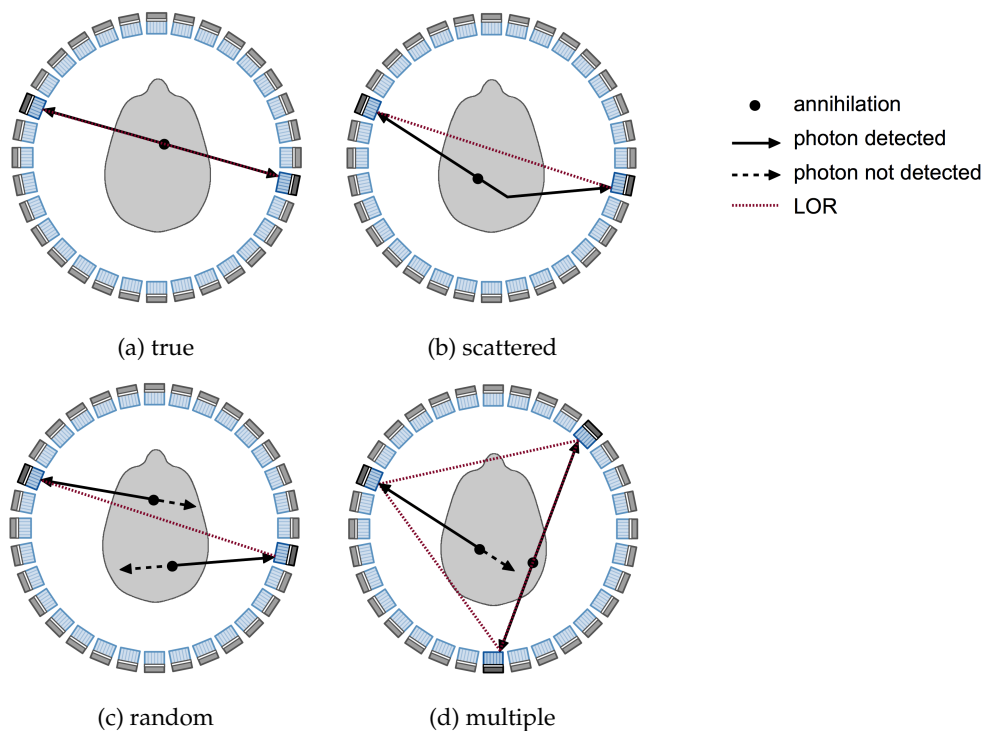


FIGURE 2.1: The coincident events comprise true (a), scattered (b), random (c) and multiple (d) coincidences.

2.1.4 Data Corrections

In order to calculate the point of annihilation, the energies, arrival times and detection positions of the two annihilation photons are fed to the reconstruction algorithm. Of course, the reconstruction is not trivial at all for several reasons: first, photons can interact with matter before leaving the patient, second, the point of annihilation is not known, only the line of response which contains this point, and third, the properties of the detector blocks have to be considered.

Therefore, several corrections have to be applied to the raw data; they will only be described briefly, since they are not the focus of this thesis.

Since the photons can interact with the different tissue types in the head (soft matter, bone, blood vessels, air cavities ...) and MRI hardware (e.g. head coil), an attenuation map is created to account for this effect. Different approaches are possible, such as using a transmission scan (source is rotating around the patient's head), calculation of the attenuation map from a CT scan, or in case of PET-MRI the MRI-derived calculation. An overview can be found in the work of *Rota Kops and Herzog* [24].

The attenuation map can be further used in combination with an emission map and a model of the scanner geometry, in order to estimate the distribution of scattered events in the object analytically [25]. Other methods exist for the correction of scattered events and are described in the book *Physics in Nuclear Medicine* by Cherry et al. [23].

Another important factor for image degradation is the presence of random coincidences. Two methods exist for random correction, the 'delayed window' and 'scaled singles' method.

To estimate the rate of random events, a second time window can be set with a delay that is larger than the window width. Photons that fall within this delayed time window have different arrival times that are significantly larger than the time window, so that they are definitely not belonging to the same annihilation event. The event rate acquired in this mode can be used to assess the random rate, and to subtract random events from the coincidence data.

A second method to correct for random coincidences is called 'scaled singles' method: the singles rates of two detectors, S_i and S_j , and the coincidence time window τ are used to estimate the random rate of the LOR for detectors i and j : $R_{ij} = 2\tau S_i S_j$ [26].

A normalisation scan is necessary to account for different detector block response functions. Depending on the orientation (i.e. the angle between the two detector faces) and position of a LOR inside the scanner geometry, different LOR sensitivities can be observed. To correct the different sensitivities, all detector blocks are irradiated equally, e.g. with a rotating rod source, to calculate the ratio of measured and expected events. These factors also indirectly account for different photodetector gains and crystal light yields [27].

When the raw data has undergone all corrections, it can be evaluated by a reconstruction algorithm. While the filtered back projection has been the gold standard for a long time, nowadays iterative methods such as maximum likelihood expectation maximisation (MLEM) and the accelerated ordered subset EM (OSEM) are usually used for reconstruction. However, a description of these algorithms would be beyond the scope of this thesis; an introduction can be found in the work of *Defrise et al.* [28].




FIGURE 2.2: The parallax error in one transaxial direction in different crystal blocks. While it is quasi not existing in the centre field of view (FOV) (limited to crystal size), the parallax error becomes larger towards the periphery of the FOV. The LORs in the areas marked in grey are assigned a fixed LOR (black lines), which can differ a lot from the true LORs (red lines).

2.1.5 The Parallax Error

A correct determination of the LORs is very important, especially when annihilation photons hit detector modules under oblique angles. In these cases, the so-called parallax error can occur. The photons are not necessarily stopped in the first crystal they encounter, but can also pass through them into another crystal of this detector block. As indicated in figure 2.2, a coordinate in the front part of the hit crystal is usually used for the LOR calculation in many reconstruction algorithms, so that the depth of interaction leads to mispositioning errors (also called the parallax error). The effect of the parallax error is especially visible when the source has a radial offset from the scanner axis, because the amount of photons entering at oblique angles becomes larger with the radial offset. The parallax error therefore lowers the spatial resolution with increasing radial offset, which is often referred to as radial elongation [29]. If the correct interaction position including the depth of interaction (DOI) inside the crystal would be known, the LORs could be calculated more precisely, which would reduce the parallax error.

2.1.6 Other Factors Limiting Spatial Resolution

Apart from the parallax error, there exist more physical effects which can limit the spatial resolution of a PET detector. A detailed description was provided by *Moses* [8], so that only a brief overview will be given here.

When a positron is emitted from the tracer, it does not necessarily form positronium or annihilate with an electron from the surrounding tissue immediately. By propagating through the tissue, the spatial resolution can be degraded by 0.54 mm FWHM for ^{18}F and up to 6.14 mm FWHM for ^{82}Rb , depending on the positron's

kinetic energy after emission. Moreover, the annihilation photons are not always emitted under a fixed angle of 180° , but show a mean acollinearity of 0.5° (depending on the positron's and electron's kinetic energy prior to annihilation). According to Moses [8], the acollinearity can degrade the spatial resolution by 1.3 mm in brain imaging.

The crystal size d of a pixelated PET detector also limits the spatial resolution: the detector response function for two opposing detector blocks has the shape of a triangle with a FWHM of half the crystal pitch $d/2$. However, crystal size cannot be minimised arbitrarily, since the technical implementation would be challenging and not economical (more details in the following section).

2.2 PET Detectors

Generally speaking, PET detectors have two main tasks: stopping the annihilation photons of 511 keV, and providing information about the photons' interactions, i.e. time, coordinates or crystal ID within the detector block, and amount of deposited energy. Such a typical detector block consists of a scintillator material, which emits photons in the visible spectrum when irradiated with high-energy particles, and a photodetector, which absorbs and quantifies the scintillation light.

Actually, this principle of detection has already been used in the 1950s, when Curran and Baker counted scintillation light in a zinc sulfide (ZnS) screen with a photomultiplier tube. The idea of using a scintillator for particle detection is even older: Crookes application of watching a ZnS screen with bare eye to detect α -particles is considered to be the first of its kind [30].

2.2.1 General Detector Concept

Nowadays, the same basic principle of detection is used. *Lecoq et al.* described the scintillation process in detail in reference [20], the following description will be based on this reference.

Scintillators

When radiation interacts with a scintillator material, effects of basic electromagnetic interactions like Compton scattering or photoelectric effect can take place. Also the pair production of electron and positron is possible, although energies higher than 1.022 MeV are required, so that Compton scattering and photoelectric effect are the dominant processes for a 511 keV photon. The Compton- and photoelectrons undergo Compton scattering as well, and thereby lose energy in the keV range. Inorganic scintillators like Lutetium Oxyorthosilicate (Lu_2SiO_5 , LSO) have a crystal lattice structure, where electrons are occupying selected energy bands. Electrons from the valence band can get excited by the secondary electrons, and occupy allowed states in the conduction band or excitator band (closely below conduction band),

depending on the excitation energy. At the same time, this process leaves a hole in the valence band, which is usually filled at room temperature. Depending on the electron's energy E , either a loosely bound electron-hole pair (also called exciton) is created ($E \lesssim E_g$), or the electron propagates independently from the associated hole through the conduction band.

Both electrons and holes quickly lose energy by coupling to phonons (lattice vibrations). Different processes can take place afterwards: electrons and holes can get captured by traps in the band gap, generating phonons with the excess energy ($E_c - E_{tr}$ for electrons, $E_{tr} - E_v$ for holes). Excitons can recombine at luminescence centres or lattice defects. These luminescence centres are required for scintillator materials: when a freely moving electron recombines to a lower energy state under emission of a photon, and the photon's energy corresponds to the band gap size (in the range of 5 eV to 10 eV in insulators), the probability for bulk absorption within the crystal is very high. Doping the crystal with impurities provides additional energy levels within the forbidden energy band, so that the emitted light has a larger wavelength: the scintillator becomes transparent for its own scintillation light. Since the amount of scintillation photons produced is proportional to the deposited energy in the crystal, measurements of these scintillation photons can be used to determine the energy of the high-energy photon that just interacted within the crystal. The possible processes are depicted in figure 2.3.

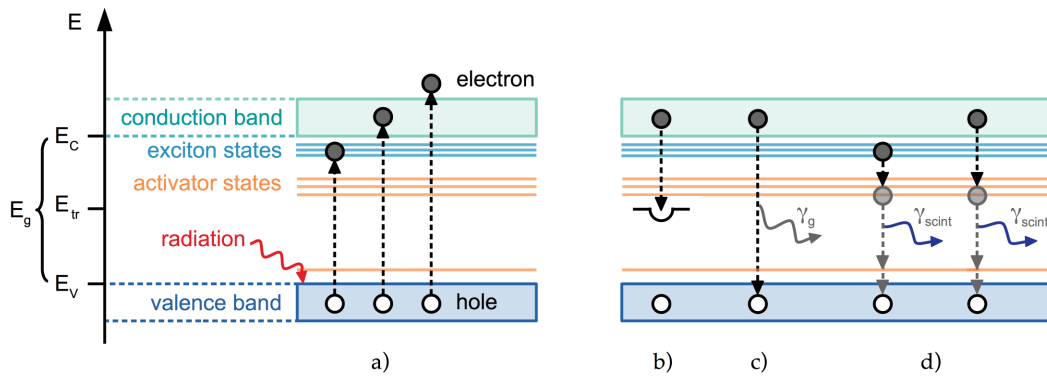


FIGURE 2.3: The scintillation process described in the band model. a) Possible excitations of an electron, leaving a hole in the valence band. b) Electron caught in a trap, the remaining energy is transferred to the crystal lattice (phonon). c) De-excitation of an electron under emission of light γ_g with E_g -self absorption is probable. d) Recombination via activator states, emission of scintillation light γ_{scint} .

When it comes to the choice of a suitable scintillator for an application, transparency and light yield of a crystal are important parameters. For PET, especially high density ρ and high effective atomic number Z_{eff} are important to increase the probability for stopping the high-energy photons by photoelectric effect ($\propto Z_{eff}^3/E^3$) [23]. For a high temporal resolution, a short decay time of the excited states is required. Of course, the economic aspect plays a role in detector design, so that low production costs and uncomplicated crystal growth are desired.

Commonly used scintillator materials for PET scanners are bismuth germanate ($\text{Bi}_4\text{Ge}_3\text{O}_{12}$, BGO), Ce-doped lutetium oxyorthosilicate ($\text{Lu}_2\text{SiO}_5\text{:Ce}$, LSO:Ce) and (similarly) Ce-doped lutetium-yttrium oxyorthosilicate ($\text{Lu}_{1.8}\text{Y}_{0.2}\text{SiO}_5\text{:Ce}$, LYSO:Ce), just to name a few. Their basic properties are displayed in table 2.2, together with those for NaI, which is often used as reference due to its early discovery and extensive use in single-photon emission computed tomography (SPECT) [23].

TABLE 2.2: Comparison of properties of common scintillators used for nuclear imaging. This table was published in S. R. Cherry, J. A. Sorenson, and M. E. Phelps, "Radiation Detectors" in *Physics in Nuclear Medicine*, Fourth Edition, Oxford: Elsevier LTD, 2012, p. 101, Copyright Elsevier 2008 [23]. Data for LYSO:Ce was added according to reference [31].

Property	NaI:Tl	BGO	LSO:Ce	LYSO:Ce
Density [g/cm^3]	3.67	7.13	7.40	7.11
Effective atomic number	50	73	66	65
Decay time [ns]	230	300	40	41
Photon yield [per keV]	38	8	20-30	33.8
Index of refraction	1.85	2.15	1.82	1.81
Hygroscopic	Yes	No	No	No
Peak emission [nm]	415	480	420	420

In current PET scanners, LSO and LYSO are preferred due to their fast decay constant, so that the coincidence time window can be minimised and random rejection improved [32]. Moreover, the high light yields and lower refractive indices are beneficial for detector design. NaI has an even higher light yield, but the lowest Z_{eff} and a comparatively long decay time. Additionally, it is hygroscopic and fragile, which requires careful sealing and handling inside a detector block.

Light Sharing

After the interaction of the high-energy photon, the created scintillation photons are guided towards the photodetection unit. To prevent photons from escaping the crystal, its surfaces are usually coated with a reflective material (e.g. Teflon, 3M ESR vikuiti, titanium dioxide paint, Lumirror, Tyvek, barium sulfate). The surface which is facing the photodetector remains uncoated and is coupled to the detector with a glue or grease, that has a similar refractive index as the lightguide or entrance window of the photodetector (MeltMount, Grease, OP-20) [33]. In principal, two detector designs are possible, as depicted in figure 2.4: either the scintillator material has the shape of a single cuboid (monolithic design), or it is cut into many small crystal needles, which are arranged in an array and optically separated (pixelated design).

In both cases, either a position sensitive photodetector or an array of photodetector pixels has to be used to measure the light distribution. A dedicated algorithm is then




FIGURE 2.4: Two possible detector layouts: either a monolithic crystal is directly glued to a photodetector array, or a crystal array is used in combination with a lightguide, which distributes the scintillation light over multiple sensor pixels.

applied to calculate the position of the scintillation event from the measured light distribution (further details in section 2.2.6). Intuitively, the easiest case would be a one to one coupling of crystal needles and detector pixels, because the event position would directly correspond to that of the triggered pixel. Considering that each crystal requires its own readout and digitisation channel, this approach comes with high costs and complexity, so that only few PET systems make use of this technology (e.g. the Philips Vereos scanner [34]). Instead, most PET systems make use of the light sharing principle: the scintillation light is spread over multiple detector pixels, so that the event position can be calculated as the weighted mean of selected pixel signals. In monolithic scintillators the photons propagate within the whole crystal volume, resulting in a broad light distribution on the detector pixels. When an array of crystals is used, a lightguide has to be inserted between crystals and detector (usually a glass sheet) to spread the scintillation light over at least four pixels to apply the weighted mean algorithm. The event position (a_x, a_y) can be calculated from the detected charges $q_{i,j}$ on the detector pixels (i, j) , because the physical positions (x_i, y_i) of the N_x and N_y detector pixels are known from its geometry:

$$a_x = \frac{1}{Q} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} x_i \quad \text{and} \quad a_y = \frac{1}{Q} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} y_j \quad , \quad (2.3)$$

where Q is the sum of all detected charges q_i :

$$Q = \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} \quad . \quad (2.4)$$

Since Q is proportional to the deposited energy of the photon E_γ , equation (2.4) can be used for energy estimation.

This technique of position determination is called Anger logic, after its inventor Hal O. Anger [35]. It relies on the assumption that the crystal with highest measured number of photons is closest to the crystal where the scintillation event took place, and that the signal on the adjacent pixels is inversely proportional to the distance to the event position. It is also referred to as centre of gravity (CoG) algorithm.

Scintillation Light Detection

Collecting the greatest possible amount of the scintillation light and converting it into a measurable electric signal is the task of the photodetector. Until the 2000s, the gold standard for photon detection in PET devices have been photomultiplier tubes (PMT). Being bulky, expensive and not compatible with magnetic fields, they got replaced by the more compact avalanche photodiodes (APD), which were first used in an animal scanner in 1994 [36]. The high quantum efficiency and MRI compatibility but at the same time low gain encouraged further developments, including silicon photomultipliers (SiPM), which consist of Geiger-mode APDs (GAPDs) and which will be explained in detail in section 2.2.4. To optimise light collection, the sensitive spectral range should match the emission spectrum of the scintillator. As mentioned before, the amount of emitted scintillation photons is assumed to be proportional to the deposited energy of the high-energy photon. Figure 2.5 shows a typical energy spectrum of a single LSO crystal.

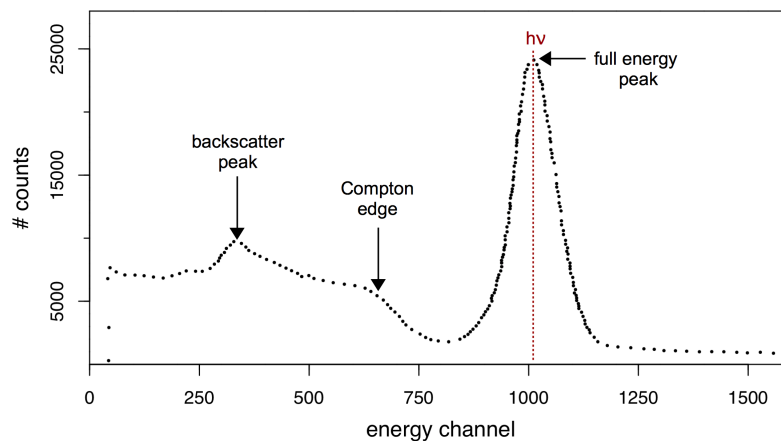


FIGURE 2.5: Energy spectrum of a single LSO crystal, which was coupled to a light diffuser and a photomultiplier. The spectrum was recorded using a ^{22}Na source. Data redrawn from reference [37].

When a photon deposits all its energy $h\nu$ in the crystal by means of photoelectric effect, this event will contribute to the "photopeak". This peak is also referred to as full energy peak, because the same effect can be observed when a photon undergoes multiple Compton scattering processes in the same crystal, and all scattered photons are absorbed within this volume. In case a scattered photon escapes, the event energy will be below $h\nu$, and thus contribute to the events between Compton edge and full energy peak. Single scattering processes give rise to a broad continuum in the pulse height spectrum (Compton continuum), which is cut off at the already mentioned Compton edge. This edge marks the maximum transferable energy in Compton scattering, i.e. when the photon is scattered backwards (180°) and the recoil electron propagates in the direction of incidence. The peak with lowest energy is the backscatter peak. It is formed by events where a photon is scattered back into the crystal from the detector electronics, shielding or housing, after having lost a

major part of its energy in the scattering process. Since the spectrum was originally recorded with a ^{22}Na source, there are also events with energies higher than $h\nu$, which are caused by the γ emission at 1275 keV [38] (\sim channel 2570) and signal pile-up in the detector (both not shown here). The data has been originally published by *Moszyński* in reference [37], where the full energy spectrum can be found. Further information on detector response functions for photons of different energy ranges was provided by *Knoll* in reference [39]. Suitable photodetectors will be described in detail in section 2.2.4.

For the described combination of scintillators and photodetectors, there exists a limit of energy resolution, which is caused by multiple effects. A detailed description can be found in *Inorganic Scintillators for Detector Systems* by *Lecoq et al.* [20].

The intrinsic energy resolution δ_{sc} of a scintillator is related to its non-proportional response, which means that the number of scintillation photons is not proportional to the deposited energy. During the photoelectric absorption of an incoming high-energy photon, "hot" electron-hole pairs are produced, which subsequently cause emissions and absorptions of X-rays and Auger electrons during their thermalisation. Since the secondary electrons can have different energies, the summed amount of produced scintillation photons along these single steps varies [40]. Furthermore, inhomogeneities in the crystal lattice can lead to local light yield differences (δ_{inh}). In case of LSO, *Balcerzyk et al.* determined an intrinsic energy resolution of $(8.9 \pm 0.6)\%$, which includes the non-proportionality and light yield inhomogeneity [41].

An additional degradation of the energy resolution δ_{tr} occurs during the transfer of scintillation photons to the photodetector and their subsequent detection, which can be summarised in the term light collection efficiency. This includes imperfections of the reflector surfaces surrounding the crystal, the optical coupling of crystal and detector, the inhomogeneous quantum efficiency of the detector etc. [40].

Lastly, there are two uncertainties, which are related to the photodetector: the dark noise contribution δ_{n} , and the statistical uncertainty of the photodetector. The latter can be described by Poisson statistics and is proportional to the inverse square root of the number of created photoelectrons N_{pe} : $\delta_{\text{stat}} \sim 1/\sqrt{N_{\text{pe}}}$ [20], [42]. The five effects, which degrade the energy resolution of a scintillator, all add in quadrature:

$$\left(\frac{\Delta E}{E}\right)^2 = \delta_{\text{sc}}^2 + \delta_{\text{inh}}^2 + \delta_{\text{tr}}^2 + \delta_{\text{n}}^2 + \delta_{\text{stat}}^2 \quad . \quad (2.5)$$

2.2.2 Choosing Components for a PET Design

In general, detectors can be arranged in different ways to establish a PET device. The most common geometry is a full ring scanner, where multiple detector blocks form a cassette and 20-30 cassettes are arranged in a circular structure. Each detector block consists of a scintillator (either array of crystals or monolithic), a lightguide (where applicable) and a photodetection unit. Usually, multiple blocks share a readout unit, which does pulse pre-processing such as energy thresholding, before

passing the information to the coincidence detection unit. An advantage of the ring geometry is the high solid angle coverage, which can be achieved by minimisation of cassette interspacing. The drawback are the relatively high production costs, since many crystal pixels are necessary for such a geometry. It has been attempted to overcome this drawback by building detector segments (planar or curved), which rotate around the patient in order to cover the full angle. However, this approach extends measurement time and does per se not yield the same sensitivity as a fully covered ring structure. Sensitivity, i.e. the ratio of detected photons and β^+ decays of the tracer, does not only depend on the geometric fill factor of the device, but also on the intrinsic detection efficiency, which is related to the scintillator material and thickness. In the next section, common crystal geometries will be described.

2.2.3 Crystal Geometries

As mentioned before, either a monolithic crystal block or an array of crystal pixels can be used for a PET detector, while both have advantages and disadvantages. A typical detector with a monolithic crystal has been shown in figure 2.4a. The large bulk volume of 15–20 mm length and 10–25 mm thickness is easy to manufacture [43], and thus a cheap geometry with high packing fraction providing high sensitivity. After the scintillation process, photons can spread within the whole crystal volume, and create a light distribution on the sensor, which is characteristic for the respective 3D event position in the crystal. In principle, DOI information is provided intrinsically in the light distribution, but especially at the outer edges a degradation in event positioning is a common problem due to light reflections at the crystal side faces. This degradation is becoming worse with crystal thickness [44]. Moreover, the large light spread can lead to enhanced pulse pile-up. Therefore, monolithic designs are not preferred for clinical scanners [44]. Another problem is the long and extensive calibration procedure, because the whole crystal has to be scanned with a pencil beam in a narrow grid, in order to interpolate the light distribution depending on event position. Even though calibration can be accelerated up to 2-3 h for a crystal of 2 cm length [43], the amount of data to be processed during a PET scan would be large. The data processing remains challenging, since the scintillation light spreads widely over multiple pixels, so that the signal to noise ratio (SNR) is low in some pixels.

When it comes to calibration and positioning, the advantage of a pixelated scintillator (figure 2.4b) becomes apparent: the detector pixel with highest signal is closest to the crystal where scintillation took place, so that the event position can already be calculated with standard Anger logic. In contrast to monolithic crystals, the calibration can be achieved without a collimated and carefully aligned pencil beam, but with a flood source, that irradiates the whole detector block at the same time. A disadvantage is the lower sensitivity: the reflective coating, which separates the crystal pixels optically, occupies space and reduces the geometric fill factor. Moreover, the cutting, coating and gluing process of the crystal arrays is elaborate and expensive.

When crystal pixel size is chosen, the aspect ratio influences strongly the light output, so that crystal pixels should not be thin and long in the ideal case [45].

Measuring Depth of Interaction (DOI)

Different designs have been presented with pixelated scintillators to determine the depth of interaction (DOI) of the scintillation event. Designs using monolithic crystals will not be described any further, because focus will be set on pixelated designs in this work; a review article by *Ito et al.* [9] contains further information on monolithic designs. Figure 2.6 gives an overview over the different options.

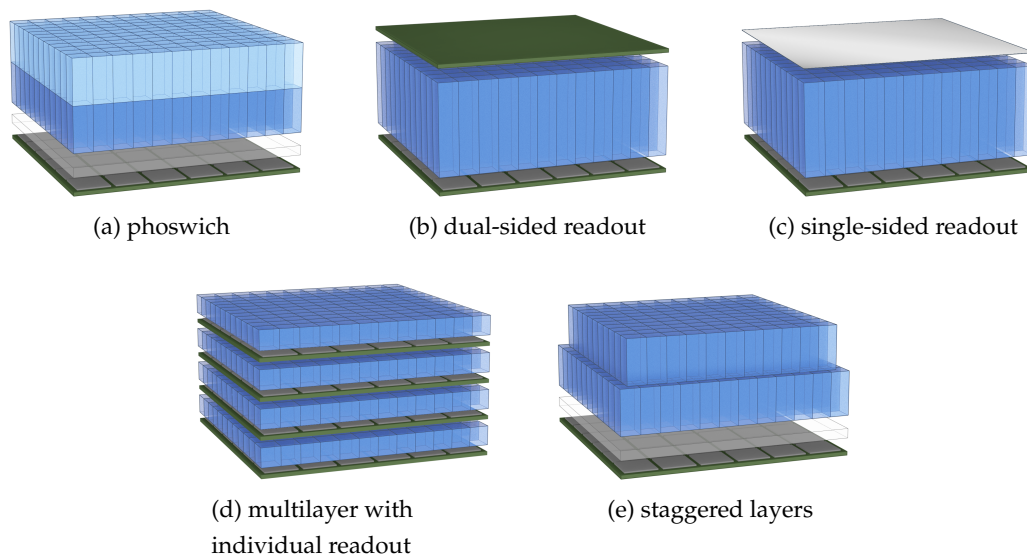


FIGURE 2.6: Different DOI capable detector geometries.

Figure 2.6a shows the concept of the phoswich ('phosphor sandwich') detector, where layers of scintillator materials with different decay times are stacked on top of each other. By analysing the pulse shape of a scintillation signal, it is possible to assign the event to the respective layer, so that DOI information is limited to the layer thicknesses. This concept has already been used for the ECAT HRRT scanner in 2002 [14], but research is still ongoing in this field [46]. A disadvantage are the relatively high production costs for the crystal layers.

Another approach, shown in figure 2.6b, has been used for the ClearPEM ([47]): LYSO crystal matrices ($2 \times 2 \times 20 \text{ mm}^3$ each crystal) are coupled to arrays of Avalanche Photodiodes (APDs) on both ends (dual-sided readout), so that the DOI information can be calculated from the asymmetry of light collection of both APDs, up to a resolution of 3 mm. *Yang et al.* reached a resolution of 2 mm FWHM with LSO crystals ($0.9225 \times 0.9225 \times 20 \text{ mm}^3$ each) and position-sensitive APDs [48]. Clinical use is usually limited due to high costs for twice the amount of APDs and the respective pulse processing units. One option would be to rotate two detector modules around the object (as *Yang et al.* did), but this inevitably decreases the sensitivity.

Trying to overcome this drawback, single-sided readout systems are focused in ongoing research (figure 2.6c). *Pizzichemi et al.* presented a detector, where DOI information is gained by reflecting the scintillation light at the top of the array back into the crystal array, in order to read out the direct and the reflected signal with a silicon photomultiplier (SiPM) at the other end of the crystal array. The ratio of both signals allows the determination of DOI with a resolution of 4 mm [49], [50]. *Ito et al.* achieved 4.3 mm average DOI resolution by inserting triangular shaped reflectors between the crystals, in order to manipulate light spread depending on the DOI [51].

A very intuitive but also elaborate approach has been realised by *Vandenbroucke et al.* [52]: multiple crystal arrays are stacked on top of each other, while each layer has its own super-thin position-sensitive APD (PSAPD) of less than 300 μm thickness. In contrast to other detector blocks, this geometry is irradiated from the side (edge on), not from the top (face on), so that DOI information is acquired along the longitudinally separated crystal pixels. With a crystal pitch of 1.0 mm intra-layer and 1.3 mm inter-layer, a corresponding DOI resolution of 1.0 mm has been achieved. However, this geometry requires a large amount of readout channels and is expensive in production (both crystal production and the electronic components). This makes it feasible for organ-specific PET systems, but not for a full detector ring dedicated to brain imaging. A schematic drawing of this setting is shown in figure 2.6d.

Using stacks of crystal arrays has been tested in various configurations, even a combination of two phoswich layers and another crystal layer made of one of the materials, shifted by half a crystal pitch in lateral dimensions [53]. Designs with two [54], four [55], [56], and eight crystal layers ([57], also phoswich) have been proven feasible. To explain the principle of staggered layer designs (depicted in figure 2.6e), the generation and interpretation of flood images will be briefly explained.

Flood Images

In section 2.2.1, the calculation of an event position from the measured light distribution using Anger logic has been explained. Usually, the detector pixel with the highest signal is selected as the "hot pixel", and the signals of the direct neighbour pixels are used for position calculation (eight neighbour pixels in the centre, five at the edges and three in the four corners). Irradiating a crystal array homogeneously leads to scintillation events in each crystal, so that Anger logic can be applied to calculate event positions for each crystal. This results in a 2D crystal position map, also called flood image or flood map, which contains an accumulation point for each crystal. An example for a flood image of a single crystal array is shown in figure 2.7. Since the light distribution cannot be measured completely at the outer edges of the photodetector, the Anger positions are shifted towards the centre. If a scintillation event is detected, its 2D position can be calculated with Anger logic, and the corresponding crystal can be identified, e.g. by searching for the spot in the flood map with smallest Euclidean distance (advanced positioning algorithms will be described in section 2.2.6).




FIGURE 2.7: Flood map of a crystal array with 30×30 crystals. An accumulation point can be identified for each crystal. Data redrawn from reference [58].

To gain DOI information, another layer can be added on top, but shifted by half a crystal pitch in both lateral dimensions, as shown in figure 2.6e. The scintillation light can then spread from the top layer to the bottom layer and further inside the lightguide, resulting in a broader light distribution on the photodetector. Since the layer is shifted relatively to the underlying one, the centroid of light dispersion is shifted accordingly in the flood map. This allows the user to distinguish events from both layers, and to acquire DOI information in this way. DOI information is limited to the thickness of the individual layers, so that refinement is attempted by using multiple, thinner layers. However, this entails some problems.

1. The flood image is a 2D histogram, and the intensity of accumulation points is proportional to the number of scintillation events, that took place in the respective crystals. To achieve points of similar intensity, similar amounts of photons would have to deposit energy in every layer, which is not possible due to the exponential attenuation in the crystal material.
2. The individual points have to remain separable, otherwise the hit crystal cannot be determined by means of the Anger position. Precise alignment of the shifted layers is absolutely necessary for good separability.
3. The performance of Anger positioning deteriorates towards the edges and in the corners of the detector block, because the light distribution cannot be measured completely.
4. The influence of the involved photodetector should not be underestimated: gain differences of the pixels can distort the expected regular spot pattern.

2.2.4 Photodetectors

Photodetectors are an essential part of a PET scintillation detector, since they are used to convert the weak scintillation light coming from the crystal into a measurable electric signal. The basic functionalities have been well described in literature, so that only a brief overview will be given.

Photomultiplier Tubes

Until the mid 2000s, photomultiplier tubes (PMT) have been considered the gold standard for photon counting in PET. The basic principle of PMTs has been described many times, and will be summarised based on reference [39].

In a PMT, photons impinge on a photo cathode and cause the release of photoelectrons due to photoelectric effect. A focusing grid accelerates them towards a chain of dynodes, which is held at increasing potential, so that multiple secondary electrons are emitted at each dynode. In this way, an exponential multiplication can be achieved, leading to gains of 10^6 to 10^7 . But to keep the multiplication constant over time, the operating voltage of 1 to 2 kV has to be highly stabilised, because the gain strongly depends on the dynode voltages. Another disadvantage is the bulky geometry, compared to other detectors, and the lack of MRI compatibility: due to the magnetic fields, electrons get deviated as soon as they leave the photocathode or dynode and the focusing and multiplication effect cannot apply anymore. The quantum efficiency (QE), which describes the probability that an incoming photon gives rise to an output signal, typically amounts to 25 to 30 % at maximum sensitivity wavelength.

PIN Diodes

Much higher quantum efficiencies can be achieved with solid-state photodetectors, which have been described in detail in a review paper by *Renker* [59]. The summarised information on PIN diodes and single photon avalanche diodes will be based on this publication.

A PIN diode consists of a semiconductive layer between two highly doped n+ and p+ regions. In the depletion zone of a PIN diode, the energy of scintillation photons is high enough to release electron-hole pairs, so that quantum efficiencies between 60 to 80 % can be achieved. On the one hand this is beneficial for the signal to noise ratio (SNR), but on the other hand a PIN diode does not provide any intrinsic gain, so the device requires high signal amplification with very low noise. Another advantage is the sleek design, which allows for use in magnetic fields and cheap large scale productions.

The design has been further refined, resulting in the so-called avalanche photodiode (APD) in the 1990s. APDs were extended by an additional, highly doped p+ or n+ layer next to the depletion layer, in order to increase the electric field strength in the multiplication zone. Moreover, the diode is operated at reverse bias voltage,

close to the breakdown voltage. The acceleration of primary electron-hole pairs is then strong enough to provide enough energy for creation of secondary electron-hole pairs by impact ionisation, making single photon counting possible. However, the internal gain still only amounts to 100 to 500 [59], necessitating readout electronics with external gain of 10^4 to 10^5 . To achieve constant performance (especially gain) in a measurement, operation at stable bias voltage and ambient temperature is inevitable.

Single Photon Avalanche Diodes

As an alternative to the described linear mode, APDs can be operated at a reverse bias voltage above breakdown voltage, which is called Geiger mode. The devices are then called single photon avalanche devices (SPAD), because a single photon can trigger a self-sustaining avalanche in this mode of operation, yielding gains of 10^5 to 10^7 . The avalanche has to be stopped by a dedicated quenching circuit, after the avalanche current has given rise to the output signal. The output signal is standardised and independent of the energy, which was deposited in the multiplication zone. In analogy to a Geiger counter, SPADs are also called Geiger-APDs (G-APDs). Quenching can be done passively or actively with a dedicated quenching circuit. Despite better timing performance speaks in favour of active quenching, passive quenching has the advantage of simple implementation with a single series resistor. This leaves more space on the substrate and enhances the geometric fill factor ϵ_{geom} . The Geiger discharge itself is fast (rise time of several hundred ps [60]), while the recharge time can take several 100 ns, depending mainly on the quenching resistor R_Q and cell capacitance C : $t_{\text{decay}} \sim R_Q C$ [59]. During the recovery phase ("dead time"), a SPAD cannot discriminate multiple impinging photons. When combined with a scintillation crystal, several thousand optical photons have to be measured in a short time span, e.g. ~ 15000 photons per 511 keV photon in case of LSO, while a typical pulse integration time is 300 ns. The number of optical photons and pulse integration time depend strongly on the used scintillator.

Silicon Photomultiplier

To circumvent this drawback, so-called multi-pixel photon counters (MPPC) or silicon photomultipliers (SiPM) have been built by connecting thousands of small micro-cells in a parallel circuit. Each micro-cell consists of a G-APD and its related quenching resistor. Every G-APD contributes to the total (analogue) output signal and usually has dimensions in the range of $20 \times 20 \mu\text{m}^2$ to $100 \times 100 \mu\text{m}^2$ [61]. The overall photon detection efficiency (PDE) amounts to $\sim 40\%$ [61], although high quantum efficiencies of 50 to 80% are possible (wavelength-dependent) [59]. But also the geometric fill factor ϵ_{geom} and the probability that a photon causes a Geiger discharge ϵ_{Geiger} contribute to PDE [61]:

$$\text{PDE} = \text{QE} \times \epsilon_{\text{geom}} \times \epsilon_{\text{Geiger}} \quad . \quad (2.6)$$

The geometric fill factor ϵ_{geom} describes the ratio of active cell area and the total SiPM area including quench resistors, bond gaps etc. where no avalanche can be triggered. However, not only the beneficial characteristics like high QE are inherited from SPADs: as each semiconductor detector, SPADs have an internal noise source, which is referred to as dark counts. Electron-hole pairs can be generated thermally and then trigger an avalanche, even though the SPAD is not irradiated at all. The frequency of dark counts increases with temperature (e.g. order of MHz/mm² at room temperature, kHz/mm² at -20°C) [62]. On the contrary, the probability for afterpulsing decreases with increasing temperature. Afterpulsing occurs when carriers get trapped in metastable energy states during an avalanche. Several nanoseconds after the original avalanche has been quenched, they are released and trigger an additional, time correlated avalanche. However, this avalanche is only detectable when the lifetime of a trapped carrier exceeds the recharge time of the SPAD, and this lifetime decreases with increasing temperature [63].

Another degrading effect which occurs in SiPMs (but not in SPADs) is optical crosstalk. Ideally, photons contribute only to the micro-cell signals, in which they triggered an avalanche. But optical photons can also trigger secondary avalanches in neighbouring micro-cells, when they were emitted during the avalanche process (efficiency for photon emission with energies larger than 1.14 eV is $\sim 3 \times 10^{-5}$ for an electron passing the junction [64]). The probability that such an optical photon initiates a secondary avalanche in a neighbouring cell is between 1 and 50%, depending on the overvoltage, gain and position where they were produced [65]. Approaches to suppress optical crosstalk, e.g. by inserting optical trenches between the cells, usually deteriorate the geometric fill factor and are thus depreciated.

Another disadvantage common for APDs and SiPMs is the approximately exponential dependence of their gain on temperature and overvoltage. According to *Lecomte* [61], the relative change of gain is 30%/V at gain 1000, and temperature changes by $\sim 15\%/^\circ\text{C}$. This dependency requires careful operation of APDs and SiPMs at both very stable temperature and overvoltage. The signal of all micro-cells is summed to a single analogue output signal, which is analysed regarding shape and amplitude. Digital SiPMs, as developed by Philips Digital Photon Counting [66], digitise the pulses at a very early stage, i.e. on micro-cell level by using CMOS logic circuits to count the avalanches from discharged cells. The number of discharged micro-cells on a pixel is output together with a timestamp generated by a time-to-digital-converter (TDC). Since this technology is applied in the present work, the next section will describe its features in more detail.

2.2.5 Philips Digital Photon Counting

The following information on the digital photon counter (DPC) from Philips Digital Photon Counting (PDPC) has been compiled from the references [66]–[68], if not indicated otherwise.

In 2009, *Frach et al.* presented their concept of a digital SiPM for the first time: while

the principal network of G-APDs is used as in analogue SiPMs, each micro-cell is equipped with its own comparator and an active quenching and recharge circuit. Furthermore, each cell is connected to a counter and (via a configurable trigger network) to a time-to-digital converter (TDC).

The advantage of the novel design is digitisation in a very early phase, superseding additional hardware for signal filtering, pulse shape analysis and digitisation. The risk of affecting the analogue signal (mV range) along the interconnections by interference can be minimised. This enables single photon triggering, which is not possible with analogue signals, since the single photon trigger is hard to distinguish from electronic noise or the baseline with dark counts.

Another advantage are one bit memory cells, which are connected to the SPADs in order to enable or disable them. This allows the user to record and upload a so-called inhibit map to a flash memory chip on the back of the sensor: cells with excessively high dark count rates can be disabled to reduce the overall dark count rate (DCR). The cumulated DCR can be lowered from 14286.04 kcps at 100 % enabled cells to 3652.37 kcps by disabling 10 % of the cells with highest DCR (values obtained at 20 °C operating temperature in [66]).

Unlike analogue SiPMs, SPADs can only count exactly one photon until recharge. The total photon count on a pixel is output together with a timestamp, which is set according to the selected trigger scheme. Before explaining customisable trigger and validation schemes, the sensor architecture will be briefly described in the following.

Sensor Architecture

The sensor architecture is split hierarchically into different levels: a full tile consists of 4×4 dies, which hold 2×2 pixels. In the model DPC-3200-22-44, each pixel contains 3200 micro-cells, which are organised in 2×2 subpixels with 32×25 micro-cells each ($59.4 \mu\text{m} \times 64 \mu\text{m}$ cell size). The cell rows in a subpixel are called row-trigger-lines (RTLs) and form groups of three or four rows for the validation scheme.




FIGURE 2.8: The sensor architecture of a Philips dSiPM shown in a photograph of a tile: it consists of 16 dies with 4 pixels and 16 subpixels each. A subpixel is further divided into eight row-trigger-lines.

Moreover, the chip is equipped with an FPGA on the back to configure and synchronise the dies. A temperature sensor can be used to control bias voltage and to adjust it accordingly during event acquisition.

Event Acquisition

The general acquisition sequence of a DPC die is shown in figure 2.9. This sequence applies for all dies, which operate independently, but share the same system clock. In the beginning, all micro-cells are charged and ready to detect a scintillation event. When a sufficient number of photons is registered to fulfill the trigger condition, the event timestamp is set (t_0) and the validation phase is started. It is noteworthy that multiple dark counts can accumulate over time and thus fulfill the trigger condition. Therefore, a validation scheme is applied to allow the user discrimination of dark count events and true scintillation events. To prevent dark count triggers, RTL refresh can be enabled, so that all micro-cells are recharged every 10 ns, when no trigger has been detected in that interval.

Only when the selected validation scheme is fulfilled, the sensor goes to the integration and recharge phase to be ready for the next scintillation event. The bottom row of figure 2.9 shows the time scale of the event processing pipeline. When a trigger has been set, 10 ns to 15 ns are used to switch to the validation mode. The time intervals for validation, integration and recharge can be set by the user, while the readout time is fixed to 680 ns.

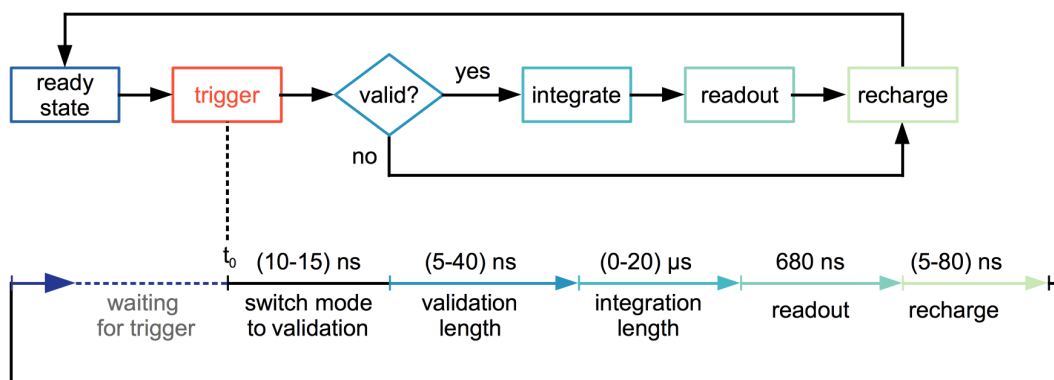


FIGURE 2.9: The event acquisition pipeline of a Philips dSiPM: when the trigger has been detected, a validation scheme is used to distinguish true scintillation events from dark counts. After integration and readout, a global recharge sets the die back to ready state. Redrawn from reference [66].

Trigger Logic

The user can select four boolean connections of the four subpixels of a pixel, listed in table 2.3.

Trigger scheme 1 denotes the lowest threshold, because a single photon in any of the subpixels suffices to set the trigger. Scheme 4 is the opposite case, where every subpixel has to detect a photon to set the trigger. Schemes 2 and 3 are intermediate

TABLE 2.3: Overview of the selectable trigger schemes. Adapted and republished from reference [69]. ©IOP Publishing Ltd and Sissa Medialab srl. Reproduced by permission of IOP Publishing. All rights reserved.

Scheme	Interconnection of subpixels (SP)	Minimum number of fired cells	Expected number of fired cells
1	$SP1 \vee SP2 \vee SP3 \vee SP4$ $[(SP1 \vee SP2) \wedge (SP3 \vee SP4)]$	1	1.0
2	\vee $[(SP1 \vee SP4) \wedge (SP2 \vee SP3)]$	2	2.33(± 0.67)
3	$(SP1 \vee SP2) \wedge (SP3 \vee SP4)$	2	3.0(± 1.4)
4	$SP1 \wedge SP2 \wedge SP3 \wedge SP4$	4	8.3(± 3.8)

thresholds, as specified in table 2.3. The expected number of discharged cells at the time the condition is fulfilled was determined by *Tabacchini et al.* [69], and is also given in table 2.3. When the required number of photons is registered in the respective subpixels, the trigger is set, and the exact time when the scheme is fulfilled is stored in the timestamp. The acquisition sequence is started hereafter, although 10 ns to 15 ns are required to switch to the validation mode.

Validation Scheme

When the trigger has been set, the sensor goes to the validation phase, which can be configured to last between 5 and 40 ns. During this time, more photons cause micro-cell discharges. The spatial distribution of these cell discharges is then compared to a selectable pattern, the validation scheme, in order to validate that it is a scintillation event. For this purpose, each subpixel is divided into eight RTL groups with three to four RTLs each (see figure 2.8 and 2.10). As soon as a single RTL detects a photon, the logic gate of the RTL group is set to logic high. The signals of all gates in a subpixel are eventually handed to the top-level pixel gate.

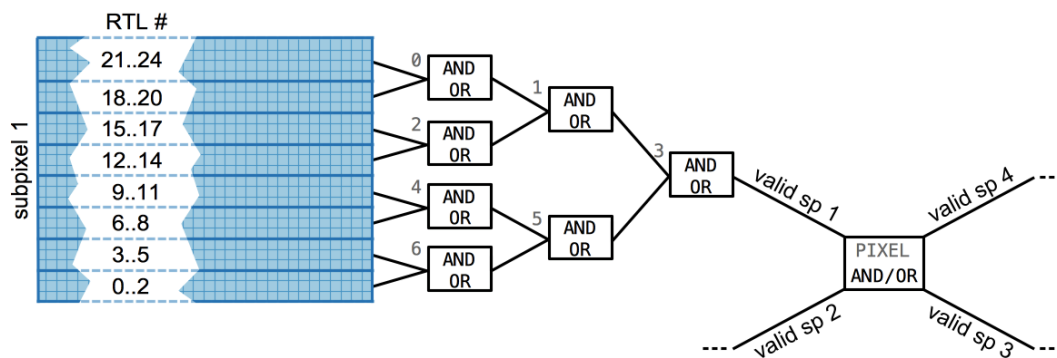


FIGURE 2.10: Each subpixel is divided into eight RTLs. To validate an event, the geometric distribution of photons in the RTLs is compared to a selectable pattern, which can be modified by setting gates 0-6. Redrawn from reference [66].

For fulfilment of the highest threshold (0x00), each RTL group of a subpixel has to record at least one discharge, to set the subpixel gate to valid. In case of the lowest threshold (0x7F), one discharge on the whole subpixel suffices to validate the subpixel. Since it does not make any difference if one or multiple cells discharged in a RTL group, the validation threshold is stochastic, so that table 2.4 gives minimum and average numbers of cell discharges, that can be registered when the trigger scheme is fulfilled.

TABLE 2.4: Overview of the selectable validation schemes (adapted from references [66], [69]).

Scheme	Validation pattern	Subpixel connection	Minimum number of fired cells	Average number of fired cells
1	0x7F : OR	OR	1	1
2	0x77 : OR	OR	2	4.7(\pm 2.1)
4	0x55 : OR	OR	4	17.0(\pm 6.2)
8	0x00 : OR	OR	11	53.0(\pm 15.0)
16	0x55 : AND	AND	16	54.0(\pm 19.0)
32	0x00 : AND	AND	32	132.0(\pm 40.0)

After validation of an event, the whole die is set to integration mode. In a configurable time interval of 0 ns to 20 485 ns further photons are collected. In the subsequent readout phase of 680 ns, the states of the micro-cells are read and recharged row by row. Rows which have not yet been processed can still collect more photons. During this process, the number of cell discharges is counted at pixel level. The recharge phase takes 20 ns by default, but can be customised by the user (0 ns to 80 ns). At the end of the acquisition cycle, the number of counted photons on the four pixels are sent out to external data processing along with the trigger timestamp. In case the event could not be validated, the die is put to recharge mode directly.

Neighbour Logic

The independently operating dies can also be connected in a master and slave hierarchy. When a master die registers a valid scintillation event, the vertically and horizontally direct neighbouring dies go to the integration phase as well, without waiting for their own trigger and validation. This can be helpful when the light distribution is highly focused, so that the signal of neighbouring dies could not exceed the trigger threshold. However, slave dies will ignore the master trigger signal, when they are currently recharging. Moreover, up to 20 ns can be necessary to transfer the master trigger signal to the slave dies, so that photons from this time interval will not be recorded by the slave dies. This can be a limitation for fast scintillation crystals. In case of LSO, which has a rise time of $\tau_r = 80$ ps and a decay time of $\tau_d = 39.6$ ns [70], nearly 40 % of the photons would therefore not be recorded when neighbour logic would be enabled.

2.2.6 Positioning Algorithms

It has been mentioned before, that dedicated positioning algorithms are applied to determine the crystal, where a scintillation event took place. These make use of the measured light distribution on the photodetector.

Before event positions can be calculated, a calibration step is necessary. The starting point is a two-dimensional flood histogram, where the single event positions are calculated with Anger logic (equations 2.3). The photodetector pixel with the highest photon count is regarded as the *hot pixel*, which is closest to the crystal where scintillation took place. The direct and diagonal neighbours are used for Anger calculation, i.e. eight neighbours for a hot pixel in the centre region, five neighbours at the border and three neighbours in the corner of the sensor tile.

Depending on the trigger and validation schemes, it is possible that not all pixels record a signal of the event. Also the independent operation of dies can lead to this case, because a die might still be in the readout or recharge phase (dead time), when neighbouring dies already register the next event. Applying the centre of gravity algorithm to the incomplete data set nonetheless yields positioning errors, which can be in the range of the crystal pitch: *Lerche et al.* investigated the positioning error for an array of LYSO crystals with 1 mm pitch, by taking complete sets of measurement data and removing randomly $n = 0 \dots 6$ SiPM signals. The deviations of the calculated Anger positions for the respective cases showed, that already a single missing SiPM signal can yield a positioning error larger than the crystal pitch.

Even more severe is the case for four missing SiPMs, because more erroneous than correct positions have been acquired [58]. *Schug et al.* overcome this problem by extrapolating the missing signals [71], which enhances the sensitivity, because incomplete event signals can be processed, though.

In the ideal case, each crystal can be assigned an accumulation point in the flood image. The calculated event position can then be matched to a crystal, e.g. by segmentation of the flood image and nearest position search or via a look-up table. These approaches work well in combination with single crystal arrays, but have drawbacks in multilayer geometries, where a considerably higher amount of crystals is used. Especially at the borders and at the corners of the tile, accumulation points in the flood image are focused towards the tile centre. This effect can cause very little distances or even overlap of the points (depending on individual die gains), so that they have to be distinguished by more parameters than just the 2D position, e.g. energy or spot size.

Alternatively, neural network based algorithms can be used to build the crystal look-up table, as shown by *Hu et al.* for the Siemens Inveon scanner [72]. However, missing data in the measured photon distribution remains a challenge, since all possible combinations of missing SiPM signals would have to be considered in the training network to adequately handle missing data during the event positioning.

Maximum Likelihood Based Positioning

An alternative to the described algorithms is maximum likelihood positioning, although this is predominantly used with monolithic detector blocks (as e.g. realised by *Espana et al.* [73]). It has been shown by *Lerche et al.* [58] that application to a crystal array for a preclinical PET system is also feasible.

The algorithm will be briefly described, based on the references [58], [74], if not denoted otherwise.

The general idea is to develop a model, that is trained on an existing data set, and thus capable of describing the observed (and maybe future) data sets by an appropriate choice of parameters. The "likelihood" describes how well the model (with a set of parameters) and the observed data match, or in other words, how likely the model with the set of parameters generated the data set.

Let $p(\mathbf{q}|S)$ denote the probability density function (PDF), that a set of charges $\mathbf{q} = (q_1, \dots, q_k)$ is measured on SiPMs given a scintillation event with S created scintillation photons (in total). Since the measurements are statistically independent, the PDF can be rewritten as the product of individual PDFs:

$$p(\mathbf{q} = (q_1, \dots, q_k)|S) = \prod_k p_k(q_k|S) \quad (2.7)$$

In reality, the inverse case has to be considered: the charge pattern \mathbf{q} is observed, given that S scintillation photons were created. The goal is to find the right set of parameters and thus the PDF which most likely generated the observed data set. For this purpose, the Likelihood function \mathcal{L} is defined

$$\mathcal{L}(S|\mathbf{q}) = p(\mathbf{q}|S) \quad (2.8)$$

and maximised, this process is also called *Maximum Likelihood Estimation* (MLE). Maximising the product in 2.7 can be computationally demanding, so that calculations with a sum are preferred instead. This conversion is possible when computing the natural logarithm of \mathcal{L} (log-likelihood), and also justified, because the logarithm is monotonically increasing, thus the maxima of both functions coincide. Since the measured photon counts follow the Poisson distribution [75], equation 2.8 can be extended to

$$\mathcal{L}_i(i, S|\mathbf{q}) = \ln \prod_k \frac{e^{s_{i,k}} s_{i,k}^{q_k}}{q_k!} \quad (2.9)$$

In equation 2.9, $s_{i,k}$ is the expected amount of scintillation light in crystal i , while S denotes the sum of all scintillation photons. Therefore, $s_{i,k}$ can also be described by the energy-normalised matrix $c_{i,k}$, containing the light fraction detected by SiPM k after scintillation in crystal i : $s_{i,k} \approx c_{i,k}S$. With simple algebraic manipulation equation

2.9 can be rearranged to the following form:

$$\mathcal{L}_i(i, S|\mathbf{q}) = \sum_k q_k \ln [c_{i,k} S] - c_{i,k} S - \ln(q_k!) \quad . \quad (2.10)$$

Finding the maximum of $\mathcal{L}_i(i, S|\mathbf{q})$ yields the crystal i^{ML} which was most likely hit:

$$i^{\text{ML}} = \arg \max_i \mathcal{L}_i(i, S|\mathbf{q}) \quad . \quad (2.11)$$

Being independent from i , the term $\ln(q_k!)$ in equation 2.10 can be omitted. The total energy can then be updated:

$$S^{\text{ML}} = \frac{\sum_k q_k}{\sum_k c_{i^{\text{ML}},k}} \quad . \quad (2.12)$$

Eventually, the following sequence is used for iterative maximum likelihood estimation.

1. Begin with the sum of detected charges as energy prior:

$$S^0 = \sum_k q_k$$

2. Calculate the log-likelihood (2.10) and the crystal index (2.11), then update the total energy according to 2.12.
3. Repeat 2. with the updated energy from 2.12.
4. If $\mathcal{L}_i(i, S|\mathbf{q})$, i^{ML} and S^{ML} have not changed, stop the iteration. Else if the maximum number of iterations is reached, stop as well.

Afterwards, the energy can be converted from photon counts to an energy in keV. *Gross-Wege et al.* showed that MLE is more robust regarding light distributions with missing SiPM signals, both for spatial and energy resolution [76]. Sensitivity could thus be increased, since incomplete events do not have to be filtered out, but can be processed with MLE. A further advantage of this algorithm over conventional CoG is its flexibility: the model used for MLEM can be adapted in such a way that it accounts for system parameters which have an influence on its spatial resolution. Moreover, it could be trained to reject events that do not match the selected statistical model (e.g. scattered events) [75].

2.3 Hybrid Imaging

Compared to other imaging modalities like computed tomography (CT) or magnetic resonance imaging (MRI), PET yields a low spatial resolution, and furthermore a different kind of information, i.e. information on metabolic processes, instead of anatomical or structural information. The combination of CT or MRI with PET is therefore an obvious approach to benefit from the strengths of both techniques. However, the combination is technically challenging: while a combined PET/CT device

was first developed between 1995 and 1998 [77], the first combined simultaneous human whole body PET/MRI scanner was developed more than 15 years later [78]. The main reason is the high static magnetic field B_0 (typically 1.5 T to 3.0 T in clinical environment, but even up to 9.4 T in research), which made the use of (at that time common) photomultiplier tubes impossible.

The static magnetic field in an MRI system is used to split energy levels of spin $1/2$ nuclei according to the Zeeman effect. In the thermal equilibrium state, the Boltzmann distribution of the two possible spin alignments (in case of ^1H) leads to a small net magnetisation in direction of B_0 . Radio frequency (RF) pulses can be applied with the nucleus-specific Larmor frequency to disturb the equilibrium state and flip the net magnetisation vector away from B_0 (excitation). There are characteristic relaxation times T_1 (longitudinal) and T_2 (transversal) for tissues, when the net magnetisation vector recovers to equilibrium state. By means of a receiver coil the magnetisation can be detected as an induced electric current. To assign spatial information to the signal, additional linear gradient fields are superimposed for frequency and phase encoding¹.

2.3.1 MRI compatibility of the PET insert

A PET insert for an MRI scanner is therefore exposed to the strong static magnetic field, the rapidly switching gradient fields and emission of high power RF. Especially the strong permanent magnetic field B_0 makes the use of ferromagnetic materials (often used in pins and sockets) impossible: on the one hand, ferrite-containing electronics become ineffective at a static magnetic field of already 0.3 T; on the other hand, they distort the homogeneity of the magnetic field. To ensure diagnostic image quality, field variations of only 0.1 ppm (i.e. 300 nT in 3.0 T MRI) are tolerable. A variation of this size can already be caused by a 0.1 mm^3 ferrite volume at 2.5 cm distance from the object [80]. Also para- and diamagnetic materials can cause problems, but are considered compatible as long as the magnetic susceptibility does not differ more than 10^{-2} from the susceptibility of water ($\chi_{\text{H}_2\text{O}} = -9.05 \times 10^{-6}$ at 37°C), and as long as they are not located close to the object. LSO, fused silica glass and Copper are compatible and do not interfere with the magnetic field homogeneity (a detailed list of the respective susceptibilities was compiled by *Schenck* in reference [81]).

Photomultiplier tubes cannot work in the strong static magnetic field, because photoelectrons can get deviated before reaching the chain of dynodes. Hybrid imaging with MRI therefore necessitates the use of APDs or SiPMs, as used in the GE Signa, Siemens mMR or Philips Vereos scanner. Since silicon has a susceptibility of $\chi_{\text{Si}} = -4.2 \times 10^{-6}$ [81], it is MRI compatible. Using the silicon based detectors has another advantage: they are compact, which is beneficial due to limited space inside the magnet bore hole. SiPMs further have better time resolution, which can increase the signal to noise ratio [70].

¹The focus of this thesis is on PET detectors, so the working principle of MRI is not discussed in detail. For further reading, see e.g. [79].

Apart from magnetic susceptibility, the electric conductivity has to be taken into account as well: eddy currents can be induced by the switching gradient fields on the surfaces of conductive boards causing heating and vibration. Power supply and printed circuit boards (PCB) can also be affected, leading to signal alteration or additional noise [82].

The RF pulses can furthermore cause interferences in the PET electronics, lowering the count rates during MRI acquisition [5]. Dedicated RF shields exist, often consisting of conductive layers that surround a PET detector cassette. Since they can in turn give rise to eddy currents themselves, various approaches using slitted or structured layers of copper or carbon fiber composites exist. Detailed information can be found in references [83] and [84].

2.4 Monte Carlo Simulations

A widely recognised modelling tool for medial imaging is the Monte Carlo (MC) simulation method. The essential part of simulating an event in PET is tracking a particle along its path, from emission at the source through a patient or phantom, until it is registered in the detector block. The underlying physics processes are defined by probability density functions (PDFs), so that the particle trajectories can be determined by sampling random numbers from these PDFs. The accuracy of a MC simulation strongly depends on the underlying model, which defines the computation of energy loss, and on the data bases or parameterisations, which are used to interpolate the cross-sections [85], [86]. To create a set of random numbers, the seed of the simulation (set by the user or defined automatically) is fed to a random number engine like Ranlux64, Mersenne Twister or James Random.

At this point, the benefit of MC simulations becomes obvious: they allow the user to investigate physics processes in detail, e.g. the amount of scattered photons in a phantom, or alternatively to exclude processes from the simulation, or even to test different materials. This contributes to a better understanding and improvement of the model. In general, parameters, which are impossible to realise in an experiment, can be chosen, and the influence on image resolution can be studied in detail, because the simulation allows to record a ground truth for all events.

The disadvantage of this approach is the large number of events, which is required for minimisation of the statistical error. This is also the reason why the MC method became relevant with the development of computers, although the approach has been used before, e.g. by Enrico Fermi in the 1930s (unpublished work) [87]². Moreover, only those effects can be modelled which are already fully understood and can be accurately described by a model.

²The modern MC method was first applied during the second world war on the electronic computer ENIAC, while Stanislaw Ulam and John von Neumann were working on neutron diffusion in matter. Inspired by the famous casino of Monte Carlo, their colleague Nicholas Metropolis suggested *Monte Carlo* as a code name for Ulam's "concept of lucky numbers" [87].

Nowadays, Monte Carlo simulation is the method of choice when it comes to the design or characterisation of new detectors for nuclear imaging. To give some examples, investigation of the scatter fraction [88], influence of event positioning algorithm on the overall scanner performance [89], new concepts for Cherenkov-based PET detectors [90], [91] or a combined experimental and simulation study for a monolithic detector [92], [93] have been simulated. Moreover, models for common PET systems are validated, so that parameters like the FOV or timing performance can be varied to understand how the current device could be improved [94], [95].

2.4.1 GATE: Geant4 Application for Tomographic Emission

A powerful toolkit for Monte Carlo simulations is GATE, a "Geant4 Application for Tomographic Emission" [16]. It is an open source software maintained by the Open-GATE collaboration and the users, who can contribute to the software development in a Git repository. While GATE was first implemented as a toolkit for PET and SPECT [96], further dedicated modules were developed for optical imaging [97] and for radiotherapy [86], addressing a large worldwide community. GATE uses the Geant4 kernel for the simulation of particle propagation in matter, but provides a user-friendly macro language to set up the simulation and its geometry in an easily understandable way.

The following paragraph is based on information obtained from the User Guide and the GATE source code³.

Structure of a Simulation

To set up a GATE simulation, the geometry and allowed physics processes have to be specified first. Predefined models like cylindricalPET or scanner can be selected to model a full PET ring or single detector blocks. The simple macro language allows the user to create his own virtual phantom by standard shapes (cube, sphere, etc.), or to load a voxelised virtual phantom. The individual volumes are defined by their shapes, dimensions, materials and position in the simulation *world*. A hierarchical order can be implemented by defining mother and daughter volumes, e.g. a whole detector cassette, which consists of multiple detector blocks, made of an array of crystals each. Each level of the *sensitive detector* can be assigned a readout level, so that in this example the readout could take place on cassette, block or crystal level. The physics processes can either be added manually, or by applying a predefined physics list from Geant4.

An important step in the middle of the simulation is the initialisation: at this point, the cross section tables for the physics processes in the used materials are calculated, and neither the geometry nor the physics processes can be changed anymore. Afterwards, the source is defined in the init mode. Moreover, the digitiser is defined,

³User guide: <https://opengate.readthedocs.io/>
Git repository: <https://github.com/OpenGATE>.

which is used to model effects of the readout electronics like crosstalk, dead time or an energy blurring. The digitiser also combines individual *hits* to pulses, and later *singles* and *coincidences*, according to the set readout level. Before the simulation can be run, data output (ASCII, ROOT tree or binary) and experiment settings like the simulation time and the random generator have to be specified. Assuming that there are time periods during which the simulated system is static, the whole simulation time can be split into multiple slices. Each time slice is treated as a G4Run, and each emission of a primary particle of the source (e^+ , γ , two annihilation photons) is simulated as a G4Event.

Digitisation of Hits

When a primary particle is emitted from the source, its interactions in the phantom and later in the detector are stored as hits. Each primary particle is assigned an event ID, which is passed on to all secondary particles, that are created along the trajectory. In fact, it is possible to store every single hit inside a detector, but this requires a large amount of disk space and extends the computation time. Usually the hits are converted to singles and coincidences, before an output manager converts them into ROOT or binary output. Being important for section 6.4, the digitisation workflow will be briefly explained.

At the end of each event, i.e. when the particle and all secondary particles left the *world* or were stopped, the *digitiser* is called. All hits are stored in a CrystalHitCollection (CHC), which is passed to the hit converter. Hits with non-zero energy are converted to pulses and added to the pulse list called "Hits". A GatePulseProcessor-Chain processes the digitiser modules sequentially, as given by the user (e.g. noise, crosstalk, energy blurring, adder for singles calculation). The standard *singlesadder* combines all pulses with the same event ID within the same readout volume (crystal or block) to a single: its position is the centroid of all hit positions, weighted by the deposited energy, and the singles energy corresponds to the sum of all energy depositions during the hits. Instead of the standard adder, a module called *adderCompton* can be used to store the first interaction position in the readout volume instead of the centroid weighted by energy depositions. However, the standard pulse adder is usually preferred, since it is more similar to the readout of scintillation detectors (these determine the event position from the measured scintillation light distribution, and lead to the detection of the energy-weighted centroid of all interactions).

Afterwards, a coincidence sorter determines the coincidences from the singles, and stores them by means of the DigiMaker. It is then possible to convert the singleDigiCollection and coincidenceDigiCollection to ROOT, text or binary output. The pulse lists are then deleted to enable processing of the next event.

Chapter 3

Motivation and Thesis Outline

The objective of this thesis is the design and characterisation of a novel human brain PET insert, which can be integrated into a magnetic resonance tomograph. A major part of the present work has been carried out in the scope of the *Helmholtz Validation Fund* (HVF), which has been granted in 2015 to develop a next-generation BrainPET for 7T MRI.

3.1 Requirements for a BrainPET Scanner

PET scanners dedicated to brain imaging have enhanced image quality, i.e. higher spatial resolution and sensitivity compared to whole body scanners, which is advantageous for clinical and research studies. This has been demonstrated with the high-resolution research tomograph (HRRT) [98]: while the spatial resolution of whole body scanners is 3.5 mm to 5 mm [6], [7], 2.3 mm could be achieved with the HRRT in the centre field of view (FOV), and 3.2 mm at the edges, enabling the visualisation of signals from small structures within the brain. The comparatively small size of the crystals ($2.1 \times 2.1 \times 7.5 \text{ mm}^3$) plays a major role in the enhanced spatial resolution [15].

Nevertheless, there are (to the best knowledge of the author) currently no dedicated brain PET scanners commercially available. In 2008, one of four prototypes for a Siemens BrainPET combined with 3T MRI has been installed at the Research Centre Jülich, the other three devices were located in Tübingen, Boston and Atlanta [15]. The Siemens BrainPET provides a spatial resolution of 3.0 mm in the centre FOV (tangential, radial and axial direction), but at a radial offset of 75 mm the radial image resolution degrades to 4.39 mm [15]. Since the detectors cannot yield any depth of interaction (DOI) information, a correction of the parallax error is not possible. The HRRT has a larger ring diameter and is equipped with two layer phoswich detectors, therefore this drawback can be overcome [98].

The aim of the Helmholtz Validation Fund project "BrainPET 7T" is to develop an ultra-high field MRI compatible, dedicated brain PET scanner, which has a high sensitivity, high spatial resolution, and allows the acquisition of DOI information.

Depth of Interaction Encoding and Parallax Error

Figure 3.1 shows three views of reconstructed PET images of a brain phantom, that have been simulated in the Siemens BrainPET¹. The coronal and sagittal phantom images clearly show that the cortex can be well resolved in the regions near the scanner axis (z axis), which corresponds to the centre of the transverse field of view. These structures appear more and more blurred, when they are located farther away from the z axis (marked with red arrows in figure 3.1). This effect can also be observed in the transverse plane (figure 3.1c). The scanner axis is pointing through the centre of the image, which is marked by the black circle. While the structures in the centre of the plane appear clearly separated, the structures blend into each other at the outer rim of the brain.

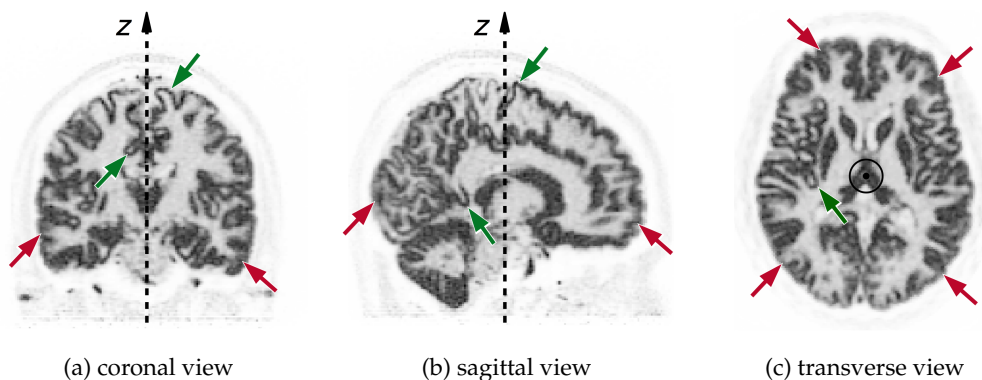


FIGURE 3.1: Coronal (a), sagittal (b) and transverse (c) view of the reconstructed images of a simulated phantom (BrainWeb [99]): while the grey matter structures are well resolved in the central transverse field of view (green arrows), off-centre structures appear blurred due to the degradation caused by the parallax error (red arrows). The simulated and reconstructed data was provided by Hancong Xu².

This degradation of spatial resolution is caused by the parallax error and appears as inhomogeneous spatial resolution of the cortex surface in patient images [15]. Studies concerning the outer cortex regions are affected the most by this circumstance: this includes studies using labelled neurotransmitters or neuroreceptor ligands, that have shown to be present in different layers of the cortex [100]. One example are PET studies using ^{18}F -Flumazenil to find epileptogenic foci: γ -aminobutyric acid A (GABA_A) is a principal neurotransmitter for inhibitory neurons, and alterations in the corresponding receptor density can be associated with various diseases of the central nervous system. ^{18}F -Flumazenil is a radioligand that specifically binds to the benzodiazepine binding site of GABA_A receptors. A decreased level of GABA_A receptors can therefore be detected by the reduced binding of Flumazenil, which is believed to be an indicator for the presence of an epileptogenic focus [101]. Since the thickness of the human cerebral cortex is partially only 2 mm to 3 mm, a high spatial resolution of

¹The phantom is taken from *BrainWeb*, a simulated brain database [99].

²Institut für Neurowissenschaften und Medizin (INM-4), Forschungszentrum Jülich, <https://www.fz-juelich.de/inm/inm-4/>

the PET image in these regions would be highly beneficial for receptor-related studies. Additionally, studies on the glucose metabolism with Fluorodeoxyglucose (FDG), e.g. progression of Alzheimer's disease and other neurodegenerative disorders investigated with amino acid-based tracers, would benefit from the improved spatial resolution in the cortex: since quantitative images of the glucose metabolism are used as a measure for synaptic functionality, these studies require a homogeneous spatial resolution to accurately calculate the tracer uptake in each region of interest. Due to the parallax error, the resolution is degraded in the cortex area, so that the partial volume effect lowers the quantitative accuracy [102].

As pointed out in section 2.2.3, different approaches exist to minimise the parallax error in PET scanners. Except for the Siemens HRRT scanner, DOI concepts have only been implemented in preclinical scanners so far, e.g. in the Bruker Albira scanner [103] or MADPET-II detector (two scintillator layers with individual readout) [104]. However, four prototypes for a brain dedicated PET scanner with DOI encoding were developed and should be mentioned in this context. The jPET-D4 made use of four crystal layers, which contained different amounts of cerium doping and reflector arrangements ([105], [106]). A device developed by Hamamatsu photonics acquired DOI information from four crystal layers, too, but with independent front end readout of each crystal layer ([107]). Another detector concept with DOI encoding using four staggered crystal layers has been presented in 2007 [108], but feasibility for a complete PET scanner has only been tested for a preclinical setting. The remaining two brain dedicated PET prototypes, MINDView and Oncovision CareMiBrain, were based on monolithic crystals and SiPM readout [109]–[111]. The crystal sizes and achieved spatial resolutions are listed in table 3.1. There exist further detector concepts, which have not yet been assembled as complete brain PET scanners or where unconventional scanner geometries were envisioned. A review of existing approaches can be found in [11] and [112].

An important aspect of the next generation BrainPET detector will therefore be its envisioned DOI capability, which is not yet state-of-the-art technology for commercially available human brain PET scanners.

Spatial Resolution

To ensure a high spatial resolution, a crystal pitch of 2 mm is envisioned (instead of 2.5 mm in the current BrainPET). The common crystal pitch for contemporary commercial human PET devices lies between 3 mm and 4 mm. A lower crystal pitch between 1 mm and 2 mm can only be found in preclinical systems. The spatial resolution (and at the same time energy and time resolution) could be further improved by employing more recent photodetectors with larger sensitive areas. Silicon photomultipliers (SiPMs) have proven to be well suited, since they combine a large sensitive area with excellent timing resolution, as can be observed in the Philips Vereos scanner (307 ps) [32].

Basic properties of currently used PET scanners are compared in table 3.1.

TABLE 3.1: Overview over different clinical and preclinical scanners. The spatial resolution is given as axial \times transverse resolution. Data compiled from the respective sources. *Measured at 1 cm radial distance from the centre FOV. **Tapered monolithic crystals with 40 – 50 mm width.

	Crystal size [mm ³]	Photo-detector	Axial FOV [cm]	Spatial resolution at centre [mm]
Human PET				
Philips Vereos (PET/CT) [34]	4.0 \times 4.0 \times 19.0	dSiPM	16.40	3.99 \times 3.99
Siemens Biograph Vision (PET/CT) [6]	3.2 \times 3.2 \times 20.0	SiPM	26.10	3.50 \times 3.50 – 3.60*
Siemens ECAT HRRT (PET) [14]	2.1 \times 2.1 \times 7.5	PMT	25.20	2.60 \times 2.40
3T PET/MRI				
GE Signa [113]	4.0 \times 5.3 \times 25.0	SiPM	25.00	3.55 \times 1.78 – 1.84*
Siemens mMR [114]	4.0 \times 4.0 \times 20.0	APD	25.80	4.10 \times 4.00*
Siemens MR-BrainPET [15]	2.5 \times 2.5 \times 20.0	APD	19.20	3.00 \times 3.00
Small animal				
Siemens Inveon (PET) [115]	1.51 \times 1.51 \times 10.0	PSPMT	12.70	1.15 \times 1.46 – 1.49
Mediso Nanoscan PET 122S/MR [116]	1.12 \times 1.12 \times 13.0	APD	10.00	0.70 \times 0.70
Bruker Albira (PET) [117]	40 – 50 \times 40 – 50 \times 10.0**	MA-PMT	8.00	0.70 – 1.20 \times 0.70 – 1.20
Brain dedicated prototypes				
jPET-D4 (PET) [105], [106]	2.9 \times 2.9 \times 7.5	PSPMT	26.00	3.00 \times 3.00
Hamamatsu Photonics (PET) [107]	1.2 \times 1.2 \times 3.0 – 8.0	MPPC	20.16	1.80 – 2.10 \times 1.80 – 2.10
Oncovision CareMiBrain (PET) [109]	50 \times 50 \times 15.0	SiPM	15.00	1.94 \times 1.93 – 2.34*
MINDView (3T PET/MRI) [111]	50 \times 50 \times 20.0	SiPM	15.20	0.70 – 1.70 \times 0.70 – 1.70
Next generation				
HVF BrainPET (7T PET/MRI)	2.0 \times 2.0 \times 8.0 – 10.0	dSiPM	24.40 – 26.30	TBA

Sensitivity

The peak point source sensitivity of the Siemens BrainPET is 6% [15]. Further improvements can be achieved by enhancing the solid angle coverage, i.e. minimising unused space between the detector blocks. Moreover, thicker detector blocks can be used to enhance the number of detected events N . An increase of the sensitivity can be translated into an improvement of the signal to noise ratio (SNR), since $\text{SNR} \propto \sqrt{N}$ [10].

UHF MRI Compatibility

The feasibility of combining PET and ultra-high field MRI has already been demonstrated, e.g. in preclinical scanners [118] and with a Siemens MR-BrainPET insert located at the research centre Jülich [119]. However, there is no hybrid device using ultra-high field (UHF) MRI commercially available, that is suited for application in humans, although the clinical benefit from higher field strengths is apparent: working with 7 T or more enables spatial resolution in the sub-millimetre range, enhances the intrinsic SNR and increases the contrast in T1 and T2* weighted images (since the underlying physical processes are field strength dependent) [120].

Moreover, MR spectroscopy (MRS) at 7 T has proven to enhance the differentiation of brain tumours and metastases, and enables the measurement of so-called X-Nuclei like ^{23}Na or ^{31}P with sufficiently high SNR and spatial resolution [13]. These advantages of UHF MRI are beneficial for acquiring both structural and functional information.

3.2 Detector Concept in this Work

The detector for the next-generation brain PET will comprise new technologies to achieve higher sensitivity and spatial resolution.

To acquire DOI information, a staggered three layer geometry has been chosen, similar to the one suggested in a four layer concept for a preclinical scanner [108]. As shown in figure 3.2, the bottom layer is planned to consist of LSO pixels with dimensions $2 \times 2 \times 10 \text{ mm}^3$, each wrapped with a specular reflector (3M ESR vikuiti [121]) in the lateral dimensions. The second layer will have shorter crystals ($2 \times 2 \times 9 \text{ mm}^3$ planned), and will be shifted by half a crystal pitch (1 mm) in x direction. The top layer with crystals of the planned dimensions $2 \times 2 \times 8 \text{ mm}^3$ will be shifted by half a crystal pitch in y direction, and one crystal pitch in x direction. The complete block will be covered with a reflector as well, to prevent scintillation light from escaping the block (see figure 3.3).

The optimum size of the crystal array (16×16 , 20×20 or 24×24 crystals in the bottom layer) has to be determined— this is the topic of chapter 4.




FIGURE 3.2: Crystal geometry for the next-generation brain PET: three staggered layers are supposed to yield DOI information, being composed of $2 \times 2 \times 8\text{--}10\text{ mm}^3$ LSO crystal pixels. Finding the optimum crystal array size is one of the objectives of this thesis.



FIGURE 3.3: A crystal block prototype, consisting of 24×24 crystals in the bottom layer. The whole block is covered with a specular reflector, and a lightguide is glued to the bottom crystal array.

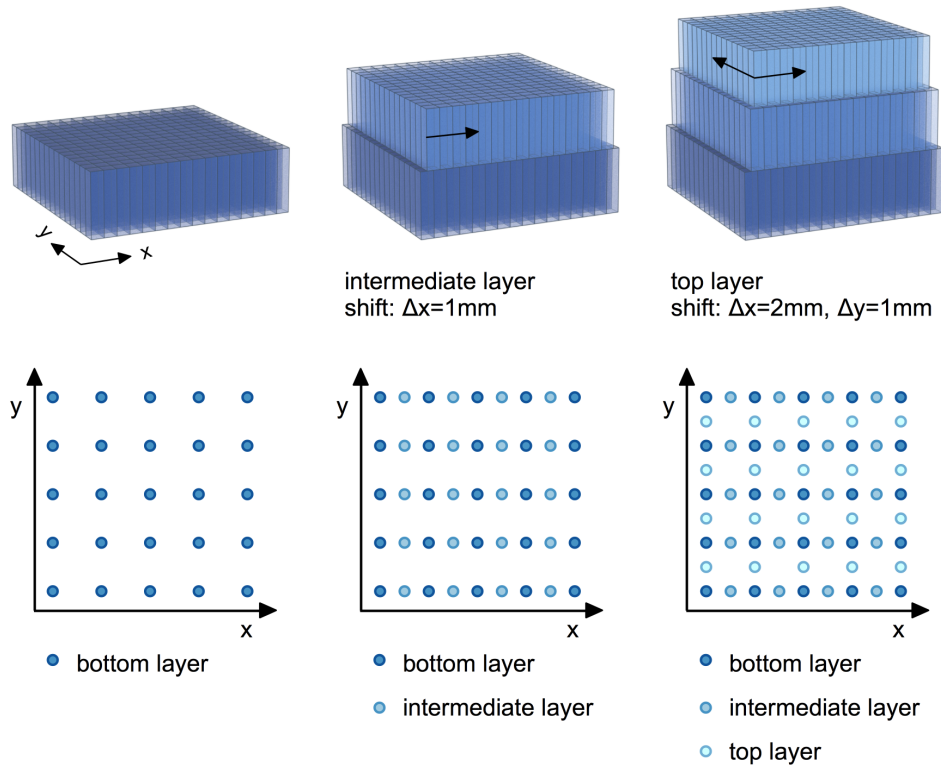


FIGURE 3.4: Top row: The three crystal arrays are shifted relative to each other. The intermediate layer is only shifted by half a crystal pitch in x direction, while the top layer is shifted in both directions (one crystal pitch x and half crystal pitch y , relative to the bottom layer). Bottom row: A schematic drawing of the spot pattern, which is expected to be found in the flood map (only a small section is shown).

The different layer thicknesses are supposed to approximately balance the number of detected events in the individual layers and therefore to ensure good point separability in the flood histogram. The expected pattern of the flood histogram is shown in figure 3.4: a regular spot pattern is expected for each layer. Due to the relative shifts of the three layers, the individual spot patterns will be shifted likewise, so that each crystal in the block can be identified by a single distinctive spot in the flood histogram. The overall thickness of the scintillator block amounts to 27 mm, which is expected to enhance the sensitivity compared to other scanners (e.g. 20 mm for the current BrainPET). As indicated in section 2.3, digital SiPMs will be used to account for MRI compatibility; with dedicated trigger and validation schemes, 250 ps to 400 ps timing resolution are expected. Optical coupling between crystal stack and dSiPMs will be ensured with a lightguide made from float glass (Borofloat, $n = 1.47$ [122]).

3.3 Objective of this Thesis

Despite having a detector concept for the new brain PET, the final dimensions of the detector blocks are not yet decided.

Employing **smaller or larger detector blocks** is expected to have significant influence on sensitivity, furthermore the **thickness of the lightguide** has to be optimised. Thicker lightguides allow broader light spread over the DPC tile and could thus enhance centre of gravity position calculations. Too large light spread is not desired, since it could increase the dead time and also lead to signal cut offs, when the signal of pixels far away is not sufficiently high to pass the trigger and validation thresholds [123]. Therefore, an optimum lightguide thickness has to be found in this thesis.

Closely related is the challenge of **decoding the flood histogram** of a three layered detector stack: it requires a sophisticated positioning algorithm. The feasibility of applying maximum likelihood estimation maximisation algorithms (MLEM) to a single layer detector has been proven before [58]. In this thesis, adaptation for this special scintillator geometry will be attempted. Moreover its efficiency will be evaluated and compared to the commonly used centre of gravity algorithm with respect to the number of correctly identified scintillation events.

Not all of these questions can be answered with evaluations of measured data. Simulations are essential, especially since they allow for the investigation of quantities that are not accessible in a measurement, as for instance a ground truth for accuracy evaluation. In this thesis, a combination of both methods is envisioned, including **extension of the simulation software (and hardware)**, if necessary for the respective research question.

One aspect that should be kept in mind is a possible commercialisation of this brain PET scanner. In that sense, the optimum detector geometry and readout have to be realised at reasonable production costs.

Chapter 4

Optimisation of Detector Block Configuration

For the next-generation brain PET, three different crystal array sizes were considered: the dimensions ranged from 16×16 crystals up to 24×24 crystals in the bottom layer, but all lead to a similar PET insert configuration. Table 4.1 contains detailed information on the three different options. All crystals will have a pitch of 2 mm and will be made of LSO, covered with a specular reflector (3M ESR vikuiti). The layer which is coupled to the lightguide will be 10 mm thick, the intermediate layer 9 mm and the top layer 8 mm.

TABLE 4.1: Overview of different possible crystal array sizes, which could be used for the new scanner. For comparison, the dimensions of the currently used Siemens BrainPET are given.

	Layer	Crystal array size	Scanner layout	Axial FOV [cm]
Option I	Bottom (10 mm)	16×16	34 Cassettes	26.30
	Intermediate (9 mm)	15×16	8 Blocks each	
	Top (8 mm)	14×15	272 Blocks in total	
Option II	Bottom (10 mm)	20×20	28 Cassettes	24.50
	Intermediate (9 mm)	19×20	6 Blocks each	
	Top (8 mm)	18×19	168 Blocks in total	
Option III	Bottom (10 mm)	24×24	24 Cassettes	24.40
	Intermediate (9 mm)	23×24	5 Blocks each	
	Top (8 mm)	22×23	120 Blocks in total	
BrainPET	Single layer (20 mm)	12×12	32 Cassettes	19.25
			6 Blocks each	
			192 Blocks in total	

A common measure for the performance of a PET scanner is the sensitivity. It is defined as the ratio of registered true coincidences and the source activity, and therefore gives an estimate of how many annihilation photons can be detected. A standardised procedure to measure the sensitivity was defined by NEMA (National Electrical Manufacturers Association): the NEMA standard suggests to place a 70 cm long line source filled with ^{18}F solution in the transaxial centre of the scanner, and to assess the rate of true coincidence events at a given source activity. The measurement is repeated with five aluminium sleeves of different diameters, which are slid over the line source, so that six measurements with different amounts of attenuation are acquired. The geometric sensitivity without any attenuation can then be extrapolated [124].

The three possible detector block configurations differed in the arrangement of the crystals (i.e. the size of the arrays), and therefore led to different geometric coverages. These coverages were expected to have a major influence on the system sensitivity, which should be maximised for the new brain PET scanner. To determine the optimum detector block configuration, sensitivity profiles were simulated using GATE. This simulation allowed for investigating the sensitivity of the different crystal geometries, while neglecting the influence of the readout electronics and attenuation. In this way, the pure influence of crystal array size on the ratio of annihilation events and true coincidences could be studied.

4.1 Simulation of Sensitivity in GATE

Methods

Figure 4.1 shows the implementation of the simulation with scanner option III in GATE. A line source was placed along the z axis of the air filled scanner, emitting two 511 keV annihilation photons "back-to-back" per event, while the source length matched the axial scanner FOV¹. A source activity of 1 MBq was set for a simulation time of 30 s. To model the photon interactions inside the LSO crystals, the physics list `emstandard_opt3` was selected, because it contained physics processes which are relevant for simulations of electromagnetic processes without a magnetic field².

Four digitiser modules were set to implement the readout electronics of each block.

1. Readout: The scanner geometry was implemented as "cylindrical PET", which is a predefined framework in GATE to model the scanner readout; the readout level was set to block level. This means that all individual hits of an event inside a crystal block

¹Of course it would also have been possible to simulate an ^{18}F source in water, but this was not necessary since the positron range was not relevant for this simulation. Moreover, the source would have to be defined in such a way that the positrons annihilate inside the source volume but also encounter the least possible attenuation. Skipping the positron annihilation process also has the positive effect of reduced computation time of the simulation.

²<https://geant4.web.cern.ch/node/146#opt3>




FIGURE 4.1: A line source is aligned to the z axis of the scanner, emitting pairs of annihilation photons. As shown for scanner option III, the detector block signals of each cassette are calculated independently.

were combined to a single, while its position corresponded to the energy-weighted centroid and its energy was the sum of all hits.

2. Energy blurring: After combining the hits to singles, the energy of each single was sampled from a Gaussian distribution. Since the scintillation process was not explicitly simulated, this technique was applied to approximate the influence of different light yields of crystals, the light collection efficiency, intrinsic energy resolution and the quantum efficiency of the sensor. The summed energy of the hits was used as the mean of the Gaussian distribution during the sampling process, while the standard deviation was calculated from the specified energy resolution (12% at 511 keV).

3. Energy window: Similar to a real measurement, an energy window was set to filter scattered photons: singles with energies between 380 keV and 634 keV (i.e. twice the FWHM of the full energy peak at the set energy resolution of 12%) were stored temporarily and then sorted by a coincidence sorter.

4. Coincidence sorter: As long as the two photons were not originating from adjacent detector cassettes (`minSectorDifference 1`), singles were combined to a coincidence, when they arrived within a time window of 5 ns. In case of multiple coincidences, the hits of all involved photons were discarded (`killAllIfMultipleGoods`).

GATE offers many additional parameters for modelling, e.g. dead time, noise or the electronic signal processing. These parameters were not used, because the corresponding system parameters were not yet known and also not relevant for the optimisation of the detector geometry.

True coincidence events can be selected by making use of the ground truth information, which is provided by the `eventIDs` of the coincident photons: if both `eventIDs` match, the photons originate from the same event, i.e. the same annihilation,

and are thus a true coincidence event. Vice versa, if the eventIDs differ, the coincident photons are regarded as a random event and discarded. The true coincidences were sorted into a histogram according to their point of creation, in order to evaluate the simulated data similar to the NEMA standard. In a measurement, counts from oblique LORs are assigned to the transaxial slice, where the LOR intersects the scanner axis (single slice rebinning). Usually, the bin size is equivalent to the voxel size in axial direction (e.g. 1.25 mm for the Siemens BrainPET [15]), but to give comparable values for all scanner designs, 1.0 mm bin size was used. The resulting histogram corresponded to an axial sensitivity profile.

Results

In comparison to measurements, this simulation overestimated the absolute scanner sensitivity, e.g. a peak sensitivity of 12.6% was acquired in case of the Siemens BrainPET instead of 6.0%, as measured by *Herzog et al.* [15]. This effect can be mainly deduced to the ideal source that was simulated: the photon pairs emitted with an energy of 511 keV did not pass any attenuating source volume, leading to an overrepresented full energy peak. In contrast, positrons from an ^{18}F source in water would first have to annihilate and subsequently leave the source volume, which enhances the amount of Compton scattered events. Moreover, the simulation did not take the patient bed, detector housing, scanner gantry, etc. into consideration, and therefore underestimated Compton scattering. However, an investigation of the relative improvement of sensitivity, which might be achieved merely with a modified crystal geometry, was the focus of this study. Therefore, all effects apart from the geometry were assumed to remain the same, and all sensitivity profiles were scaled by the factor 1/2.1 to match the peak of the simulated sensitivity profile of the Siemens BrainPET with the measured point source sensitivity.

Figure 4.2 shows the simulated axial sensitivity profiles of the three scanner designs, and – for comparison– the result of the same simulation with the Siemens BrainPET geometry. All datasets showed the typical triangular shape of the sensitivity profile [125], indicating that the most registered true coincidences originated from the axial centre of the scanner. The ‘overall sensitivity’, i.e. the ratio of registered true coincidences and the number of emitted photons, is printed in table 4.2. Moreover, the peak sensitivity is given for all scanner designs.

Discussion

The sensitivity of the Siemens BrainPET had been determined in two measurements: *Herzog et al.* measured a sensitivity of 6.0%, using a point source in the centre of the scanner and single slice rebinning [15]. *Kolb et al.* conducted a measurement according to the NEMA procedure, but used 3D slice rebinning, which resulted in a peak sensitivity of 7.2% in the centre FOV [126]. The simulation results were scaled

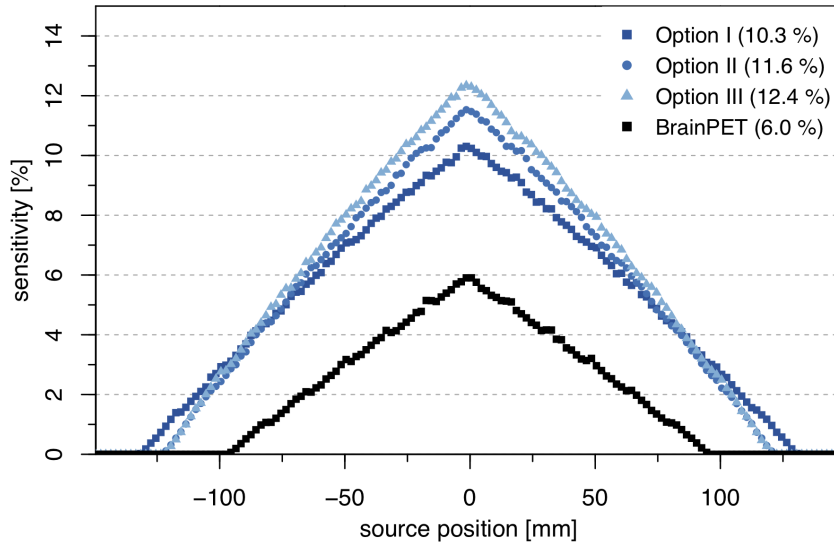


FIGURE 4.2: Axial sensitivity profiles for the three new scanner design options and the Siemens BrainPET for comparison. The increase in sensitivity due to less dead space and a longer axial FOV is clearly visible. For each scanner design, the peak sensitivity is given in the legend.

TABLE 4.2: Ratio of detected true singles/coincidences and emitted photons for each scanner design. A change in geometry influences the sensitivity significantly.

Scanner	Overall sensitivity		Peak sensitivity
	Singles	Coincidences	Coincidences
BrainPET	18.7 %	3.0 %	6.0 %
Option I	27.6 %	5.5 %	10.3 %
Option II	29.2 %	6.1 %	11.6 %
Option III	30.7 %	6.6 %	12.4 %

to the measurement of Herzog et al., because the point source measurement yields a more accurate estimate for the peak sensitivity.

The overall sensitivity of the BrainPET and the newly developed scanner configurations differed significantly, ranging from 3.0 % to 6.6 %. This was reducible to the thicker crystals (20 mm vs. 27 mm), the change of the axial FOV (192.5 mm vs. 244 mm or more) and the gaps between the modules of each cassette (BrainPET: 2.5 mm, new scanners: 1.0 mm). The elongated axial FOV allowed the detection of more LORs, and the tightly packed crystal arrays diminished the dead space in each cassette.

From both figure 4.2 and table 4.2 it seems that option III would be the best choice for a new PET scanner, judging by the overall sensitivity. This corresponds to the option with fewest, but largest detector blocks. The increased sensitivity was not only apparent for the detected singles, but especially for the detected coincidences, which were nearly doubled for option III. Option I showed the worst performance of the new designs, although this scanner has the largest axial FOV. At the same

time, the detector blocks are smaller, so that more blocks are necessary to cover the scanner volume. This leads to many gaps in-between, creating more unused space and eventually reducing the sensitivity. Option II has intermediately sized blocks, but the same axial FOV as option III, so that the corresponding sensitivity lay between the other two options (6.1 % vs. 5.5 % and 6.6 %). It seems that the block size (and resulting little dead space) has the major influence on sensitivity.

Furthermore, the detector block thickness was expected to contribute to the higher sensitivity. In the new design options, the scintillator thickness was increased to 27 mm, compared to the Siemens BrainPET with 20 mm scintillator thickness. Within a 20 mm thick LSO block, approximately 82 % of the incoming photons are stopped (linear attenuation coefficient of 0.8658 cm^{-1}) [29]. This value corresponds to the intrinsic detection efficiency of a singles detector. For coincidence detection, two single detectors register independent events, so that the coincidence efficiency is the product of both efficiencies (67 %). A block thickness of 27 mm LSO increases the singles detection efficiency to 90 % and the one for coincidences to 81 %, so that an increase of sensitivity was also related to the thicker detector blocks.

4.2 “Is it really worth it?”- A cost-benefit assessment

Although option III would yield the highest sensitivity, it might not be first choice. Each scintillator crystal has to be cut and covered with reflecting foil individually, so that the price usually increases with the number of crystals. Table 4.3 gives an overview over the number of crystals and a price comparison for the scanner designs (quotes from different companies).

TABLE 4.3: The different scanner designs with the respective scintillator volumes and number of crystals. Option IV was additionally considered, being a presumably cheaper alternative.

Scanner layout	Scintillator volume		Number of crystals		Price
	Block [cm^3]	Scanner [cm^3]	Block	Scanner	
BrainPET	18.0	3456.0	144	27648	429,300\$
Option I	25.6	6963.2	706	192032	800,500\$
Option II	40.6	6824.8	1122	188496	702,500\$
Option III	59.1	7092.5	1634	196080	729,600\$
Option IV	53.2	6387.4	1105	132600	445,600\$

The three-layered detector concepts have comparable scintillator volumes, between 6800 cm^3 and 7000 cm^3 in total. However, the number of crystals varies by more than 5000 crystals, making option II the cheapest design due to the lowest number of crystals. Since the scintillator material makes up half of the costs for

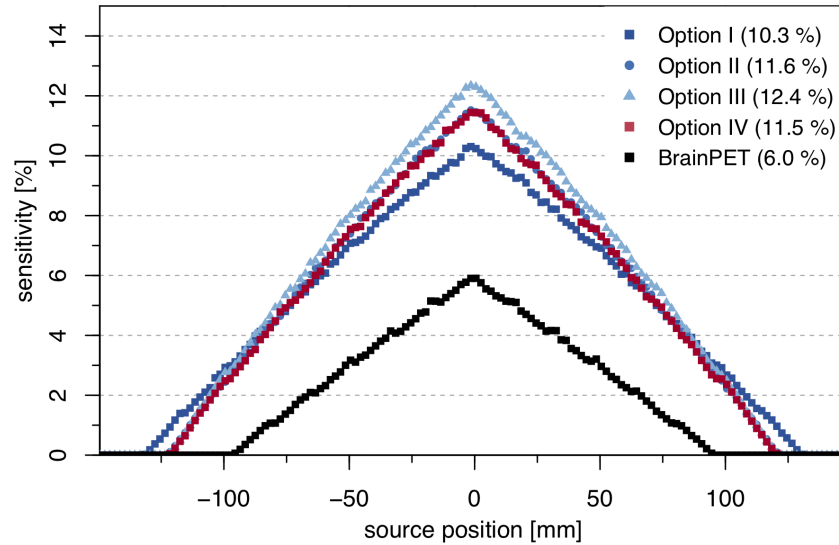


FIGURE 4.3: Axial sensitivity profiles for the three new scanner design options, and the two layer option (IV) for comparison. The sensitivity profile is equivalent to the one of option II with intermediate sized blocks. For each scanner design, the peak sensitivity is given in the legend.

this PET scanner, the question arose whether three crystal layers are really required. Moreover it has been described previously that the benefit of more than two crystal layers is limited [127].

Methods

An additional option (IV) was included, as already mentioned in table 4.3: only two crystal layers are stacked with a relative shift of half a crystal pitch in both x and y direction. The layer thickness is increased to 13 mm in the bottom and 11 mm in the top layer, resulting in a total stack height of 24 mm. This scanner geometry has significantly less crystals (132600 instead of ~ 192000), which is reflected in lower production costs.

To investigate the achievable sensitivity of scanner option IV, the simulation was repeated for the additional geometry. The scanner diameter of option IV is the same as for option III.

Results

Figure 4.3 shows the results for the two layer option (IV) along with the previous results for options I-III. Although the detector blocks of the two-layer option are not as high as the others (24 mm instead of 27 mm), the sensitivity was very similar to that of option II with the intermediate sized blocks. Since the modules are larger, they contain 24×24 crystals in the bottom layer, like in option III. The higher fill factor of the system balanced the lower stack height, leading to a result similar to that of option II.

Discussion

The results obtained so far excluded option I from the possible scanner designs. The gaps between the small detector blocks and between the cassettes³ caused considerable unused space, therefore this option had the lowest sensitivity of all options, but was among the most expensive ones. The large detector blocks of option III showed a high axial coverage and thus high sensitivity, but their feasibility depends on the final budget for the scanner. Cheaper alternatives could be options II and IV, where the main difference is the number of crystal layers. Since both the axial sensitivity profiles and the peak sensitivities were basically the same (11.6% vs. 11.5%), the lower price would speak in favour of the two layer design. However, it had to be clarified, whether the three layer design would be superior when it comes to the reduction of DOI uncertainty.

4.3 Impact of Capability to Detect DOI Information

A commonly used phantom to measure spatial resolution in reconstructed images is the hot rod Derenzo phantom: it consists of a cylinder of water equivalent material (PMMA) and several rods of different sizes, which are embedded in triangular patterns inside the cylinder. For PET, these rods are filled with aqueous ¹⁸F solution, so that the rod diameters can be determined in a reconstructed image to estimate the smallest resolvable object size. This phantom was slightly modified to study the impact of reduced DOI uncertainty on the reconstructed image.

Methods

Different versions of the hot rod phantom exist, as it has been used in preclinical and clinical applications since the 1970s [128]. A similar phantom with larger dimensions was set up in GATE (see figure 4.4): the total diameter of the phantom was extended to 20 cm, so that the phantom size was similar to the typical head size. This was important for investigating the influence of the parallax error in off-centre regions (several cm distance from the scanner axis). The phantom length was set to 4 cm, and the rod diameters were 4.8 mm, 4.0 mm, 3.2 mm, 2.4 mm, 1.6 mm and 1.2 mm. As shown in figure 4.4, the rods were equidistantly spaced in each of the six sections, while the rod distance corresponded to twice the rod diameter. The rods contained a total activity of 32.5 MBq, distributed to the rods according to their volume, and implemented as a back-to-back photon source. Once again, the annihilation step was skipped to reduce the simulation time. Positron range, acollinearity of the annihilation photons and their attenuation inside the phantom were not simulated, in order to investigate only the influence of the scintillator arrays and the corresponding

³The distance between two detector blocks inside a cassette can be decreased to 1 mm. This flexibility is not given for the whole cassettes, because each cassette requires an individual RF shield. Shielding all PET detector blocks as a whole is not an option, because the eddy currents would become too strong.




FIGURE 4.4: The modified hot rod phantom inside scanner option III.

DOI correction in the reconstructed image. The readout parameters for the scanners remained unchanged, and the simulated time was set to 600 s.

Before the simulated data could be reconstructed, a normalisation measurement was required to balance different detector block response functions (see section 2.1.4) [27]. For this purpose, a cylindrical source was placed inside the scanner, covering the whole axial and transaxial FOV: the diameter was 183 mm, and the length of the cylinder was adapted individually to match the scanner length. Back-to-back photons of 511 keV were emitted from the full cylinder volume, in order to irradiate all detector blocks equally. All other parameters (readout level, energy window, coincidence sorter etc.) were the same as in the previously described simulations. An activity of 100 MBq was simulated over a time interval of 540 s, in order to acquire approximately 18×10^9 coincidence events. To handle the long simulation time and large data output, the supercomputer JURECA [129] was used. The simulation was split into 720 jobs, running 24 jobs on each node. A typical simulation of normalisation data required between 6000 and 8000 core-hours.

A difficult aspect of this study was to create fairly comparable data sets. Compared to the current BrainPET, the three-layer option has many more possible LORs, simply because the number of crystals is seven times higher. Therefore, larger statistics had to be simulated for the normalisation data, since a certain number of events has to be simulated for each LOR. In PRESTO, the same strategy was applied to calculate the system matrix from a virtual layer. Moreover, the same number of iterations was used for all three designs. To guarantee comparability in the simulation, the cylinder source had the same diameter and the same activity, and the length has been adapted

individually to the scanner length. The reconstruction of the simulated data was conducted by Jürgen Scheins⁴.

Results

Figure 4.5 shows reconstructed images of the hot rod phantom inside the three scanner options. Each image shows the sum of all planes in axial direction, and has been normalised individually to the sum of all voxel values of the data set. It is striking how much the intensity distribution varied between the BrainPET and the two new designs. In case of the BrainPET, the rods in the centre FOV were clearly visible, while the intensity (and therefore the separability) of the rods decreased towards the peripheral FOV.

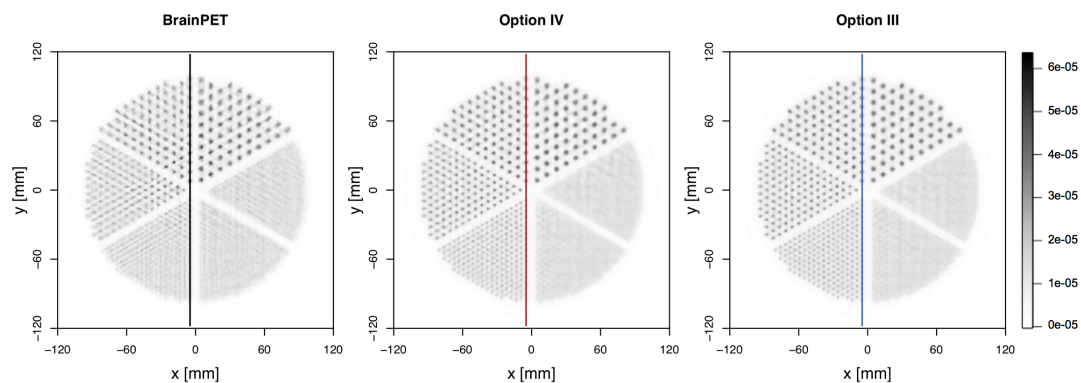


FIGURE 4.5: Reconstructed image of the hot rod phantom inside the three scanners (150 iterations, isotropic voxels of 1.25 mm^3 , one subset). Each image shows the sum of all planes in axial direction. The coloured lines indicate the position of the line profiles in figure 4.6.

The intensities were distributed more homogeneously in the new designs. This effect can be observed even more clearly in figure 4.6, which contains line profiles through the compartments with 2.4 mm and 4.0 mm rods: the peak to valley ratio was very high around the centre FOV, but for the BrainPET it decreased to less than half of the initial value at the border, making the rods appear more and more blurred. The peaks in the black line profile decreased and broadened accordingly.

⁴Institut für Neurowissenschaften und Medizin (INM-4), Forschungszentrum Jülich, <https://www.fz-juelich.de/inm/inm-4/>

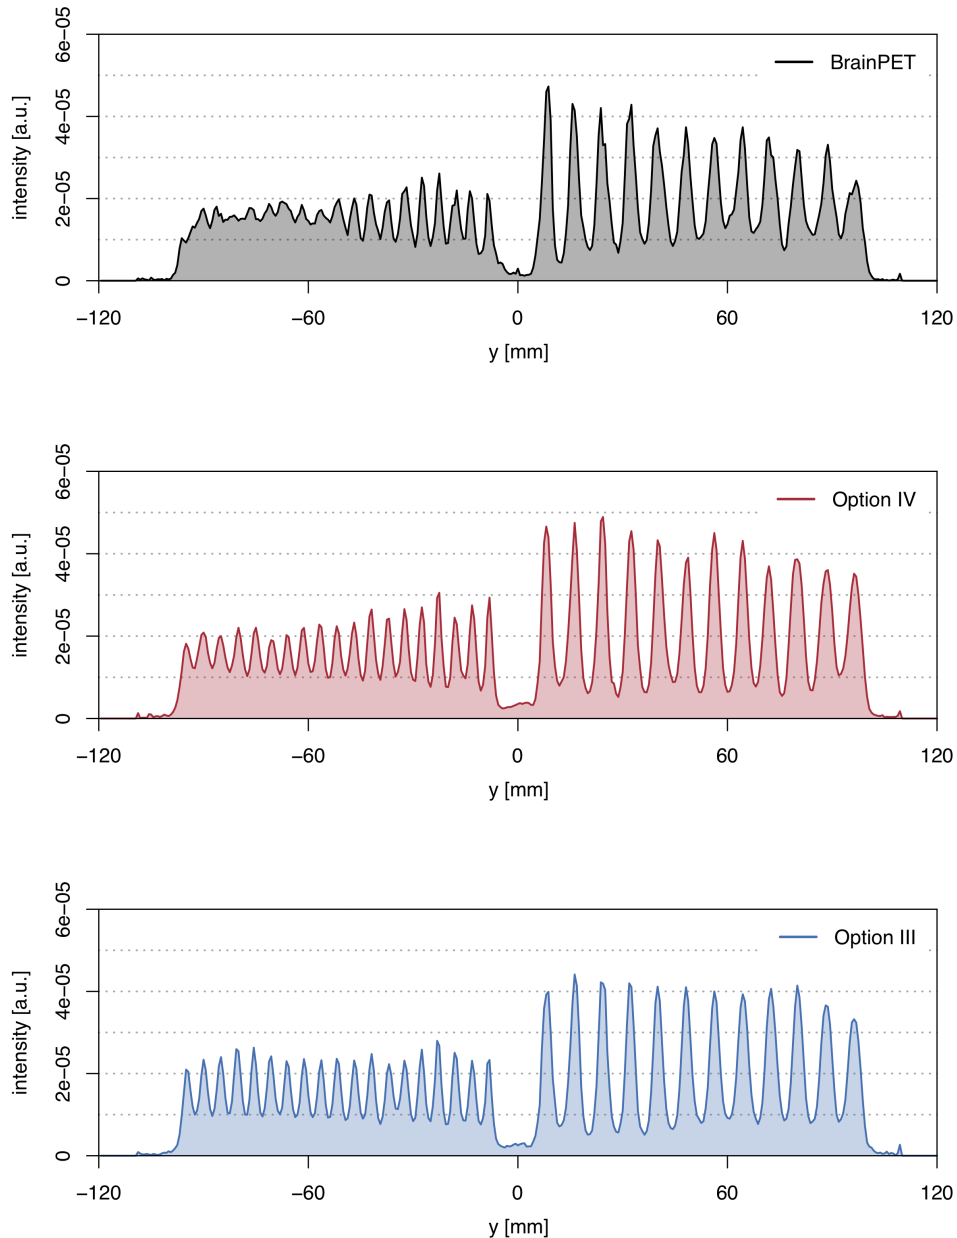


FIGURE 4.6: Line profiles through two compartments of the phantom, as indicated in figure 4.5, for the BrainPET (top), the two layer scanner design (middle) and the three layer option (bottom). The decrease of peak to valley ratio is strongly visible for the BrainPET, while the DOI consideration in option IV and III homogenises the peak to valley ratio over the FOV.

The explicit peak to valley ratios for the 2.4 mm rods are shown in figure 4.7. The rods of 2.4 mm size could be separated inside the whole FOV, when the two or three crystal layer options were used. In case of the BrainPET, the peak to valley ratio converged to 1, which corresponds to no rod separability. Option IV, which is the design with two crystal layers, yielded better separability at the peripheral FOV (6 cm to 10 cm from the centre, red line), where the impact of DOI is becoming stronger. Using three crystal layers for DOI correction (option III, blue line) improved the peak

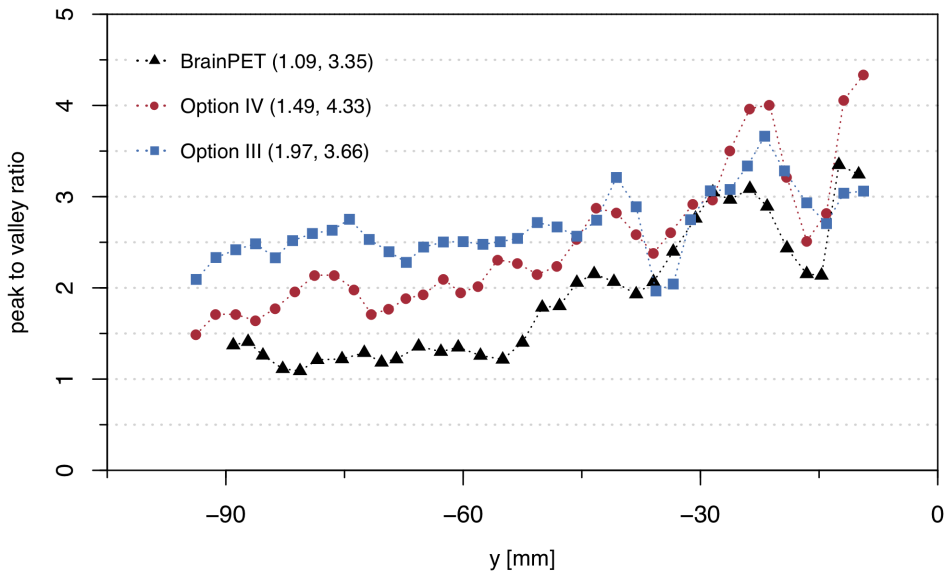


FIGURE 4.7: Peak to valley ratios of the line profiles for the 2.4 mm rods (left side in figure 4.6). The ratio has been calculated for all peaks and their neighbouring valleys, which are located between the two outermost rods. The legend contains the minimum and maximum peak to valley ratios for each scanner option.

to valley ratio even more, compared to option IV. This effect was also reflected in the higher mean value and smaller variance.

The rods of 2.4 mm could be resolved by all scanners, but the signal of the 1.6 mm rods was blurred when using the current BrainPET. Most of these rods were visible in three layer option, even when a single image slice was considered, as shown in figure 4.8. In this image slice, 26 of the 28 rods could be identified.




FIGURE 4.8: Line profiles for the 1.6 mm and 4.8 mm rods in the scanner with three crystal layers. Except for the two outermost rods, the 1.6 mm rods could be identified.

Discussion

The line profiles in figure 4.5 clearly showed the advantage of three crystal layers over two layers regarding DOI correction. The difference was best visible in the peripheral FOV (between 6 and 10 cm from the centre).

However, the expected spatial resolution of the scanner designs cannot be derived directly from the simulated phantom, since it will be further degraded by photon acollinearity and the positron range. Assuming a mean acollinearity of 0.5° (FWHM), the degradation can be estimated by the formula $0.0044 \cdot R = 0.814 \text{ mm}$ [8], where R is the scanner radius. The FWHM of blurring due to the positron range amounts to $s = 0.54 \text{ mm}$ for ^{18}F [8]. According to the estimation of spatial resolution by *Moses* [8], both factors contribute in quadrature along with other degrading factors (such as detector element width d and the parallax error), which have already been included in the simulation. By making use of *Moses'* formula, the spatial resolution of the Siemens BrainPET can be estimated:

$$G = 1.25 \times \sqrt{\left(\frac{d}{2}\right)^2 + s^2 + (0.0044 \cdot R)^2 + b^2 + \frac{(12.5 \cdot r)^2}{r^2 + R^2}} \quad , \quad (4.1)$$

where b accounts for the crystal decoding error, which *Moses* estimates to be $b = d/3$. This formula yielded a spatial resolution of $G = 2.24 \text{ mm}$ in the centre of the FOV, which degraded to 6.28 mm at a radial offset of 75 mm , and resolutions of 2.0 mm and 6.2 mm respectively, when photon acollinearity was not considered. In comparison to the line profiles in figure 4.6, where back-to-back photons without acollinearity have been simulated, the spatial resolution in the centre was similar, since the 2.4 mm rods were clearly resolved. At radial offsets, the simulations of the Derenzo phantom suggested a better spatial resolution compared to *Moses'* formula. The line profiles of the simulation were in better agreement with the spatial resolution, that has been measured for the Siemens BrainPET by two groups.

Herzog et al. determined about 3.0 mm in the centre of the scanner and 4.5 mm radial resolution at a radial offset of 75 mm [15]. *Kolb et al.* measured 1.8 mm to 2.9 mm spatial resolution at 10 mm radial offset, which degraded to 3.6 mm to 6.0 mm at a position of 100 mm radial offset [126]. The line profile from the simulated phantom in figure 4.6 showed that the 2.4 mm rods could be resolved in the centre of the Siemens BrainPET, approximately up to a radial offset of 50 mm . Rods of 4.0 mm diameter could be resolved in the whole FOV. The deviation from the measured spatial resolution was caused by the further degradation due to positron range and photon acollinearity.

Applying *Moses'* formula to the results of scanner option III yielded a theoretical spatial resolution of 1.94 mm and 1.65 mm without consideration of acollinearity. For this scanner option it was assumed that the parallax error could be mostly avoided due to the three staggered crystal layers. The line profile in figure 4.8 showed that

rods of 1.6 mm diameter could still be resolved, although a certain degradation due to the positron range would have to be expected.

The simulation results were in good agreement with the estimation of spatial resolution according to *Moses'* formula in the centre FOV, and with the measurements of *Herzog et al.* and *Kolb et al.* This means that structures between 2.0 mm and 3.0 mm (like the cortex) are expected to be distinguishable, even at large radial offsets. This would be beneficial for studies where signals from the cortex are investigated (neurotransmitters or receptor binding).

Summary

In this chapter, the optimum detector module configuration was determined from three different options varying in size. By means of a GATE simulation, it was studied how the dimensions of the scintillator arrays influence the sensitivity of a whole scanner system. The simulations showed that the scanner option with the largest detector modules (24×24 crystals in the bottom layer) would yield the highest sensitivity (12.4% at peak), since the solid angle coverage of the scanner would be maximised, and reduced gaps between the detector blocks would be achieved. However, manufacturers' quotes revealed that this option would also be a very expensive one, so that a cheaper detector option with two crystal layers was investigated as well. The original idea of using three crystal layers was to gain information on the depth of interaction (DOI), which would be limited to the three crystal layer thicknesses (10 mm, 9 mm and 8 mm). Reducing the stack to two crystal layers (11 mm and 13 mm) would yield less depth resolution, but it was questionable how far this would influence reconstructed images. Therefore, a customised spatial resolution phantom was simulated and reconstructed with PRESTO. The reconstructed images revealed a more homogeneous spatial resolution when using the three crystal layer detector. This effect was especially visible at radial offsets, which affected the human cortex area in the acquired images. Since these areas are of special interest for neuroscientific applications, the three crystal layer detector was still the preferred design, despite higher costs. The spatial resolution of the three crystal layer configuration was estimated to be 1.94 mm in the centre (according to *Moses'* approximation [8]). Therefore, structures of 2 mm to 3 mm at large radial offsets are expected to be resolved.

Chapter 5

Building Test Detectors

The simulation studies described in chapter 4 have provided general concepts for an optimised PET detector module. Still, these hypotheses had to be validated by means of measurements with test detectors.

5.1 Assembly of Prototypes


Three LYSO crystal arrays of dimension $16 \times 16 \times 9$ mm were used for the first test detector. Each crystal array consisted of 8×8 crystals with 2 mm pitch, where the intersection faces of the crystals were wrapped with a specular reflector (3M ESR vikuiti). Only a quarter of the DPC tile was covered with this crystal stack, but due to the sensor's symmetry the results could be extrapolated to the remaining quarters. Experiments with a full-size crystal stack will be described later in section 8.5 (at this stage of the project, only the three 9 mm thick array samples were available).

As has been shown earlier in figure 3.4, the three crystal layers were shifted relatively to each other by half a crystal pitch, i.e. 1 mm. To achieve these relative shifts during the gluing process, dedicated stencils were printed in a 3D printer (see figure 5.1). For each component, i.e. the lightguide and the three crystal arrays, there was a separate stencil. All stencils were stacked and attached to each other during the glueing process. At first a lightguide ($24.0 \times 24.0 \times 1.1$ mm) was glued to the DPC tile, using glue with a refractive index of $n = 1.5$ ¹, while it was aligned to one corner of the sensor (see figure 5.1a). The stencil for the bottom crystal array was attached to the lightguide stencil, so that the crystal array was precisely aligned on the lightguide, as shown in figure 5.1b. In a similar manner, the stencils for the intermediate and top crystal arrays were stacked on top, so that a crystal array was glued to the underlying crystal array with the correct shift of half a crystal pitch. After the glueing process, the stencils were removed.


The final test detector can be seen in figure 5.1c. Since the crystal layers all had the same dimensions, the upper layers were protruding towards the DPC centre (marked with red arrows). After a scintillation event, scintillation photons might escape from the open faces of a crystal array and get detected by the dSiPM— which should be avoided. Therefore, the crystal stack was wrapped with white Teflon tape, so that the

¹Scionix Holland B.V., <https://scionix.nl/>.


majority of photons could not escape the stack. The finished test detector is shown in figure 5.1d. Figure 5.2 shows the position of the crystal arrays on the DPC tile.




(a) Lightguide alignment.



(b) Alignment of the first crystal layer. The 3D printed counterpart to hold the crystal from the front side is not shown.

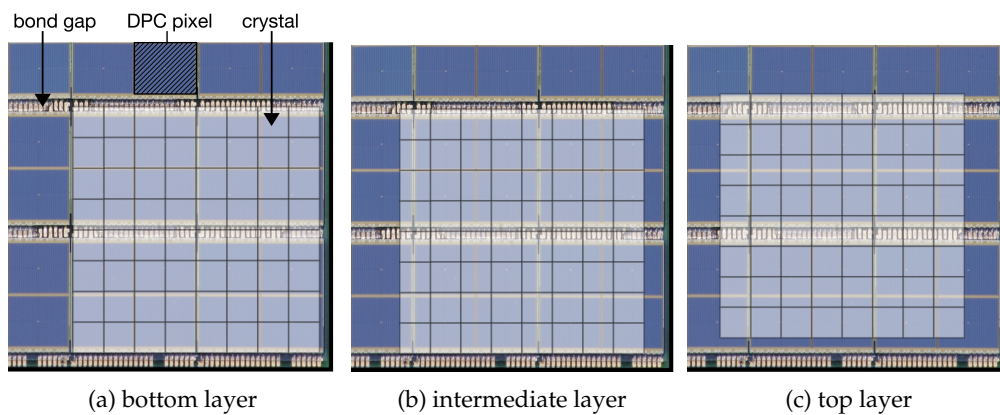


(c) The finished LYSO stack had open faces, where light could escape (red arrows).



(d) The stack was wrapped with Teflon and set on the test stand.

FIGURE 5.1: Lightguide and scintillator arrays were aligned on the DPC tile by means of 3D printed stencils.



(a) bottom layer

(b) intermediate layer

(c) top layer

FIGURE 5.2: Projected alignment of the three crystal layers on the DPC tile: the white semi-transparent squares represent the single crystals within a crystal layer, while the DPC tile is shown in the background photo. The section shows 5×5 pixels (2.5×2.5 dies) of the DPC tile.

5.2 Acquisition of Flood Maps

A dedicated test stand for DPC devices, constructed by Holger Nöldgen and Matthias Streun², was made available for the acquisition of flood maps. In this test stand, the DPC tile was attached to a cooling unit consisting of a Peltier cooler and a copper plate, as partly visible in figure 5.1d. Heat transfer pads between the DPC tile, its surrounding metal frame and the copper plate improved the cooling cycle, so that a tile temperature of 11 °C to 12 °C could be achieved.

In order to prevent external light from irradiating the sensor, a black metal cap was screwed to the copper plate, covering the complete detector unit (visible in figure 5.3). When a DPC tile was used for the first time, a dark count map was recorded in total darkness, which means that no scintillator or external radiation source was used in this measurement. Micro-cells with very high dark count rates can be identified from this dark count map. The user can generate and upload an inhibit map, which disables 10 % (or any other fraction) of the micro-cells with the highest count rates. With this technique, the dark count rate was reduced significantly, e.g. from 13.2 MHz to 3.5 MHz per die at 14.5 °C. An exemplary dark count map and the corresponding inhibit map can be found in section A.2.

Once a scintillator is glued to the DPC tile, this measurement can be repeated, since there is no emission of scintillation light, as long as no high-energy source is used. In case of LYSO crystals, the radioactive component ¹⁷⁶Lu leads to scintillation events, even without using an external radiation source. However, events from LYSO and dark counts can still be distinguished. The Lutetium decays lead to high-energy events and therefore to a high number of discharged cells, while the dark counts only affect single cell discharges. However, the dark count rate of a die is in the MHz range, while the LYSO background activity is 255 Bq cm⁻³ [130]. With 2 × 2 crystals of each layer being located above a die, the background activity of 110.16 Bq is negligibly small.

During a measurement, the Peltier cooler was set to 8 °C, resulting in a sensor temperature between 11 and 12 °C. Lower cooling temperatures would have been beneficial regarding the dark count rate, but were out of option due to possible condensation of moist air (and resulting electronics malfunction). The bias voltage of the dSiPM was set automatically, after the inhibit map was loaded from the flash memory. Trigger scheme 2 (i.e. trigger set as soon as photons are counted in two subpixels) and validation scheme 0x55:0R were set to reduce the recording of dark noise. Moreover, a photon count threshold of 16 was enabled, so that only signals above 16 photons per die were written to the output file. This reduced the output file size significantly.

When the capture mode is activated, the sensor temperature usually increases, caused by the higher power consumption of the micro-cells. To stabilise the temperature, the

²Zentralinstitut für Engineering, Elektronik und Analytik (ZEA-2), Forschungszentrum Jülich, <https://www.fz-juelich.de/zea/zea-2/>.

bias voltage was adjusted multiple times within the capture mode, before the actual data acquisition was started.

For the following measurements, a ^{68}Ge source³ was used. The line source was placed at an approximate distance of 65 mm above the crystal stack, as depicted in figure 5.3. The line source was horizontally aligned at a 45° angle relatively to the crystal stack, so that it pointed from the top right corner of the DPC tile towards the bottom left corner. However, for the experiment described in this chapter, an exact positioning of the line source is not required, since a homogeneous irradiation of all crystals is desired. The line source has been chosen due to the large range of emission angles.




FIGURE 5.3: Measurement setup with a line source: crystal stack and DPC tile were located in a light-tight black housing. The line source was positioned at a fixed distance by means of a foam cube.

Data Processing

The measured data was processed in multiple steps, mainly using ROOT (6.12.04) [132] and RStudio [133] with R (3.4.3) [134]. In this section only a short description of the applied techniques will be given.

Since all dies on a tile are read out independently, the full signal of a scintillation event was acquired by collecting adjacent entries in a time window of 40 ns. Events with less than 100 micro-cell discharges in total were discarded.

For high event rates, saturation effects can occur on a dSiPM, because every micro-cell can only fire once within an acquisition cycle. To consider this effect, the event data was corrected for saturation by applying an exponential model, which also accounts for inhibited micro-cells [66]:

³When ^{68}Ge has decayed to ^{68}Ga by electron capture, ^{68}Ga decays predominantly to ^{68}Zn by positron emission (88.88 %) or electron capture (11.11 %). The decay to an excited state of ^{68}Zn (1077 keV or even higher excited states) is possible, but the respective probabilities are only 3.25 % for the 1077 keV state and <0.4 % for higher excited states [131]. The high probability for positron emission makes it a commonly used source for PET calibration.

$$N_{\text{cor}} = -N_{\text{act}} \ln \left(1 - \frac{N_{\text{dis}}}{N_{\text{act}}} \right) . \quad (5.1)$$

To calculate the corrected number of counts on a pixel N_{cor} , the number of active (i.e. not inhibited) cells N_{act} and the number of discharged cells N_{dis} were considered.

The second correction that was applied accounts for the individual dSiPM pixel gains. Since the breakthrough voltages of the individual dSiPM pixels slightly differ, the overvoltage (which is the difference between bias and breakthrough voltage) differs accordingly [59]. This leads to different gains of the dSiPM pixels, so that the same amount of incoming photons can be interpreted differently by the dSiPM pixels. However, a calibration is not straight-forward: as a consequence of the lightsharing architecture, there are no characteristic peaks in the single pixel count spectra since the scintillation light of a single crystal is shared among multiple pixels (with individual gains). Two exemplary pixel count spectra are shown in figure 5.4a. Therefore an approximate method introduced by Lerche et al. [135] is used: since there are no characteristic peaks, the end points of the individual count spectra of the dSiPM pixels are compared. To simplify the comparison, basic transformations were applied to the histograms, such as calculation of the inverse cumulative histogram and its logarithm. In a last step, the histogram was normalised (5.4b). A measure for the relative dSiPM gains was found by determining the intersection with a threshold (e.g. at 0.2), so that the respective count spectra were adjusted to the highest gain.

This method requires an irradiation with photons of only one energy. Although Lutetium has a radioactive component, these events can be neglected due to the much higher activity of the ^{68}Ge source (2.6 MBq vs 255 Bq cm $^{-3}$). Anger logic (equation 2.3) was applied to the saturation and gain corrected signals from the pixel with the highest photon count and its full neighbourhood.

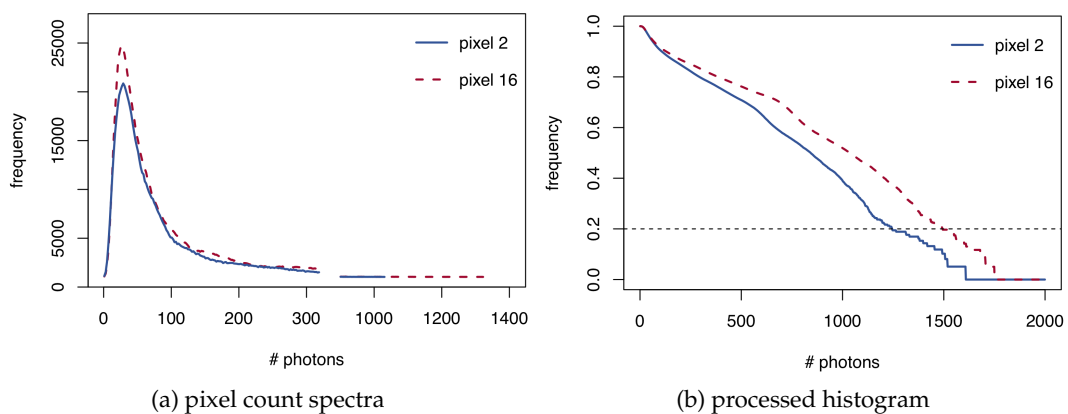


FIGURE 5.4: Approximate, relative gain estimation for the single dSiPM pixels: the count spectra for two pixel signals are shown in (a), while the corresponding gain histograms are shown in (b). A threshold determines the relative gains, and the inverse values of the relative gains are used as the relative gain correction factors.

Results

Figure 5.5 shows the resulting flood histogram for the test detector shown in figure 5.1d, using a float glass lightguide of 1.1 mm thickness.⁴ It is striking that the flood histogram appeared in well separated sections, which can be associated to the 16 dSiPM pixels. The signals in the bottom right corner corresponded to the corner where crystal array and lightguide had been aligned to the dSiPM corner. Vice versa, the spot patterns at the top and left edges corresponded to the edges of the crystal array, where the lightguide was protruding towards the dSiPM (see figure 5.1c). The sections of the flood histogram along these edges will not be discussed in detail, because this setting (lightguide larger than crystal array) will not be present in a complete detector block. The used crystal stack was only used for feasibility studies; a complete crystal array will be shown later (section 8.5).

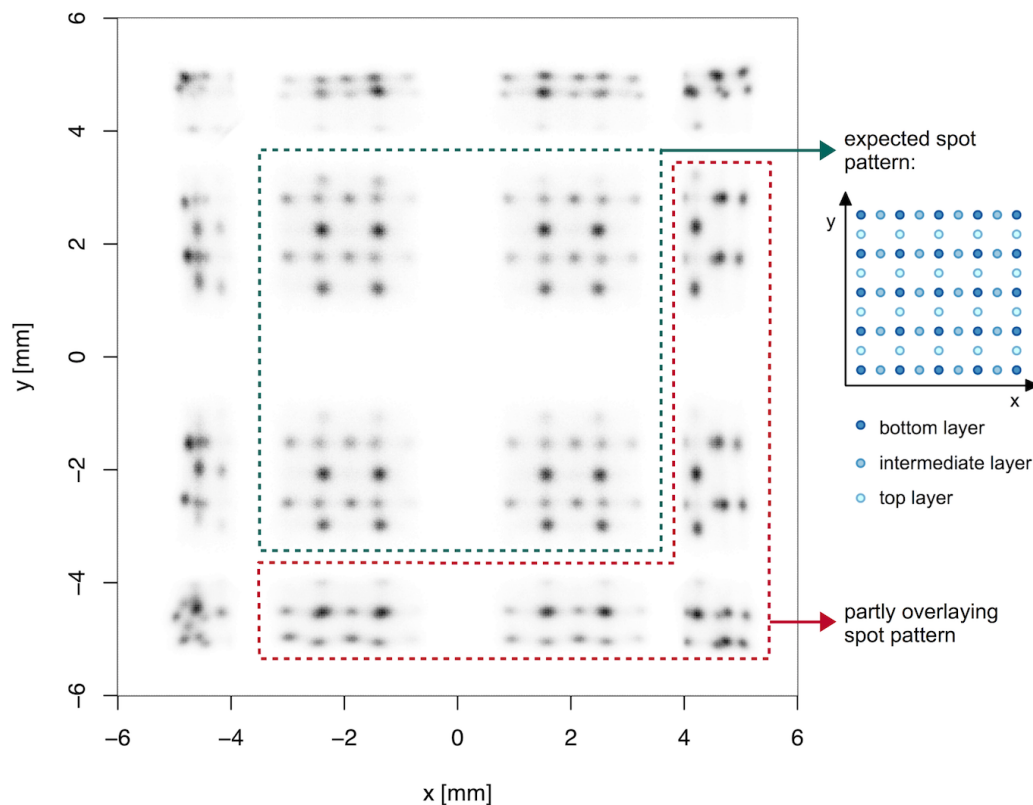


FIGURE 5.5: Flood histogram of the test detector. The three crystal layers can be clearly resolved in the centre part, while the spot pattern at the outer edges partly overlays.

Figure 5.6 shows a detailed view of the signal from a single dSiPM in the centre. The spot pattern resembled the expected one, which has been shown earlier in figure 3.4 and 5.5: the spots with highest intensity had to belong to the top layer (figure 5.6a), as this layer was shifted in y direction compared to the bottom and the intermediate

⁴Although the crystal stack is located in the right bottom corner, the coordinate system has been aligned to the centre of the stack, to ensure better comparability to the following measurements and simulations.

layer. The bottom and the top layer shared the same offset in x direction, so that the bottom layer spots could be assigned accordingly, as marked in figure 5.6b. Figure 5.6c shows the spots which belonged to the intermediate layer.




FIGURE 5.6: Detailed view of a single dSiPM signal from the centre of the flood map. The spots could be assigned to the three crystal layers and were well distinguishable.

Some crystal spots seemed to occur twice, as marked in figure 5.7. These spots presumably belonged to crystals which were either located above a bond gap or above more than one single dSiPM pixel. This was e.g. the case for a whole crystal row of the top layer, which was located above a bond gap for the largest part (see figure 5.2c). Depending on the light distribution on the sensor surface, the highest number of photons was counted on either of the underlying dSiPM pixels– the eight neighbouring pixels for the calculation of an Anger position were selected accordingly. This led to two spots in the flood map, which were assigned to the same crystal. This effect is illustrated in figure 5.7 for a crystal of the top layer (white opaque square), where the two possible hot pixels (red squares) and their eight neighbours are marked on the right side.

Discussion

In contrast to comparable flood maps from literature, the spot pattern was very regular and did not show any geometric distortions. *Gonzalez et al.* acquired flood histograms with three staggered LYSO layers (1.5 mm pitch) on an analogue SiPM [110]: although the crystal signals appeared to be separable, there were large geometric distortions visible. Moreover, the spots were not as narrow and well separated as in the flood map shown in figure 5.5. *Ito et al.* presented flood maps from two test detectors, which made use of four staggered LYSO layers of 1.5 mm pitch and 2.0 mm pitch [108]. Position-sensitive PMTs were used for their test detectors, which also led to distortions and partly overlaying spots, although the test crystal stack was located in the centre of the PMT.

While the expected spot pattern for the three crystal layers could be observed on the four centre dSiPM pixels, the spot patterns at the outer edges were hard

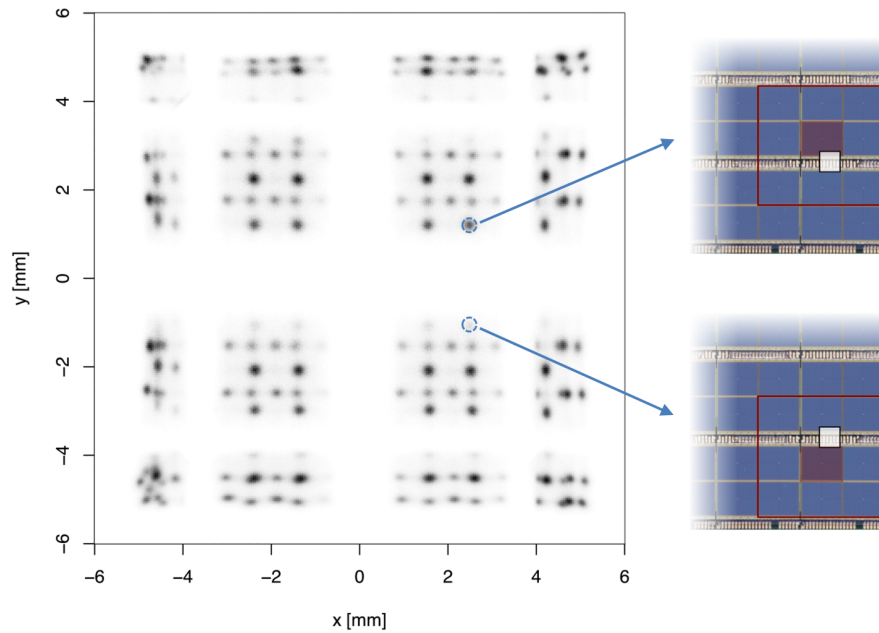


FIGURE 5.7: When a crystal is located above a bond gap or above two dSiPM pixels, two event positions are possible, depending on the maximum light detection on either of both sides.

to interpret (red box in figure 5.5). The spots of several crystals were missing and assumed to overlay with other crystals' signals, as shown in figure 5.8 (expected spot positions marked in red, not clearly interpretable spots marked in blue).

An obvious approach to study the impact of the three layers separately would be to record flood maps for one crystal layer, the combination of two and finally three layers. But since the scintillation photons are emitted isotropically, they propagate in all connected crystal layers, so that the number of crystal layers would influence the spot positions. In other words, the spot pattern belonging to the bottom layer would be altered by adding a second or third crystal layer on top, because the propagation of scintillation photons would change.

Moreover, at this stage of the project only three crystal samples and three DPC tiles were available, so that careful handling was required. The influence of the cleaning, glueing and measurement process on the devices' durability was not known, so that it was decided to investigate the flood maps by means of simulations. Further analysis of the flood histogram (e.g. energy spectra) will be shown in chapter 8.

Summary

The combination of three staggered crystal layers and digital SiPMs was tested in this chapter. Preliminary measurements were made with three crystal arrays of the dimensions $2 \times 2 \times 9$ mm and 8×8 crystals, a lightguide of 1.1 mm thickness, and a DPC tile with 4×4 dies. The flood map pattern turned out to be very complex at the outer edges of the DPC tile, because overlapping spots of the three crystal layers occurred in these areas. This is a common side-effect of pixelated and monolithic




FIGURE 5.8: Exposed view of the flood histogram from figure 5.5: the missing spots are marked in red, while not identifiable signals are marked with blue boxes.

detectors, since the scintillation light distribution cannot be completely detected at the outer edges of the photodetector array [9].

In the centre of the flood map, the expected spot pattern of the three crystal layers could be confirmed. Furthermore, it was observed that crystals located above a bond gap could give rise to two spots in the flood map, while the spot intensities depended on the exact location of the crystal. Inaccuracies of the 3D printed stencils, that were used for alignment of the crystal arrays during the glueing process, led to small displacements of the crystals on the sensor, so that either one or two spots were observed for these crystals. Since the identification of overlapping spots at the flood map edges was not straightforward, it was decided to use GATE simulations to gain deeper understanding of the flood map.

Chapter 6

Simulation of Scintillation Photons in a Detector Block

Simulations of complete scanner geometries have been presented in chapter 4, but only considering the crystal geometries, not the propagation and detection of scintillation photons. In this chapter, the focus will be on the simulation of the scintillation process inside a single detector block, in order to gain insights on the corresponding flood histograms. The timing resolution will not be investigated, since the BrainPET insert is not going to be a time-of-flight device. With the combination of LSO and DPC a time resolution of 450 ps is expected, while a timing resolution below 200 ps would be necessary to improve the performance noticeably [12].

Standard GATE simulations of PET scanners terminate particle tracks after absorption in the crystal or scanner volume (crystal housing, gantry), or after leaving the simulation world. All digitisation and signal post-processing steps are modelled by the *digitiser*, and the creation and propagation of scintillation photons is completely omitted. This is a reasonable approach for some parameters, which can be approximated with a random distribution, e.g. the individual crystal light yields. But for other parameters it makes sense to run more detailed simulations, e.g. by including the effective sensitive area of the photo sensors or the thickness of a lightguide, since they have significant influence on the flood histogram [136]. Usually, quantum efficiency, effective sensitive area, etc. are approximated by multiplying the detected energy with a factor that includes all these effects.

6.1 Simulation Setting

The main drawback of simulating scintillation photons is the enormous required computation time: for each scintillation event in LSO $\approx 13,000$ secondary photons are created, which have to be tracked inside the crystal array until they are absorbed by the sensitive detector (in this case the photo sensor). While it is feasible for monolithic crystals or pixelated detectors with a single crystal layer, this is not the case for a multilayer approach. For this reason, the supercomputer JURECA was used to simulate optical processes on a statistically sound basis.

If a full detector block of option III (24×24 crystals in the bottom layer) was simulated,

the simulation would require more than 100,000 core hours on the supercomputer (this corresponds to the current project budget for a whole year). In order to limit the computation time for a single block simulation, the size of the detector block had to be reduced.

Geometry

The simulated geometry is shown in figure 6.1: the bottom crystal array consisted of 8×8 crystals, and the other two layers of 8×7 and 7×6 crystals respectively. The crystal dimensions were $2 \times 2 \times 9$ mm. A lightguide of 1.1 mm thickness was inserted between the crystal stack and 100 μm thick cover glass of the dSiPM. The dSiPM volume contained several daughter volumes, namely the die, pixel and subpixel volumes. The subpixel areas were declared as the *sensitive detector*, which is meant to absorb and count the scintillation photons. For this purpose, the built-in digitiser module *optical adder* was used; it integrates all optical photons of the same scintillation event in the same sensitive detector (i.e. on subpixel level). But before this module was called, the module *quantum efficiency* was applied: to mimic a reduced detection probability for the photons, their energy was multiplied with the specified factor. For the following simulations, a QE of 40% was assumed.




FIGURE 6.1: The simulated detector geometry and the corresponding surfaces. The parameters of the surfaces 1-9 are defined in table 6.1.

Materials

To simulate scintillation, some basic scintillator properties had to be defined. The emission curve can be described by two exponential decays with a slow and a fast time constant. The respective emission spectra and time constants can be defined with `FASTCOMPONENT`, `FASTTIMECONSTANT`, `FASTSCINTILLATIONRISETIME`, and `SLOWCOMPONENT`, `SLOWTIMECONSTANT`, `SLOWSCINTILLATIONRISETIME`, while their ratio is set with `YIELDRATIO`. The number of scintillation photons, which is produced per energy unit, is sampled from a normal distribution. The expectation value of optical photons per MeV is set with `SCINTILLATIONYIELD`, and the variance is defined in `RESOLUTIONSCALE`.

For LSO, the following values were applied [41]: scintillation yield of $26\,000\text{ MeV}^{-1}$, resolution scale 4.41 (this corresponded to an intrinsic energy resolution of 8.9% at 511 keV), 39.6 ns fast time constant, 80 ps fast scintillator rise time [70] and a yield ratio of 1, which means that the slow component was neglected.

Surfaces

The surface model used in these simulations was the UNIFIED model, which has been first introduced by *Levin and Moisan* in 1996 [137] (as an extension of the work of *Nayar et al.* [138]). A surface is modelled as an ensemble of micro-facets, whose normal vectors follow a Gaussian distribution. The user can set the surface roughness by varying the standard deviation σ_α of the angular distribution. Moreover, the probabilities for different types of reflection have to be specified.

The values for σ_α and the reflection probabilities were initially chosen according to [33], [93], [139] or information provided by manufacturers. Table 6.1 contains basic properties of the surfaces, which had to be defined in "both directions", i.e. from one volume to the adjacent one and vice versa. The lateral crystal surfaces were polished and coated with a specular reflector (3M ESR vikuiti). The reflective foil itself was not implemented as a volume, but rather as a reflective crystal surface with 98.5% reflectivity [140]. The crystal faces which were in contact with glue were defined as rough ("ground") surfaces. Both glass layers were polished, while the lateral faces of the lightguide were "as cut", which has been considered to be rough. The dSiPM chip surface was assumed to be nearly perfectly smooth– surface roughness and reflectivity of the pixel material were not available from the manufacturer, and therefore had to be estimated.

TABLE 6.1: Surface definitions of simulated optical interfaces and surfaces. R: Reflectivity, T: Transmittance, E: Efficiency.

	Surface	σ_α	R [%]	T [%]	E [%]
1	Pol. crystal \leftrightarrow pol. crystal	1.3	98.5	1.5	0
2	Pol. crystal \rightarrow air	1.3	98.5	1.5	0
3	Crystal layer \leftrightarrow crystal layer	12.0	5	95	0
4	As cut crystal \leftrightarrow as cut crystal	12.0	98.5	1.5	0
5	As cut crystal \rightarrow air	12.0	98.5	1.5	0
6	Crystal layer \rightarrow lightguide	1.3	15	85	0
7	Lightguide \rightarrow cover	1.3	15	85	0
8	Cover \rightarrow dSiPM pixel	1.3	5	0	95
9	Lightguide/cover \rightarrow air	1.3	15	85	0

Ground Truth Information

The output file of the simulation only stores information that has been acquired in the sensitive detector volumes. To record the true interaction positions of the high-energetic photons (DOI information), a new GATE actor was implemented: it recorded time, position and deposited energy of processes that involved high-energetic photons. These were distinguished from optical photons by their PDGEncoding and by making use of the GEANT4 *UserSteppingAction* [85].

Figure 6.2 shows a scene from a GATE simulation, where a high-energy photon (represented by the red line) was scattered two times inside the crystal stack (interactions 1 and 2), before the remaining energy was deposited by photoelectric effect (interaction 3). The trajectories of the scintillation photons are represented by green lines. In order to produce visible trajectories, the scintillation yield had been lowered to 500 photons per MeV for the creation of this figure; otherwise the whole crystal stack would have been illuminated in green, without distinguishable trajectories. Table 6.2 contains information on the three interactions, that has been acquired with the GATE actor.

Three types of ground truth were obtained from the actor information: the first impact of the high-energy photon (= interaction 1), the interaction with the highest energy deposit (= interaction 3), and the mean of all interaction positions, weighted by the correspondingly deposited energy.

TABLE 6.2: Ground truth information on the three interactions of the high-energy photon shown in figure 6.2. The information was recorded with a GATE actor.

	Type of interaction	Deposited Energy	Layer	Crystal	3D position
1	Compton scattering	137.05 keV	1	20	{0.13, -0.55, 13.25}mm
2	Compton scattering	117.04 keV	1	35	{5.49, 7.80, 3.21}mm
3	Photoelectric effect	256.91 keV	2	40	{4.26, 3.93, 2.48}mm

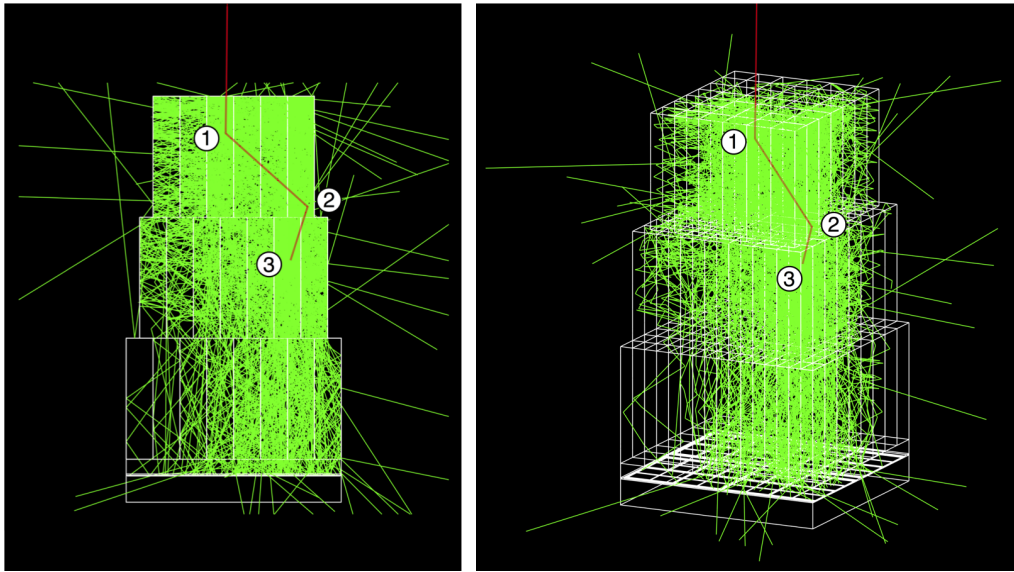


FIGURE 6.2: Visualisation of a scintillation event inside the test detector. The interactions of the high-energy photon are represented by a red line and highlighted with numbers, while the scintillation photons are visualised with green lines. The light yield was lowered to 500 photons per MeV, so that the optical photons' trajectories can still be distinguished.

6.2 Acquisition of Floodmaps

Methods

Similar to the measurements, the detector block was irradiated with 511 keV photons of a line source with 1 MBq activity. The emission angles of the source were confined to avoid photons being emitted away from the detector block, as indicated in figure 6.3. The distance to the top layer was set to 65 mm accordingly. A simulation of 240 s was run on JURECA, using 20 nodes with 24 tasks per node. Two sets of data were acquired in the simulation: first, the ground truth information of the high-energy photons (interactions in the crystal stack), and second, the number of photons that have been detected by each subpixel. For each event, the pixel with the highest

number of counted photons was used as a hot pixel to define the nine pixels that contribute to the event position calculation. The calculation of the flood histogram was analogous to the measurement, only the gain correction was neglected.




FIGURE 6.3: Setup of the simulation using a line source: the source is placed 65 mm from the top crystal layer, and emits 511 keV photons in the specified solid angles.

Results

Figure 6.4 illustrates the simulated flood histogram and the measured one from chapter 5.2. The simulated flood map showed a very regular spot pattern, which matched the (theoretical) expectation and which could be directly associated to the three crystal layers. The spots at the edges were shifted towards the centre, as the light distribution could only be partly detected due to the missing neighbouring dSiPM pixels. This is a known side effect, which occurs in every measured flood map of a pixelated detector [108]. It has been shown that the distortion depends on lightguide thickness and pixel size of the detector [136].

It is striking that the simulated spot pattern did not resemble the measured one. There were neither any gaps between the single dSiPM pixels' signals, nor was the pattern at the edges or in the corners in any way similar to the measurement. Moreover, each crystal could be identified in the simulated map, while there were overlaying crystal signals in the measurement. Since crystal identification problems motivated this kind of simulation, the simulation itself had to be revised and matched to the measurement before further investigations could be done.

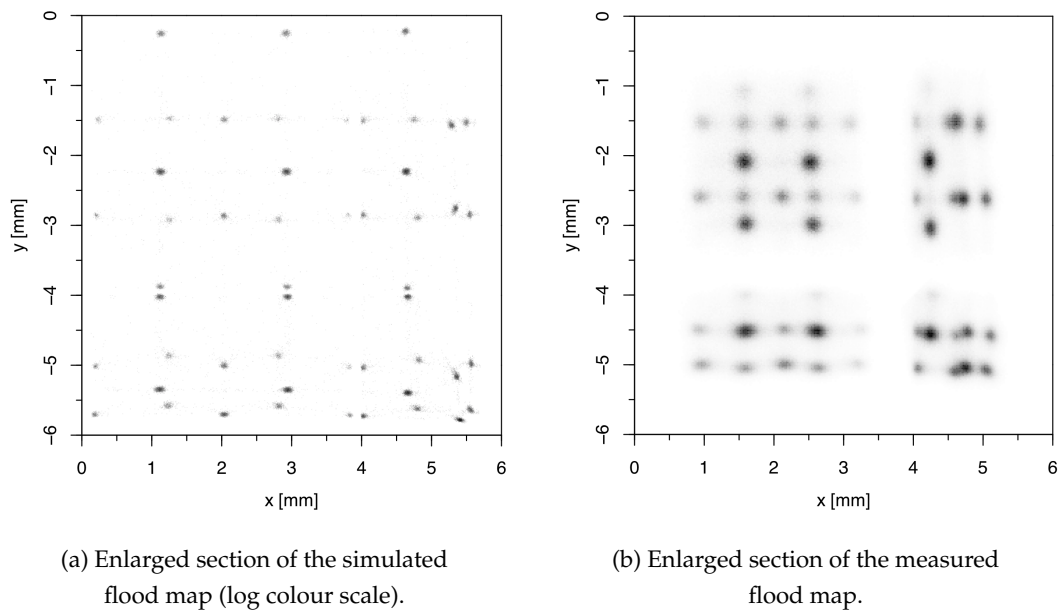


FIGURE 6.4: Simulated and measured flood map for the detector prototype with 1 mm lightguide.

Discussion

A variety of parameters has been modified to find a setting for which simulation and measurement match. In most simulations of detector blocks, the optical grease or glue between crystal and photosensor has been included due to its different refractive index [92], [93], [141], so that the addition of glue layers was an obvious approach. In our simulation, it did not have the desired impact on the flood map, probably because the abrupt change of the refractive index from LSO ($n_{\text{LSO}} = 1.82$ [142]) to glass (approximately $n_{\text{glass}} = 1.46$ [143]) is not significantly altered by a thin glue layer with $n_{\text{glue}} = 1.5$. To reduce simulation time, the glue layers were not considered any further in simulations.

A much higher impact had the reflectivity of the crystal layer coatings, the lightguide and the dSiPM. The UNIFIED model is known to be very sensitive to modifications of the adjustable parameters, since the user has to estimate surface roughness (σ_α) and the probabilities for specular or lambertian reflection, absorption and transmittance [33], [144], [145]—not all parameters can be found in literature. Moreover, the parameters have different dimensions, which makes the correct determination of the parameters even more complex. This drawback is made subject of discussion in reference [146], where the authors attempt to predict the influence of specific surface parameters on the optical spectrum by extensive simulations. Nevertheless, it was not possible to predict and adapt the parameters in a way that both simulated and measured flood maps became consistent.

Refinement of the reflectivity and transmittance values for different wavelengths did not yield the desired change in the flood map neither, but made the choice of

simulation parameters more complicated. For example it was difficult to estimate the reflectivity of a glass surface covered with glue or a roughened crystal surface from measured data in related literature. It is questionable whether the provided values really match the used crystals and glass sheets, also because surface finishes differ among manufacturers. Furthermore, care had to be taken how the different values were specified. Inaccuracies which are created by the way Geant4 interpolates between the given data points are discussed by *Dietz-Laursonn* in [147]. Since the spectral dependencies did not seem to be the most relevant parameter, monochromatic light with a wavelength corresponding to the peak emission wavelength (420 nm) was selected for the simulations, in order to avoid introduction of erroneous parameters and unnecessary complexity of the simulation.

The influence of the photons' incidence angles on the probability of reflection is often measured, but cannot be considered within the UNIFIED model [137]. Therefore, a reflectivity distribution had to be chosen which was representative for all incidence angles. To improve the underlying model, *Janecek and Moses* stored measured surface data in look-up tables (LUTs), which are used to calculate the angle of reflection depending on the incidence angle [33]. The use of measured surface data is favoured, since the incidence angle depends on the crystal's geometry and surface finish [148]. A further disadvantage of the UNIFIED model is the approach of modelling surface roughness as an ensemble of micro-facets, which always have the same height, but differ in their angular distribution of the facets [137], [145]. While this model can yield reasonable results for mechanically polished surfaces, this is not the case for ground and etched surfaces, where the micro-facet distributions cannot be assumed to be Gaussian [33]. This inaccuracy can be decreased by using measured data.

Janecek and Moses' approach has been further refined by *Roncali et al.*, who have added angular reflectivity and transmittance measurements to the model [144]. Especially the last aspect is crucial for the simulation of the specular reflector: high light crosstalk has been reported in pixelated LSO/LYSO arrays, despite using the 3M ESR vikuiti (a quantitative analysis can be found in reference [149]). *Loignon-Houle et al.* could identify reflectivity quenching of the 3M ESR vikuiti under certain incidence angles as the major cause for light crosstalk, when the reflector is used between two crystals, bonded with glue [150]. They further state that compared to other effects like light loss at pixel edges or the decrease of the 3M film reflectivity below 395 nm [140], the relatively high transparency of the reflector above incidence angles of 60° has the strongest influence on cross talk [150]. However, *Loignon-Houle et al.* demonstrated that this reflectivity quenching depends on the coupling materials, i.e. it occurs between two bonded crystals, but cannot be observed in a reflectivity measurement taken in air.

Roncali's model cannot yet be applied to bonded pixelated scintillator arrays, since the measurements are (so far) taken with single crystals in air [144]. Therefore, the UNIFIED model had to be used for further simulations, but with adapted parameters.

6.3 Simulations Including Reflectivity Quenching

The study of *Loignon-Houle et al.* [150] showed a strong dependency of the 3M ESR vikuiti reflectivity on the photons' incidence angle. The simulated reflectivity from their study at the used wavelength of 420 nm is shown in figure 6.5. The reflectivity drops to less than 70% for photons, which impinge in a very flat angle (above 65°, normal incidence at 0°). For angles smaller than 60°, the reflectivity remains very high, above 95%. The probability for the different incidence angles depends on the crystal geometry. *Loignon-Houle et al.* show that there is a considerable amount of photons, which impinges under angles where the reflector becomes highly transparent. The measurements of *Lorincz et al.* [139] and *Padera et al.* [151] are shown for comparison. When measuring the 3M ESR vikuiti reflectivity in air, values >95% can be obtained for all measured incidence angles. The reflectivity quenching cannot be observed in these measurements.




FIGURE 6.5: Simulated and measured 3M ESR vikuiti reflectivity for different incidence angles at a wavelength of 420 nm. Data adapted from references [139], [150], [151].

Methods

Since it was not possible to include angular dependency in the UNIFIED model, the overall reflectivity of the 3M ESR vikuiti between the crystals was lowered to 70%. The reflectivity at the interface towards air was 96%, as described in table 6.3. Both values had to be determined empirically: the energy spectra of single dSiPM pixels and the overlapping spots at the outer edges of the flood map were compared for different parameter settings. Reflectivities of 70% and 96% yielded simulated energy spectra and a spot pattern, that mostly resembled the ones from the measurement.

TABLE 6.3: Updated definitions for the simulated optical interfaces and surfaces. The surfaces which are not mentioned in this table remain unchanged (see table 6.1). R: Reflectivity, T: Transmittance, E: Efficiency.

	Surface	σ_α	R [%]	T [%]	E [%]
1	Pol. crystal \leftrightarrow pol. crystal	1.3	70	30	0
2	Pol. crystal \rightarrow air	1.3	96	4	0
3	Crystal layer \leftrightarrow crystal layer	12.0	5	95	0
4	As cut crystal \leftrightarrow as cut crystal	12.0	70	30	0
5	As cut crystal \rightarrow air	12.0	96	4	0

Results

The corresponding flood histogram of a simulation with reduced reflectivity of the 3M ESR vikuiti is shown in figure 6.6. The measured spot pattern could now be reproduced better with a GATE simulation of the detector geometry. Ground truth information was accessed with the GATE actor and allowed separation of the crystal signals.




FIGURE 6.6: Simulated flood map with reduced 3M ESR vikuiti reflectivity.

Figure 6.7 compares detailed views of single dSiPM signals from simulation and measurement. It is striking that the spot size was much smaller in the simulation (figure 6.7a). In other parts of the flood histogram, the crystal signals were still separable in the simulation, while some of them were *presumably* overlaying in the measurement. The spot positions were comparable, but when investigating the measured energy or number of detected photons (shown in figure 6.7c and figure 6.7d), a main drawback of the simulation became apparent. The electronic readout of the sensor could not be modelled within the GATE digitiser, which lead to extremely high photon counts in the energy peaks (around 4600 and 5500 in the simulation, around 2000 and 2500 in the measurement, marked with dotted lines).

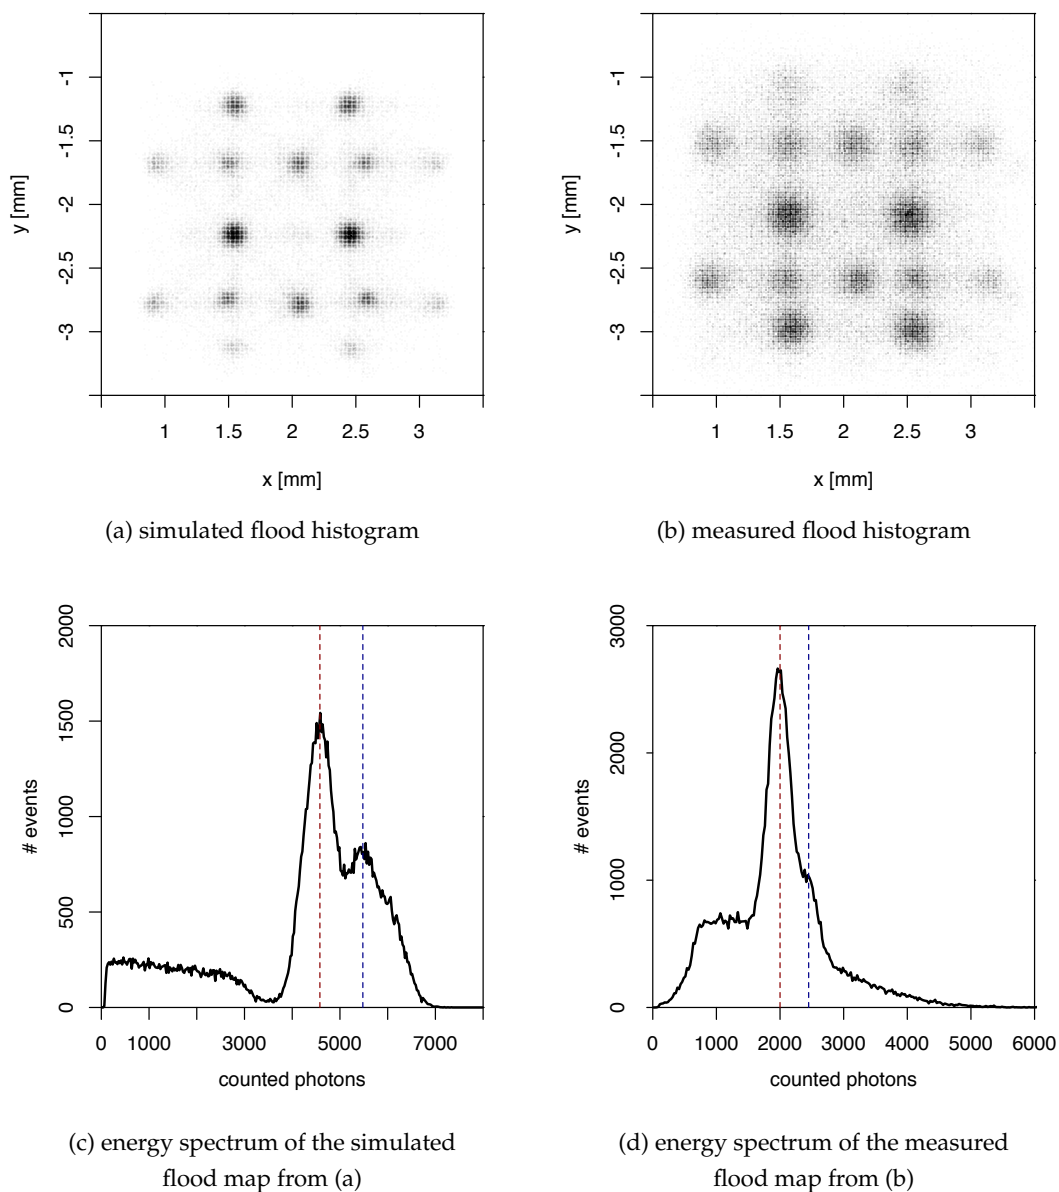


FIGURE 6.7: Exposed view of a single dSiPM signal, (a),(c) simulated and (b),(d) measured. The same number of events is displayed. For comparison, the two energy peak positions in the energy spectra are marked with dotted lines.

6.4 GATE Digitiser Dedicated to DSiPMs

In order to account for the influence of readout electronics in simulations, an additional digitiser in GATE was implemented. The standard GATE digitiser modules can be applied to simulations with optical photons, but this has several disadvantages. In simulations of complete scanner systems, optical photons are not simulated due to the enormous computational demand. Therefore, nearly all digitiser modules are dedicated to high-energy photons, e.g. the module for *quantum efficiency*: it reduces the photons' energies to the specified percentage, but still detects and counts each single photon (albeit an optical or a high-energy photon). While this might be a suitable approximation for *singles* and *coincidences*, this is not the case for optical photons and the readout electronics, where triggers can be set on single photon level. So far, trigger thresholds could not be modelled with the existing modules.

In general, the digitiser is called at the end of each event, which means that all tracks of primary and secondary particles have ended. The digitiser modules are then applied in the specified order to the list of pulses ("GatePulseList"), which have been collected during the event. As a consequence, all scintillation events are handled separately, and are clearly distinguishable. This is not the case in a measurement, where scintillation events have to be extracted from the measured raw data by clustering the recorded time stamps (e.g. within a time window of 40 ns). Depending on the detection time, it is not possible to separate all events with this approach, which leads to the detection of *pile-up* events. Within the GATE digitiser, a pile-up module can be enabled to model this effect (the pile-up time should correspond to the integration time then). But, by default, even two scintillation events with identical time stamps would still be separated by the event-based digitiser.

The recharge time of the micro-cells in a DPC could be considered by applying the module for *dead time*, but the application of this module would not make the simulation more similar to the measurement, since the duration of dead time varies. It corresponds to the readout and recharge time for a validated event, otherwise only to the recharge time, if the event did not fulfill the validation logic.

The main goal of the dedicated digitiser is to process incoming simulated optical photons in a similar way as in the DPC tile.

6.4.1 Structure of the Digitiser

The basic difference compared to the standard digitiser is the processing of hits, which are created during the simulation of optical photons. Instead of event-based processing, multiple events are collected in a buffer and then evaluated together, in order to allow mutual influences of the scintillation photons. However, as for all existing GATE digitiser modules, the new module is still called after each event. This procedure allows the user to apply additional (existing) modules. The flowchart in figure 6.8 shows the general structure of a simulation.

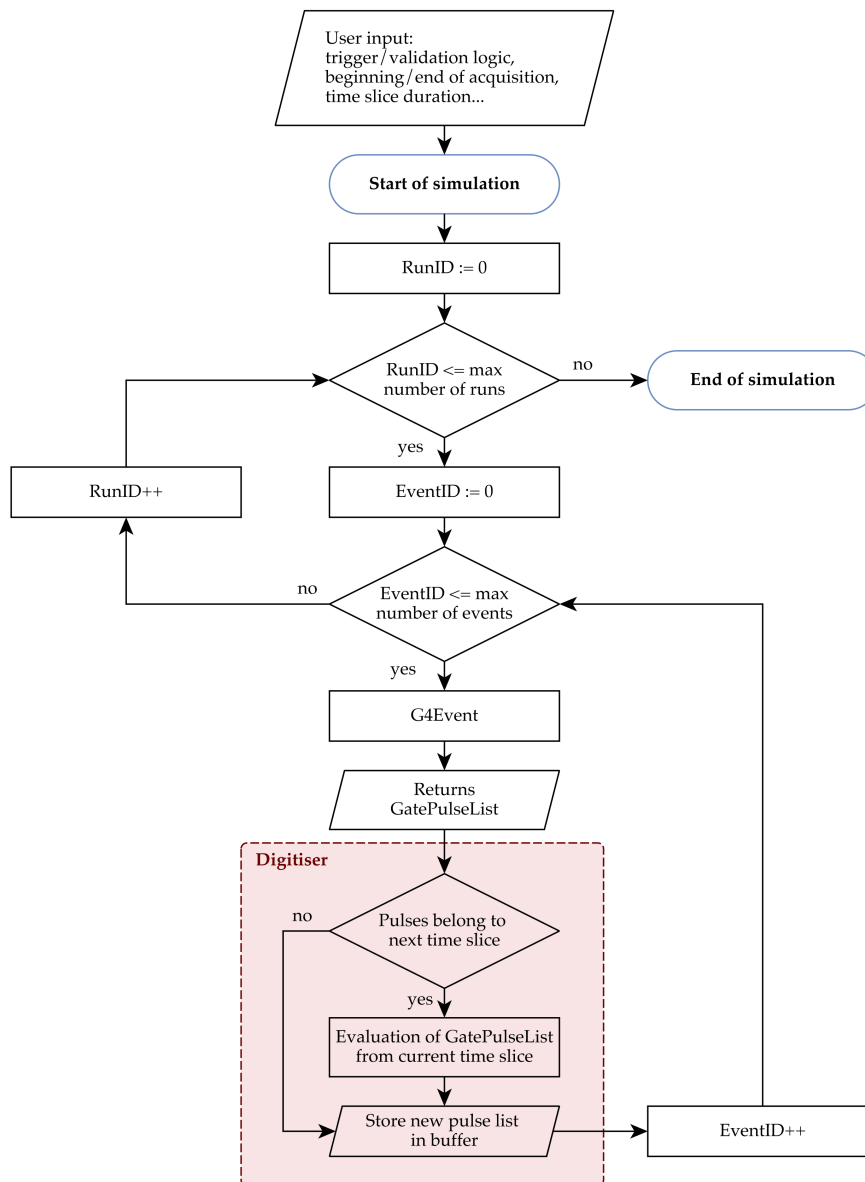


FIGURE 6.8: Outline of the simulation: the digitiser module is called after each event.

After specification of the simulation time and its division into time slices, where simulated geometry and physics are assumed to be constant, the simulation is split into multiple runs, where each run corresponds to a time slice. The number of events, which will be processed in each run, depends on the source definition and is determined by the source manager. In general, the main programme loops over all runs and events.

After each event, a list of pulses (hits) is passed to the digitiser (box marked in red). As long as the event belongs to the same time slice as the previous event, the pulse list is cached. For each pulse, a random sample is drawn from a uniform distribution. In case it is above a certain threshold (i.e. the specified quantum efficiency of the sensor), it is discarded, otherwise it remains within the pulse list. The pulses are sorted by

dieID, pixelID and subpixelID, and stored in separate vectors, to be evaluated at a later time together with the other events. When a new time slice starts, the collected pulse list is evaluated in the pipeline shown in figure 6.9.

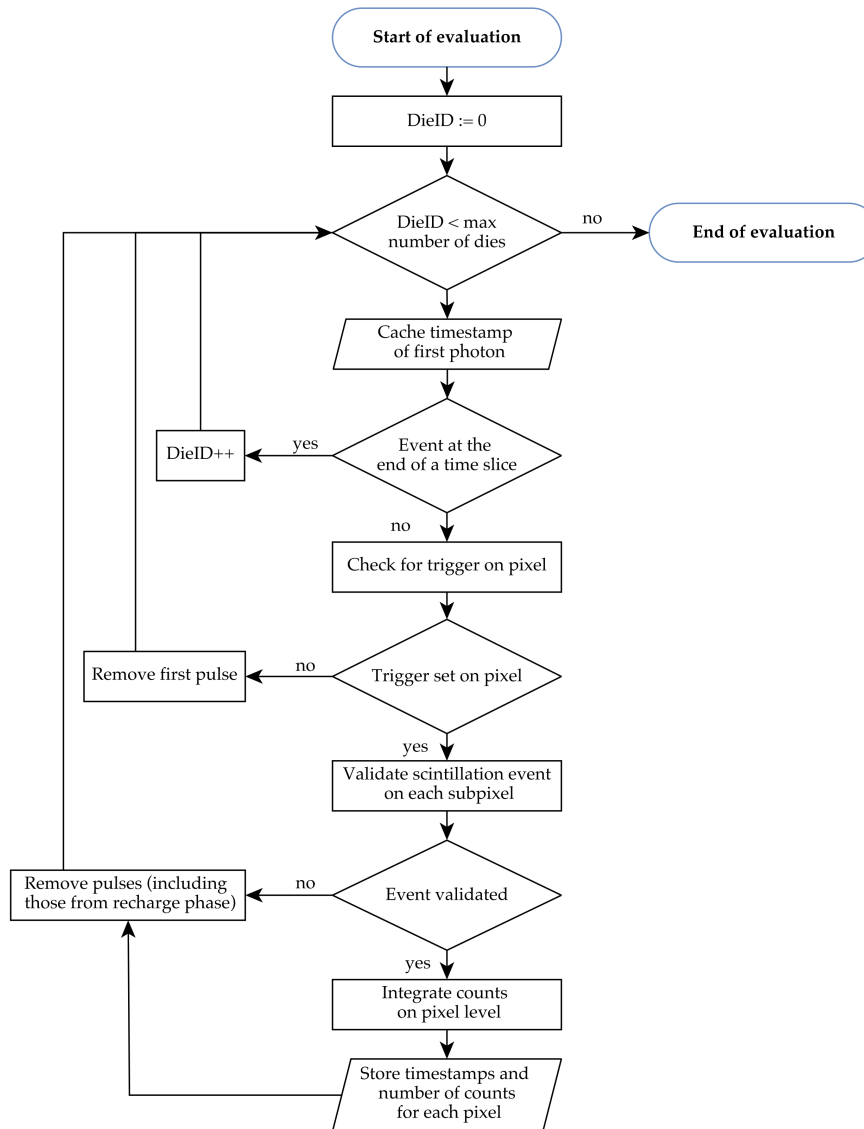


FIGURE 6.9: Structure of the digitiser module. The pulse lists of multiple events from the same time slice are evaluated together. The underlying readout scheme is shown in figure 2.9.

Analogous to measured data, the die signals are analysed separately. Before any analysis is started, the timestamp of the first photon on the die is compared to the end of the time slice. If the event occurs at the end of a time slice, and the integration and readout time overlaps with the beginning of the next time slice, the pulse list is kept as it is for combined evaluation with the next time slice. If this is not the case, the geometric distribution of the following pulses (on the first photon's pixel) is compared to the selected trigger scheme. The first photon is removed from the pulse list, when no trigger can be detected. The search for a trigger starts again, using the

'new' first photon on the die, which can be located on a different pixel. Within the described simulation, a trigger can usually be found within less than 10 iterations. The timestamp of the trigger photon is cached and used as a reference time for the validation, integration and readout intervals.

As soon as a trigger is found, the event validation process starts. For each subpixel of the trigger photon's pixel, the pulses are binned into row-trigger-lines, so that their spatial distribution can be compared to the selected validation logic. An invalid event leads to removal of the involved pulses, as well as those pulses that are lost during the recharge phase. If an event is validated, the pulses during the integration phase are counted and stored with the first photon's timestamp. During the readout phase, micro-cells stay active until readout. Since they can still detect photons, half of the readout time is added to the integration time [68]. Afterwards, the number of counted photons on each pixel is stored together with the trigger photon timestamp in the output file. The pulses are removed from the pulse list, including those from the recharge phase.

6.4.2 Validation of the Digitiser

Due to the limitations of the UNIFIED model, which have been described in section 6.2, experimental validation of the digitiser by comparison to measured data remains challenging. Therefore, another approach was chosen. *Tabacchini et al.* have developed an analytical model, which describes the probability of triggering and validation as a function of discharged micro-cells [69]. This model allows for calculation of the average number of micro-cells, which have fired until the trigger or validation pattern has been fulfilled on a pixel. *Tabacchini et al.* validated their model with measurements, where a single die of a dSiPM tile was illuminated with a homogeneous light distribution (temperature $-25\text{ }^{\circ}\text{C}$). This light distribution was created by laser pulses and an integrating sphere. For acquisition of the probability of triggering, *Tabacchini et al.* recorded dark counts in total darkness and assumed those to serve as a homogeneous illumination.

The trigger and validation probabilities were reproduced with the new GATE digitiser, while the simulation itself was kept rather simple. A monolithic scintillator block ($7.875 \times 7.15 \times 10\text{ mm}^3$, LSO) was coupled directly to a single dSiPM die and was irradiated with a 511 keV photon source. The scintillation light was collected and quantified with the dSiPM digitiser module, if the DOI was $\leq 5\text{ mm}$, to ensure a homogeneous light distribution. The surface definitions corresponded to table 6.1. Moreover, noise was added to the simulation. The dark count rate at $-25\text{ }^{\circ}\text{C}$ was reported to be 20,000 cps [68], so that an exponential noise distribution with a time constant of 0.05 ms and particle energies of 2.965 eV was chosen. Since *Tabacchini et al.* used a different naming convention for the validation schemes, table 6.4 provides a short overview of the corresponding names from the Philips DPC manual.

TABLE 6.4: Overview of the selectable validation schemes, with corresponding names from the Philips DPC manual [66] and the publication of *Tabacchini et al.* [69].

Scheme	Subpixel connection	Validation pattern Philips DPC manual	Validation pattern <i>Tabacchini et al.</i>
1	OR	0x7F:OR	1-OR
2	OR	0x77:OR	2-OR
4	OR	0x55:OR	4-OR
8	OR	0x00:OR	8-OR
16	AND	0x55:AND	4-AND
32	AND	0x00:AND	8-AND

Trigger Probability

The selected setup provides simulations at low dark count rates, so that all recorded cell discharges were considered to be relevant for the trigger. By binning these cell discharges into a histogram, the probability of triggering was directly accessed from the normalised histogram.

Validation Probability

The procedure for determining the validation probabilities for the different patterns was analogous to the trigger probabilities. For all validation patterns, the lowest trigger threshold was set (scheme 1, 0x7F).

For each trigger and validation scheme, 20,000 detected events were evaluated. The number of cell discharges before fulfilment of the trigger or validation pattern was counted and compared to the results of *Tabacchini et al.*

Since some of the results of *Tabacchini et al.* are only shown under the influence of 16 % crosstalk¹, a simplified loop to account for crosstalk was inserted into the digitiser: when a photon was detected (after application of the module quantum efficiency), a random number between 0 and 1 determined whether the photon created another one, i.e. when this random number was ≤ 0.16 . In this case, the pulse was simply duplicated.

Results

The average number of cells, that were discharged when the trigger condition was met, is shown in figure 6.10. The error bars correspond to the standard deviation. *Tabacchini et al.* did not provide these values for the measured data, but stated a good agreement of the analytical model with the measured values and those from the manufacturer. Moreover, the presented figures let the reader expect similar mean values and standard deviations. On die level (figures 6.10b and 6.10d), the determined

¹This value has been obtained from measurements in [69], as described in [68].

values were slightly lower than those from *Tabacchini et al.*, but matched within the given standard deviation, which was smaller for the GATE simulation. On pixel level (figures 6.10a and 6.10c), the deviation was larger, especially for trigger scheme 3 and 4. However, the triggered cells in the GATE simulation still lay within the standard deviation of *Tabacchini's* observations. This observation was made in both cases, when crosstalk was omitted or included in the simulation. The average number of discharged cells was slightly higher (about 1 to 3 additionally discharged cells) when crosstalk was included.

For the validation schemes, very similar observations were made, so that the corresponding figures are displayed in A.3 for the sake of better readability.

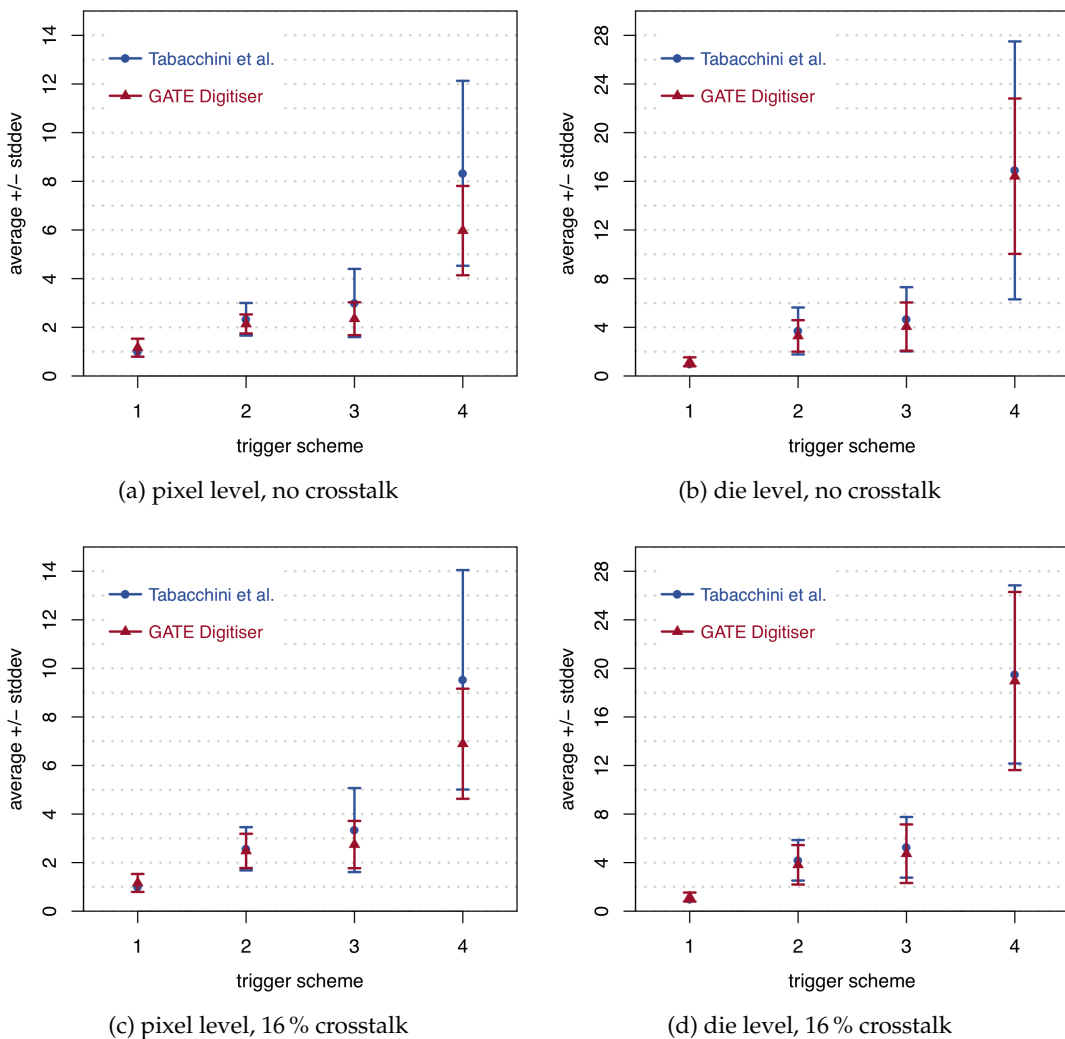


FIGURE 6.10: Number of discharged cells at the exact time when the trigger condition is reached. The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d).

Figure 6.11 shows the normalised histograms for trigger probability when using trigger scheme 2 and assuming 16% crosstalk. In *Tabacchini's* analytical model, the histogram was shifted towards higher values, which means that the probability for triggering

was higher, when more cell discharges occurred. The figures for trigger scheme 3 and 4 are shown in A.3.




FIGURE 6.11: Probability, that the trigger is set after a certain number of cell discharges, displayed for trigger pattern 2.

The influence of crosstalk is displayed for validation pattern 8-AND in figure 6.12. For both models, GATE and the analytical one, the cumulative probability of event validation is shown. When crosstalk was enabled, the probability of validation was shifted towards a higher number of cell discharges, as it was expected. In both cases, the histograms obtained with GATE were in good agreement with those from *Tabacchini et al.*




FIGURE 6.12: Cumulative probability, that validation scheme 8-AND is fulfilled after a certain number of cell discharges. The results are shown with and without the influence of crosstalk.

Discussion

In general, the values obtained with the GATE digitiser were very similar to those from *Tabacchini et al.* Two things were striking: the standard deviation of *Tabacchini's* observation was larger, and the number of discharged cells at the moment of trigger generation or event validation was larger on pixel level.

The reason for the deviation on pixel level has proven to be the source definition. In the GATE simulation, the light distribution on the dSiPM was created by scintillation events in a monolithic crystal. Care was taken that the light distribution was very homogeneous by limiting the depth of interaction to maximum 5 mm.

In case of the trigger probabilities, *Tabacchini et al.* measured dark counts to simulate uniform illumination. The dark count rates can vary among pixels and dies, depending on the sensor temperature and the amount of disabled cells in the inhibit map. Therefore, the number of discharged cells at trigger generation can be subject to high fluctuations.

In case of the validation probability, *Tabacchini et al.* used a pulsed laser with presumably 50 ps pulse width². This light source was emitting photons on a different time scale, since the scintillator has a finite rise time of $\tau_r = 80$ ps and a decay time of $\tau_d = 39.6$ ns [70]. The emission profiles of both sources are visualised in figure 6.13. Due to the faster emission in case of the laser pulse, a higher number of photons hit the dSiPM during the validation phase, compared to the scintillation pulse. The emission of photons during a shorter time influenced mainly the discharged cells on pixel level.

As a proof of principle, the scintillation pulse has been modified such that the decay curve resembled the laser pulse ($\tau_r = 20$ ps, $\tau_d = 30$ ps). Figure 6.10 was reproduced with the new artificial scintillation source in figure 6.14.




FIGURE 6.13: Emission profiles of a 50 ps laser pulse and an LSO scintillation event.

²The pulse width was not explicitly mentioned and therefore deduced from other publications of this group.

As expected, the number of discharged cells was closer to *Tabacchini's* calculated values on pixel level, and the standard deviation had increased, both for trigger and validation logic. In case of the validation logic, the number of discharged cells at the moment of event validation was lower than that of *Tabacchini et al.* on die level, but still lay within the standard deviation. The corresponding figure can be found in section A.3. It should be emphasised that GATE samples the number of emitted photons from a normal distribution, i.e. the number of emitted photons was not constant, even when the same amount of energy was deposited. This is another important difference to *Tabacchini's* approach. The shortening of the rise and decay times and the varying number of emitted photons can therefore cause high fluctuations in the number of discharged cells before the event validation.

Further adjustments of the simulation setup were not investigated, since the obtained values were already in good agreement with *Tabacchini's* approach, so that the implemented digitiser was considered to be working correctly.

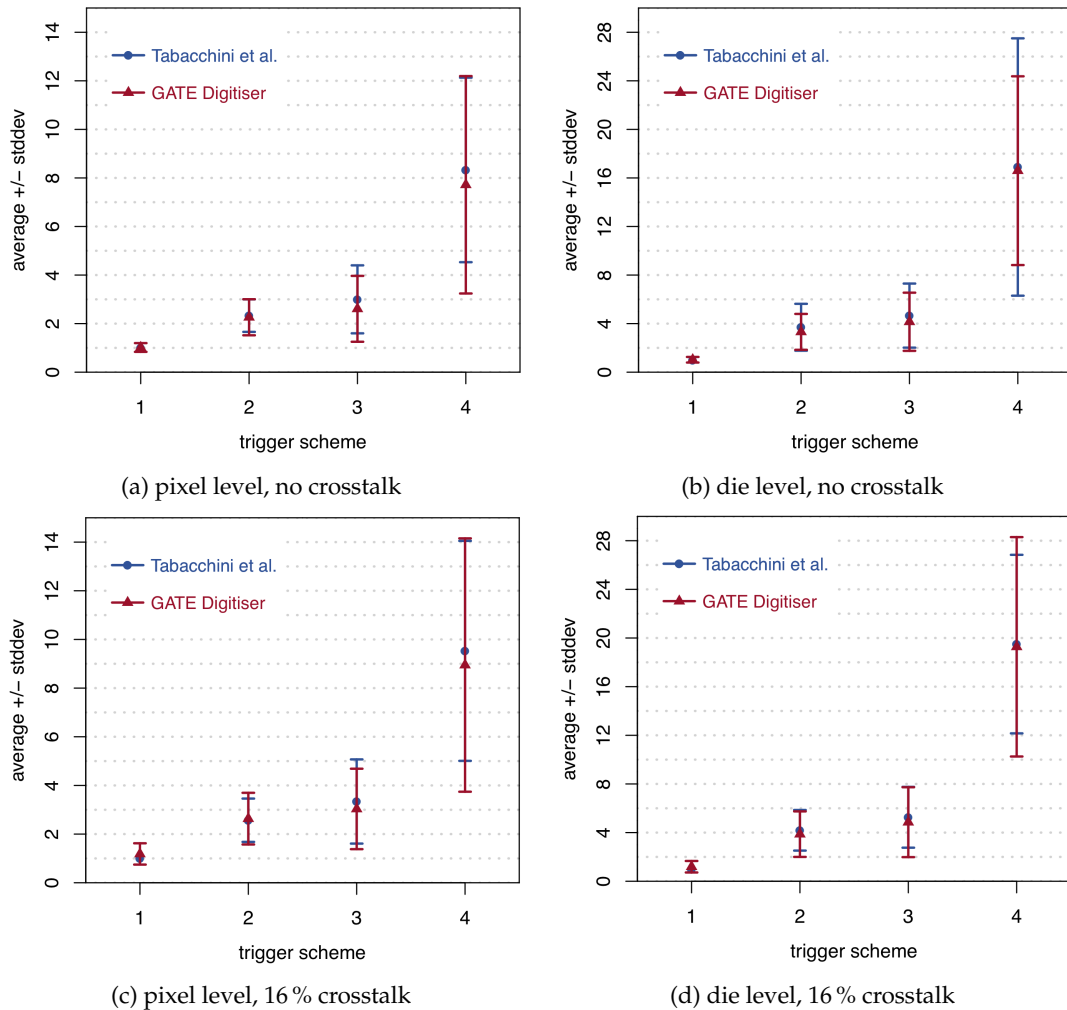


FIGURE 6.14: Number of discharged cells at the exact time when the trigger condition is reached, shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d). In this simulation, the artificial scintillation source was used for photon emission.

In conclusion, the new digitiser for GATE can model the DPC specific readout. For analogue SiPMs, there exist simulation frameworks like GosSiP ("Generic framework for the simulation of Silicon Photomultipliers") [152] and G4SiPM [153]. These aim at simulating the (electrical) SiPM response to an incoming light distribution, which can either be modelled by customisable light pulses, or which can be created externally using Geant4. The latter allows for simulation of a scintillator geometry and of the resulting electrical SiPM response. The analogue output signal is calculated after the simulation of scintillation events and after adding dark counts, crosstalk and afterpulses. With the newly developed GATE digitiser, the DPC output signal can be determined during the simulation, since the tool is completely embedded in the GATE source code.

A disadvantage is the temporary storage of scintillation photons from multiple events, which can be memory-consuming. The memory consumption can be controlled by adjustment of the source activity and the time slice duration, since these parameters determine how many events are evaluated together. A typical simulation with the three crystal layer geometry was started as follows: a line source with 1 MBq activity irradiated the detector stack during a total simulation time of 360 s, which was divided into 0.003 s time slices. 24 processes with 500 runs each were submitted to each node, which resulted in a wall clock time between 20h54min and 21h32min. Since 10 nodes were occupied, the used computation time amounted to 5085 core hours. The maximum amount of memory used by a process was ~ 5.52 GB.

6.4.3 Results: Application to Previous Simulations

The new digitiser module can be called in GATE by using the following commands:

```

/gate/digitizer/pdpc/insert dsipm
/gate/digitizer/pdpc/dsipm/setNbDies 4
/gate/digitizer/pdpc/dsipm/setTriggerScheme 2
/gate/digitizer/pdpc/dsipm/setTriggerInterval 12.5 ns
/gate/digitizer/pdpc/dsipm/setValidationInterval 40.0 ns
/gate/digitizer/pdpc/dsipm/setValidationPattern 4
/gate/digitizer/pdpc/dsipm/setIntegrationInterval 325.0 ns
/gate/digitizer/pdpc/dsipm/setReadoutInterval 680.0 ns
/gate/digitizer/pdpc/dsipm/setRechargeInterval 20.0 ns
/gate/digitizer/pdpc/dsipm/chooseVolume subpixel
/gate/digitizer/pdpc/dsipm/setQE 0.4

```

The simulation of a detector prototype, which has been described in 6.3, was repeated with the new digitiser. Figure 6.15 compares the result of the simulation using the new GATE digitiser and the corresponding measurement.

The improvement of the new digitisation procedure is clearly visible: both the spot size and the energy spectra were much more similar to those of the measurement, in contrast to the previous simulation that has been shown in section 6.3. The spot patterns in the flood histograms were similar, but differed slightly at the edges. Especially at the top and bottom edge, where crystals of the top layer were located above

a bond gap. While the spots belonging to the top row ($y = -1.2$ mm) were stronger pronounced in the simulation, these were barely visible in the measurement. The opposite case applied for the crystal signals at $y = -3.0$ mm. Multiple measurements of the same crystal stack have shown that these spot intensities strongly depend on the location of the crystals. Inaccurate alignment of the crystal arrays during the gluing process becomes directly visible in the flood map, and can therefore be assumed to be one reason for the remaining difference.

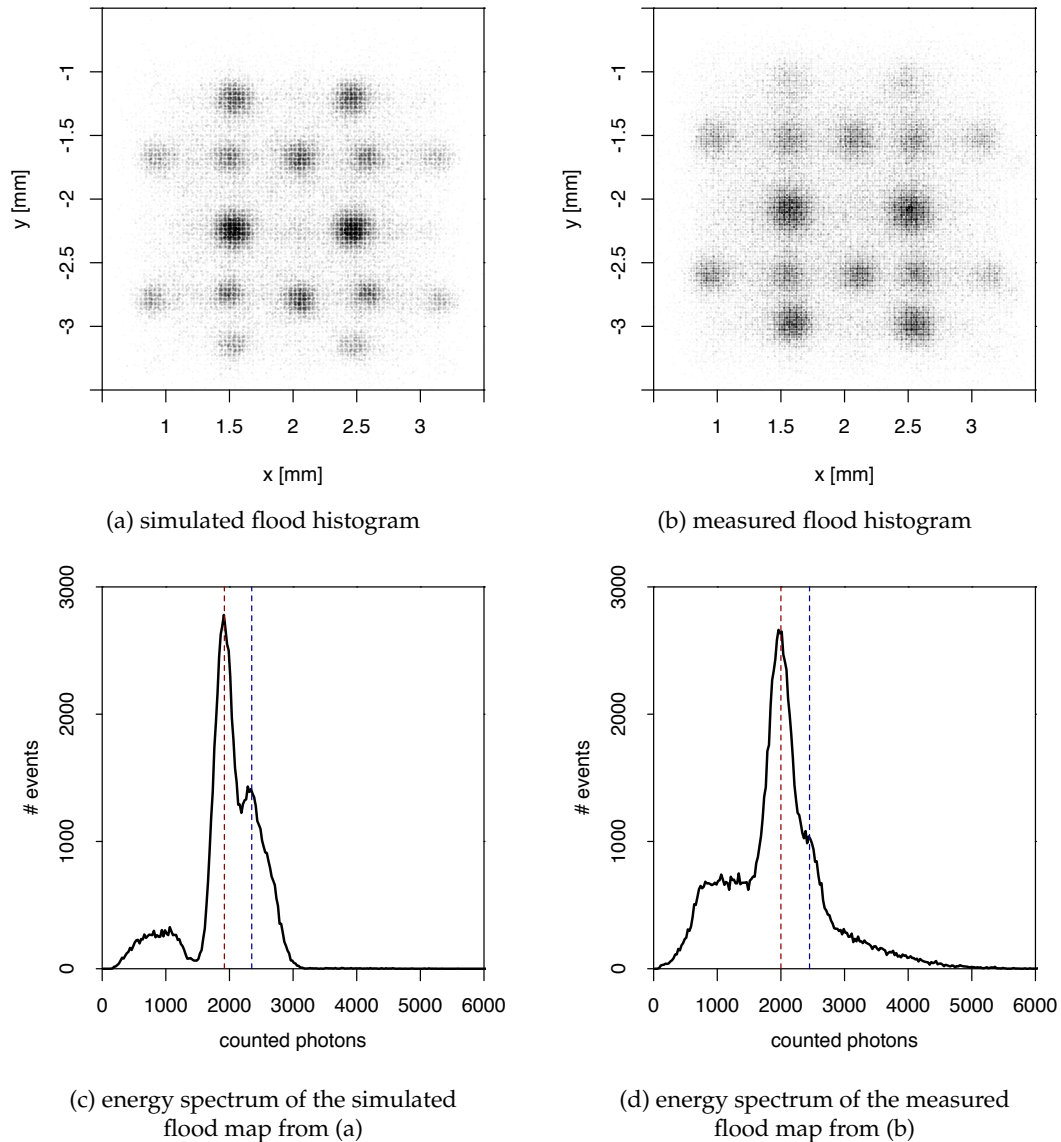


FIGURE 6.15: Exposed view of a single dSiPM signal, (a),(c) simulated and (b),(d) measured. The same number of events is displayed. For comparison, the two energy peak positions in the energy spectra are marked with dotted lines.

Both energy spectra showed two peaks, around 1900 photons and 2350 photons in the simulation (2000 photons and 2500 photons in the measurement). Moreover, there was a noticeable amount of events with energies below 1500 photons, which was more pronounced in the measurement. The lower energy can be caused by Compton

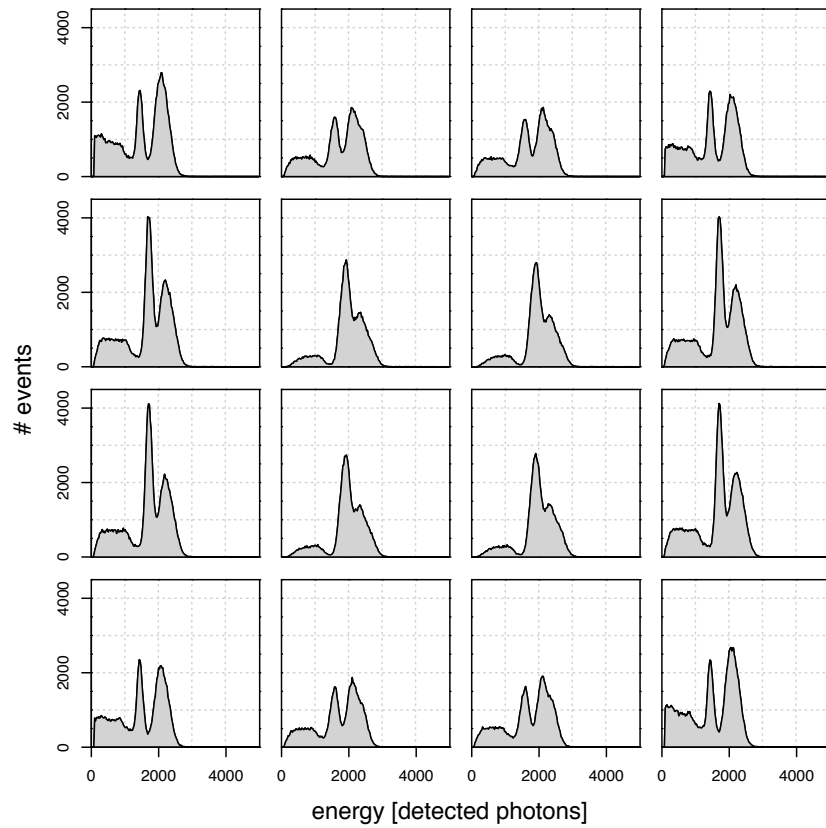


FIGURE 6.16: Simulated energy spectra of the 16 dSiPM pixels. The lutetium background radiation was not considered in this simulation.

scattering, where the photon was either scattered to a different crystal block, or where fractions of the photon's energy were deposited too far away from one another to be detected as one event ("far Compton event"). In the measurement, Compton back scatter from the testbench (e.g. cooling unit or table) was also possible, but has not been considered in the simulation. Additionally, missing dSiPM signals can cause the lower event energies, e.g. when one of the dSiPMs was still within the readout or recharge phase and therefore did not contribute to the signal.

Figure 6.16 shows the simulated energy spectra, and figure 6.17 shows the measured energy spectra for comparison (using the same number of events): while the centre dSiPMs only contained two close peaks, there was another peak at the outer edges (marked with blue arrows) in the measurement.

An effect, that has not been considered in the simulation, is the the natural background radiation of ^{176}Lu , which is contained in LSO in small quantities (2.6 % in natural lutetium). The isotope ^{176}Lu decays to ^{176}Hf , mainly via β^- decay (99.61 %, $E_{\text{max}} = 593.2$ keV), which is followed by three γ decays (93.6 % at 306.78 keV, 78 % at 201.83 keV, 14.5 % at 88.34 keV) [154]. Especially the γ decays were detected at the outer edges of the DPC tile: the low deposited energy lead to a smaller amount of scintillation photons, which was often detected at the edges, since only 4 or 6 triggered dSiPMs were required to fulfill the neighbourhood criteria. This hypothesis

is supported by the fact that the peak position was approximately half of the full energy peak. Moreover, this peak became smaller when only those events were considered, where more than four or six dSiPMs triggered (larger neighbourhood and therefore larger light spread required). This effect is depicted in figure 6.18.

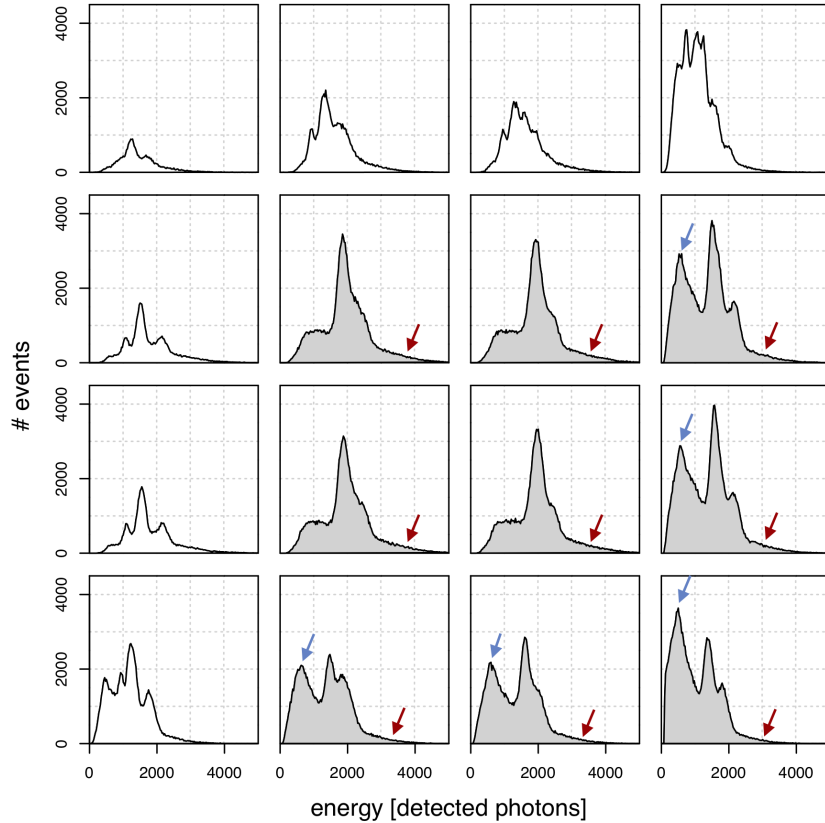


FIGURE 6.17: Measured energy spectra of the 16 dSiPM pixels. The lutetium-related peaks (blue arrows) are clearly visible at the outer edges, while they are not visible in the centre, where always 9 dSiPM signals are required for an event. The red arrows indicate events presumably originating from the 1077 keV γ emission line of ^{68}Ga . The top row and left row (white) were not considered in this analysis, since the crystal arrays of the detector prototype were smaller than the lightguide and DPC tile (alignment shown in figure 5.2c).

In contrast to the simulation, the measured energy spectrum contained events with energies higher than 3000 photons (marked with red arrows in figure 6.17). The ^{68}Ge source, that was used for the measurement, decays to ^{68}Ga , which predominantly decays to ^{68}Zn by positron emission. However, there is a probability of 3.25 % that ^{68}Ga decays to an excited state of ^{68}Zn at 1077 keV [131], so that these photons could cause the high-energy events in the measured energy spectrum. Pile-up of simultaneous or successive events could have also contributed to these high-energy events, but since pile-up was included in the new GATE digitiser, this contribution seems to be negligible.

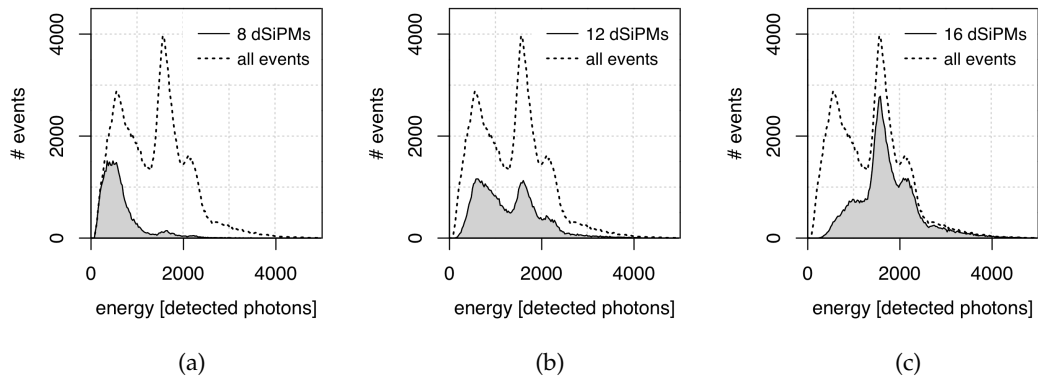


FIGURE 6.18: A single energy spectrum of a dSiPM at the right edge of the sensor tile is shown. The lutetium-related peak was dominant in events with up to 8 triggered dSiPMs (a), but decreased when more neighbouring dSiPMs (12 (b) or 16 (c)) were considered. The smaller amount of scintillation photons was not sufficient to trigger a signal of all 16 dSiPMs.

6.5 Additional Information from the Simulation

The ground truth of the simulation allows one to study the crystal layers' signals individually, and therefore to acquire deeper understanding of the measured flood-maps of the detector prototype. Figure 6.19 shows the simulated flood map, where the signals of the three crystal layers are drawn in different colours.




FIGURE 6.19: Simulated flood histogram: the event positions are encoded with the ground truth information (red: bottom layer, green: intermediate layer, blue: top layer). The 16 dSiPM signals have been normalised to the respective maximum values.

It has been mentioned earlier in section 5.2 (figure 5.7) that crystals, which are located above a bond gap, can give rise to two signals in the flood histogram: either of the underlying dSiPM pixels can count the highest number of photons, depending on

the distribution of scintillation photons on the DPC tile. Since the eight neighbouring pixels are selected accordingly, there are two typical Anger position for these crystals. Figure 6.20 depicts the three crystal layers' signals separately. Especially the flood histogram of the intermediate layer confirmed that crystals above a bond gap give rise to two signals. In the simulation, this effect lead to two spots of similar intensity: the crystals were located exactly above a bond gap, so that the probability to detect the highest amount of photons on either of the underlying dSiPM pixels was 50%. This is rarely the case in a measurement. Since the crystal arrays cannot be perfectly aligned during the gluing process, the crystals usually cover a larger fraction of either of the underlying dSiPM pixels. Accordingly, there can occur two spots of different intensity, or only one single spot.

In the simulated data set, there were 10×8 spots originating from 7×8 crystals, and 6×10 spots from 6×7 crystals in the top layer.

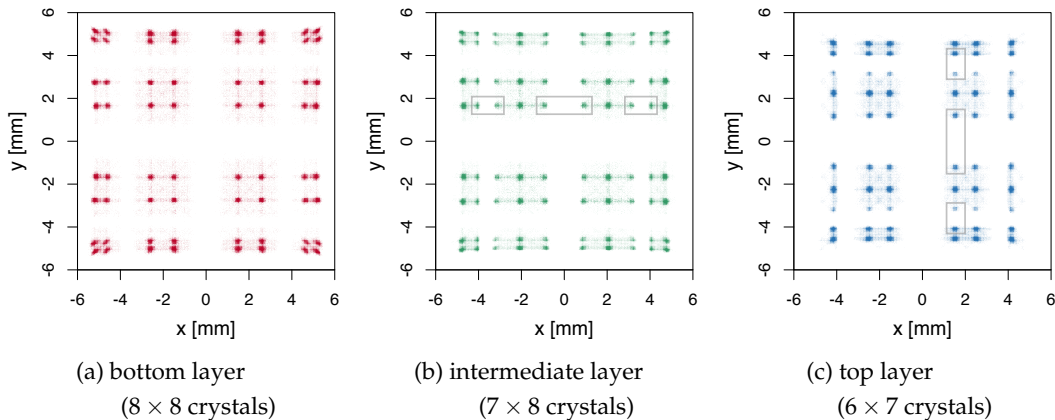


FIGURE 6.20: Separate flood histograms of the three crystal layers. Crystals above a bond gap gave rise to two spots with lower intensities in the flood histogram. Two exemplary rows of spots, which belong to only one crystal row each, are marked with grey boxes.

While the decoding of the flood histogram was straight-forward in the centre, the case was more difficult at the edges. In chapter 5, the points which could not be identified have been marked in figure 5.8. The simulation allows one to visualise the spot distributions of each layer separately, so that overlapping points can be identified.

Figure 6.21 compares the measured and simulated flood histograms of a single dSiPM at the right edge of the sensor. Originally, the two spots between the points of the top layer (blue) were assumed to originate from the bottom layer, judging by their position in the flood map. Since the bottom and top layer are only shifted relatively to each other in y direction, the crystal signals have been expected to be at similar x coordinates (see schematic drawing in chapter 3, figure 3.4). Actually both spots originated from the intermediate layer, and the bottom layer signal (red) was overlaying with the first crystal row of the intermediate layer.

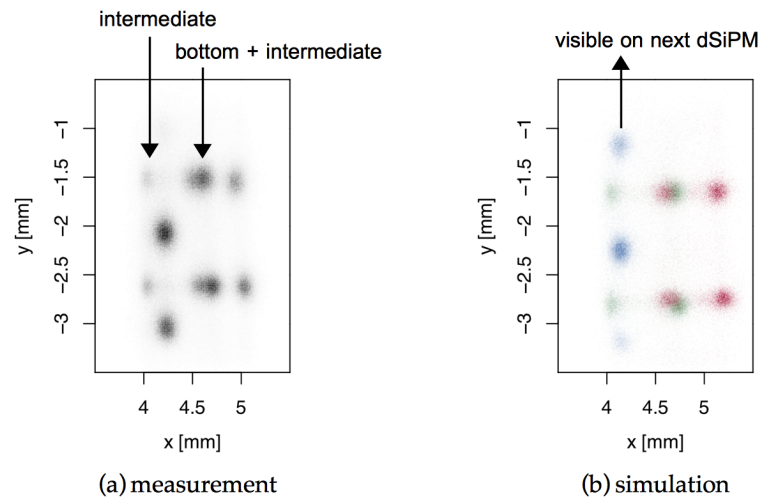


FIGURE 6.21: Exposed view of a single dSiPM signal at the right edge of the sensor, (a) measured and (b) simulated.

Figure 6.22 compares the measured and simulated signals of a dSiPM pixel at the lower edge. As previously depicted in figure 5.8, the signal of the top layer could not be clearly identified in the measured data. The simulated ground truth revealed an overlap of the top and bottom layer signal. Due to inaccurate alignment of the crystal arrays on the DPC tile, the whole crystal stack was slightly shifted in positive y direction. This shift caused the missing signal of the second crystal row from the top layer (which was visible on the dSiPM above).

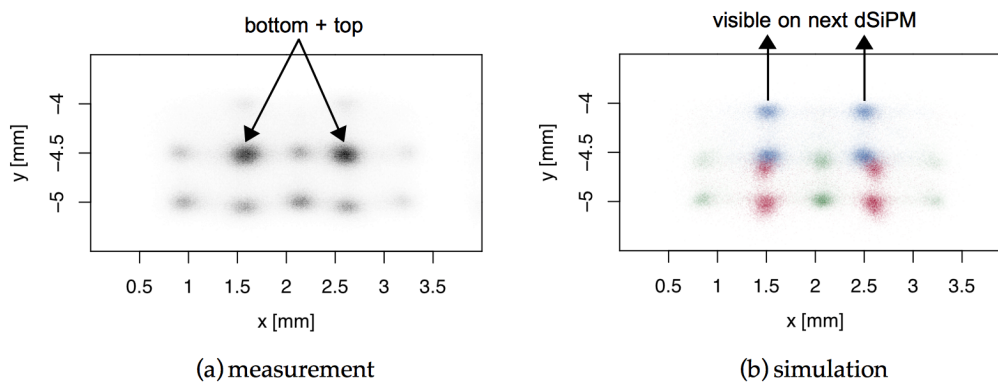


FIGURE 6.22: Exposed view of a single dSiPM signal at the lower edge of the sensor, (a) measured and (b) simulated.

The dSiPM in the corner of the sensor contained both overlaying signals of the intermediate and top layer, and of the bottom and intermediate layer, as can be seen in figure 6.23.

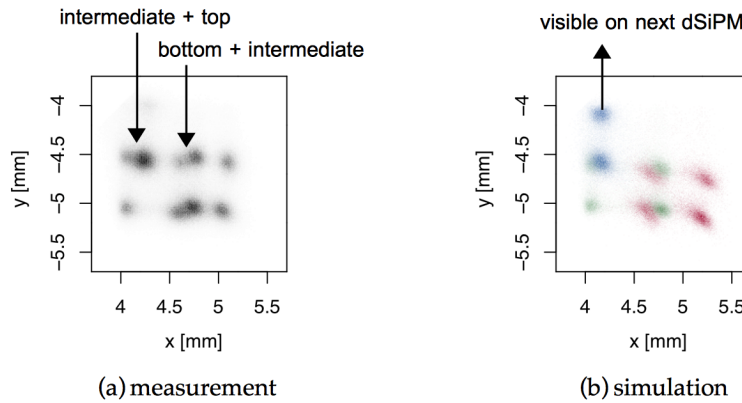


FIGURE 6.23: Exposed view of a single dSiPM signal at the corner of the sensor, (a) measured and (b) simulated.

Having decoded the flood histogram of the three layer design, further simulations will be used to find optimal lightguide dimensions and to study the identifiability of the individual crystal signals in the flood histogram.

Summary

In this chapter, the creation and propagation of scintillation photons was simulated using GATE. Despite difficulties during the implementation of surfaces, the simulated and measured flood maps could be brought to good agreement by using the UNIFIED model. The advanced surface model (Davis model [144]) cannot be applied to bonded crystal arrays so far, so that the UNIFIED model had to be used. Reflectivity quenching of the 3M ESR vikuiti between the crystals emerged as a significant effect and was considered by reducing the overall reflectivity to 70 %.

In a second step, the electronic readout of the DPC tile was implemented in the form of a new digitiser for GATE. This digitiser reproduced the measured energy spectra and spot sizes in the flood map more accurately compared to the standard digitiser. Moreover, effects like pile-up, dead time and the trigger and validation schemes of the DPC tile were considered. In contrast to related work for analogue SiPMs (e.g. G4SiPM [153] or GosSiP [152]), this digitiser was fully integrated into the GATE simulation source code. It was applied online, so that no post-processing was required to gain the dSiPM output signal. Realistic flood maps and energy spectra were acquired from the simulation along with a ground truth information. This allowed one to fully understand the acquired flood maps from the measurements, and to use the ground truth information for further detector design studies.

Chapter 7

Study on Lightguide Thickness

Many contemporary PET scanners make use of the lightsharing principle by combining a pixelated scintillator array, a continuous lightguide and a pixelated photodetector array (e.g. GE Signa [113] or Siemens mMR [78]). While crystal and photodetector specifications are often made public, the thickness and material properties of the lightguide are usually kept confidential.

Surti et al. [123] presented this concept for application in a PET detector: by using a lightguide, the light distribution of individual crystals on the photodetector array could be controlled. The impact of different lightguide thicknesses has been investigated by *Surti et al.* [123] and *Kang et al.* [136], and it was implied that there should exist an optimum lightguide thickness for a detector. In case the lightguide is too thick, the scintillation photons spread widely over the photodetector array: the low number of photons, which arrive at the outermost pixels, might not even be sufficient to generate a trigger on that die or pixel. Moreover, the amount of pile-up events increases with larger light spreads over the sensor. On the other hand, the lightguide has to be thick enough to allow enough light spread, so that the position of a scintillation event can be calculated from 8 to 12 dSiPM pixel signals [58].

Having matched simulation and measurements of the detector prototype, simulations will be used to find an optimum lightguide thickness.

7.1 Preliminary Study

In preliminary measurements, the lightguide of the detector prototype (so far 1.1 mm) has been replaced by a 2.0 mm thick lightguide of the same material (Borofloat glass). The procedures for calibration and data acquisition were the same as in chapter 5: the detector stack was irradiated for 30 min by a ^{68}Ge line source with 2.6 MBq activity, which was located 65 mm above the crystal stack.

Figure 7.1 compares the acquired flood histograms for both lightguides.

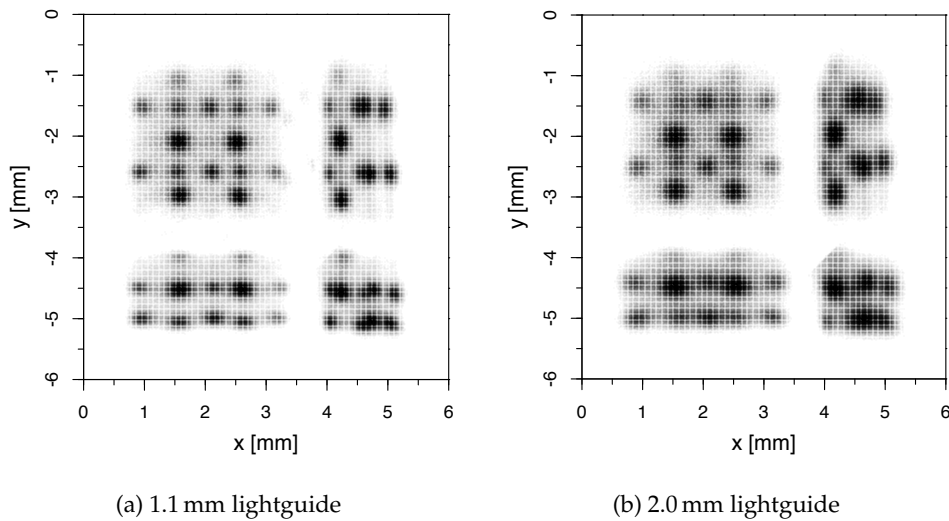


FIGURE 7.1: Flood histograms of measurements with 1.1 mm and 2.0 mm lightguide thickness. A comparable number of events is displayed in both figures.

Generally speaking, the individual crystal signals appeared to be better separable in the flood histogram using a 1.1 mm lightguide. At 2.0 mm thickness, the single spots had become broader, causing more overlap, especially at the edges of the DPC tile. A high spot separability is crucial for the calibration of a detector stack, since all individual crystal signals have to be separated during this step— this topic is subject of chapter 8. The aim of this study was to find a lightguide thickness, that optimises the spot separability both at the edges and in the centre.

A quantitative comparison of spot separability was difficult in the presented measurement, due to the different spot patterns in the centre and at the edges. Therefore, the crystal stack was moved to the centre of the DPC tile for the following study. At the centre of the DPC tile, a very regular pattern is expected from the simulation, as shown in figure 7.2. For the quantitative comparison, three representative crystals were selected on each of the four centre dSiPMs. These were entirely located above the sensitive area of the DPC tile (not above a bond gap), and gave rise to neighbouring spots in the flood histogram. In this study, the peak to valley ratios (PVRs) for these selected crystals were investigated under consideration of the different lightguide thicknesses. The PVRs were considered as a measure for the spot separability, where a higher ratio corresponded to better separability.




FIGURE 7.2: Simulated flood histogram for the detector prototype with a 1.1 mm lightguide. The peak to valley ratio was evaluated for three spots on each of the four centre dSiPM pixels (marked with black boxes), using flood maps simulated with different lightguide thicknesses.

7.2 Simulations with Different Lightguides

Methods

The simulated detector geometry was similar to the one from chapter 6: the three 9 mm thick staggered crystal layers were located on top of a lightguide, whose thickness was varied between 0.7 mm and 2.0 mm. Since Borofloat is available in 0.7 mm, 1.1 mm, 1.75 mm and 2.0 mm thickness, these values were used in the simulation. During this study, the crystal stack was aligned to the centre of the DPC tile. The dSiPM geometry itself remained unchanged, and the new GATE digitiser was used to model the DPC-specific electronic readout.

Although a line source was used for measurements, a point source had been chosen for the simulations. Since representative crystals had been selected from the centre, it was not necessary to irradiate the full detector stack in the simulation, which is already computationally demanding. The emission angles of the point source were confined, so that only the selected crystals were irradiated. Still, around 8,000 core-hours were required for the simulation of the detector block with one selected lightguide thickness. Figure 7.3 shows the simulation setup along with the emission angles of the point source.




FIGURE 7.3: Setup for simulations with different lightguide thicknesses. The point source emits 511 keV photons in confined angles to efficiently irradiate the crystals of interest.

To evaluate the simulated data, flood histograms were calculated with Anger logic. The sections for the four centre dSiPM pixels were divided into horizontal slices by using local minima of the 1D histograms in y direction (figure 7.4a). The horizontal slice, which contained the crystal signals of interest (indicated by a red arrow in figure 7.4a), was further evaluated to determine the peak to valley ratios of the bottom and intermediate layer signal.

The 1D histograms were created with a fixed bin size of 0.02 mm, as shown in figure 7.4b. Determining the peaks and valleys directly from the discrete histogram data was avoided, because noise lead to multiple local minima or maxima. Approximating the data with superpositions of two Gaussian curves required too many manual adjustments of the parameters and especially the underground signal, so that the fit functions would neither have been reproducible nor fairly comparable. Instead, the kernel density estimate of the histogram was calculated ¹. Figure 7.4c shows the density estimate of the 1D histogram from figure 7.4b. The two peak positions of interest and the valley in-between (indicated by red arrows) can then be determined.

¹The kernel density estimate is a non-parametric smoothing method, where each single data point, which would be binned into a histogram, is modelled by a smoothing kernel (in this case a Gaussian kernel) [155]. The kernel bandwidth was calculated with Sheather and Jones' bandwidth estimator [156] to prevent over- or undersmoothing. Summing up all kernels yielded the estimated probability density.




FIGURE 7.4: Analysis of the flood histogram: (a) The flood histogram was divided into horizontal slices, the slice of interest is indicated by a red arrow. (b) The 1D histogram of the selected horizontal slice is calculated, and the kernel density estimate is determined in (c) (red line). The peaks of interest and the valley in-between are marked with red arrows.

Results

Figure 7.5 shows the resulting PVRs for the three layers. Since the PVR was calculated twice for the bottom layer (once in horizontal direction and once in vertical direction), the mean value was calculated from eight PVRs. The error bars correspond to the standard errors.




FIGURE 7.5: Mean PVRs (\pm standard errors) of the four crystals from each crystal layer.

In general, the PVRs were very different for the three layers: the highest separability was given for the top layer, while the intermediate and bottom layer had significantly lower PVRs. The lowest values were obtained for the bottom layer, which was closest to the DPC tile. Observing figure 7.5 in detail, the expected tendency of an optimum lightguide thickness was visible. The PVRs of all layers reached the highest values either at 0.7 mm or 1.1 mm. At larger lightguides, the crystal signals in the flood histogram became wider and began to overlap, which decreased the separability.

Discussion

In case of the top and intermediate layer, the underlying crystal layers allow spreading of the scintillation light and therefore have a similar effect as an additional lightguide. For the bottom layer, this is not the case, so that the least events are always detected in this layer (the full neighbourhood criterion is harder to meet due to the smaller light spread). This assumption is shown in figure 7.6: the ratio of the number of all gamma interactions in a layer, and of the number of detected events that fulfil the full neighbourhood criterion is separately shown for the three layers. In case of the bottom layer, at maximum 30 % of all interactions in the layer were eventually detected as an event, where the full neighbourhood of the hot pixel triggered a signal. This ratio was much higher for the top layer, between 41 % to 45 %. The ratio increased with the lightguide thickness, but did not continuously improve the PVRs, as can be seen in figure 7.5.

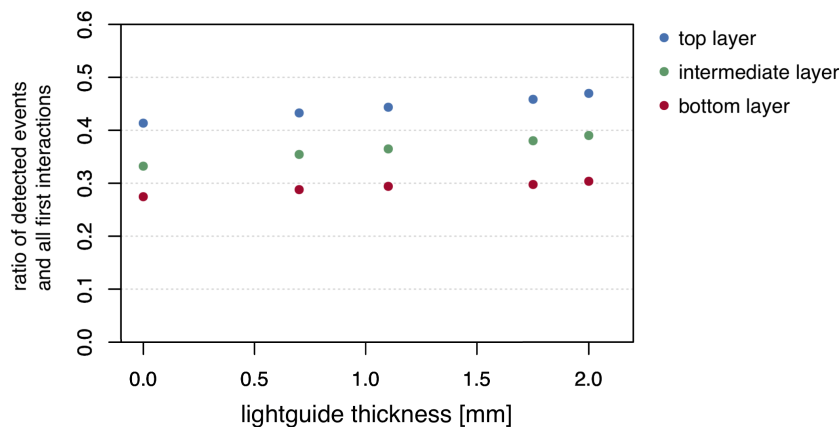


FIGURE 7.6: Ratio of the number of detected events and the number of all first interactions of the high-energy photons, shown separately for each crystal layer. The ratio increases with lightguide thickness.

A second reason for the unbalanced number of events lay in the equal height of the three crystal layers²: analysis of the GATE actor data showed that 59 % of the first interactions took place in the top layer, 28 % in the intermediate layer, and only 13 % in the bottom layer. Therefore, the different magnitudes of the determined PVRs become comprehensible.

Due to the high amount of events in the top layer, the corresponding crystal signals could always be separated from the neighbouring ones ($PVR > 12$). Therefore, the focus should be on the bottom and intermediate layers. In these layers, the best PVRs were achieved with a 1.1 mm lightguide, although the results of the 0.7 mm lightguide were identical under consideration of the standard errors.

If there was no lightguide at all, a flood map could still be recorded— according to the simulation— although there would not be any gain of information or performance expected. However, this requires partial transparency of the reflector between the

²Section 8.5 will present measurements, where the three crystal layers will have different heights to balance the number of interactions in each layer.

individual crystal pixels (as assumed in chapter 6). A corresponding measurement without any lightguide could provide evidence.

7.3 Comparison to Measurements with Different Lightguides

Methods

For the corresponding measurements, the same DPC chip has been used multiple times along with different lightguides. Once a measurement was completed, the lightguide was separated from the DPC tile by carefully sliding a 15 μm thick steel foil through the glue layer between lightguide and tile. In the same way, the glued LYSO stack has been removed as one piece from the lightguide. Therefore, multiple lightguides were measured on the same DPC tile along with the same scintillator stack.

The evaluation procedure was the same as for the simulated flood maps. Only in case of the flood histogram with a 2.0 mm lightguide, the division into horizontal slices (as shown in figure 7.4a) had to be adapted manually for two of the four dSiPMs, due to inaccurate positioning of the crystal array on the DPC tile. Apart from that, all other evaluations have been done in an automated way.

Results

Figure 7.7 shows the peak to valley ratios for the measured flood histograms with different lightguides. The result of a measurement without any lightguide (at 0.0 mm) is already contained in this figure—detailed information will follow in the next section. For a thickness of 1.1 mm, the best spot separability could be achieved: 2.84 ± 0.15 for the bottom layer, 3.45 ± 0.11 for the intermediate layer and 4.80 ± 0.14 for the top layer. However, under consideration of the standard errors, the results for the bottom layer were identical for 0.7 mm and 1.1 mm. The higher PVRs for the intermediate and top layers speak in favour of 1.1 mm, though. With thinner and thicker lightguides the PVRs were decreasing.




FIGURE 7.7: Mean PVRs (\pm standard errors) of the four crystals from each crystal layer, taken from measured data.

Discussion

In general, the measured PVRs were more similar to each other and smaller than the simulated PVRs. The mismatch of the simulation was caused by the imprecise surface definitions, that have been discussed in chapter 6 (roughness, reflectivity and transmission of the 3M ESR vikuiti, and the transition from one crystal layer into the neighbouring one(s)). In that chapter, the surface parameters have been adapted in such a way that the flood histogram pattern and energy spectra were in accordance with the measured ones. The relative spot intensities had not been considered at that time, because it would have required an even higher amount of computation time to test that many parameter configurations. Moreover, it is questionable if the UNIFIED model would be capable of reproducing the measurement data more precisely.

Although *Surti et al.* used a different detector geometry [123] (GSO crystals with 4 mm pitch, PMT readout), the results showed similarities: lower PVRs were found in measurements compared to simulations with the UNIFIED model, since the used reflectivity value of 97.5 % for the reflectors could not be matched with the true one in a measurement. This is stated to be an important reason for the deviations. Moreover, *Surti et al.* have used an uncollimated source in the measurements, while a collimated source has been used in the simulations. The source collimation is also considered to lead to an overestimation of PVRs in the simulations. Since *Surti et al.* focussed more on the difference between slotted and non-slotted lightguides (instead of extensively investigating multiple lightguide thicknesses), no detailed comparison is possible.

Measurements without any Lightguide

Additional to the measurements with different lightguides, one further measurement without any lightguide has been done. The resulting flood histogram is shown in figure 7.8.

Results

All crystal signals were identifiable in the flood histogram, and even the signal of the bottom layer was clearly visible and could be distinguished from the neighbouring spots ($PVR > 2$).

When comparing the flood histograms of measurements with different lightguides, differences in the spot pattern could be observed. Figure 7.9 compares the top right dSiPM signal in measurements with different lightguide thicknesses.




FIGURE 7.8: Flood histogram for a measurement without any lightguide. Despite the missing lightguide, all crystal signals can be clearly identified in the centre of the DPC tile.

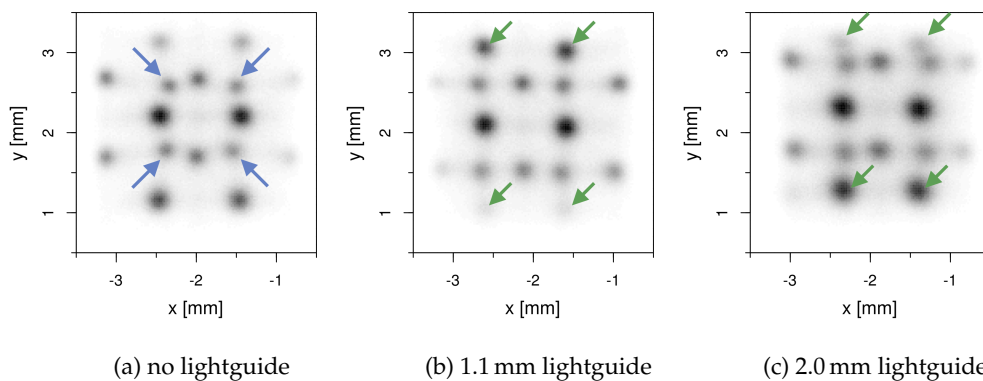


FIGURE 7.9: Flood histograms of a single centre dSiPM for measurements with different lightguides. Not using any lightguide influenced the spot positions of the bottom crystal layer, as indicated by blue arrows in (a). With increasing lightguide thickness, the spot size became larger. Since the crystal array had to be removed and glued to different lightguides for these measurements, a small shift in y direction caused different spot intensities of the top layer in figure (b) and (c), marked with green arrows.

Using no lightguide at all shifted the spots belonging to the bottom layer towards the pixel centre, as indicated by the blue arrows in figure 7.9a. These shifted spots destroyed the regular pattern, that was visible when using a lightguide (e.g. 1.1 mm in figure 7.9b). This effect was not that strongly pronounced for the intermediate layer, since the bottom crystal layer already acts as a lightguide. With increasing lightguide thickness, the pattern and the spot distances did not change noticeably. However, the spots became larger in diameter, so that the PVRs decreased accordingly. In figure 7.9b and 7.9c, the spot intensities of the top crystal layer differ. These spots are located at $y = 1.0$ mm and $y = 3.1$ mm in b), and at $y = 1.25$ mm and $y = 3.2$ mm in c). The intensity variations were caused by a small shift of the scintillator array during the

glueing process. In order to measure different lightguide thicknesses, the scintillator stack and lightguide were removed from the DPC tile after each measurement. The next lightguide was glued to the DPC tile, and in a second step the crystal stack was aligned on the new lightguide by means of the 3D printed stencils. The achievable precision of this alignment was limited, so that the small shift in y direction led to different spot intensities in the flood map.

Discussion

The fact that even the bottom layer signal was clearly visible in the flood map confirmed that the spread of scintillation light from the bottom layer was sufficiently large to trigger the full neighbourhood of the hot pixel. This is only possible, if the 3M ESR vikuiti between the crystals is partly transparent, since the influence of the DPC tile cover glass (100 μm thick) is negligibly small. The result of the corresponding simulation can be found in section A.4.

If it was as reflective as indicated by the manufacturer ($> 98\%$ [121]), the scintillation light could not distribute horizontally inside the bottom layer. Instead, it would have to be reflected to the intermediate layer and back into a neighbouring crystal of the bottom layer, in order to spread in x direction. In order to spread in y direction, the optical photons would even have to be reflected up to the top crystal layer and back to the bottom layer, because bottom and intermediate layer are aligned in y direction. But this combination of back and forth reflections between the layers have very low probability to cause such high spot intensities from the bottom layer. It is very likely that the partial transparency of the 3M ESR vikuiti also allowed horizontal spreading of scintillation light inside the crystal layers. This assumption would match the results of previously shown simulations with different vikuiti reflectivity values (see section 6.3). Moreover, it would be in good agreement with the findings from *Loignon-Houle et al.* [150], where the partial transparency of this particular reflector material has been observed in bonded crystal arrays.

Compared to the work of *Ito et al.*, where the concept of four staggered crystal layers for DOI encoding was presented in 2007 [108], the single crystal signals were much clearer separated. This applied both in the centre and at the outer edges of the DPC tile. However, *Ito et al.* used PSPMTs as photodetectors. Methods to separate overlapping crystal signals will be presented in the next chapter.

Summary

PET detectors, which make use of pixelated crystal arrays and the lightsharing principle, use a lightguide between crystal array and photodetector. The lightguide thickness has a crucial influence on the scintillation light distribution on the sensor [123], so that a variation of its thickness can be used to optimise the light spread on the sensor. Lightguide thicknesses between 0.7 mm and 2.0 mm were tested in this chapter, both in simulations and in measurements. As a measure of separability, the PVR of neighbouring spots was analysed in the centre of the flood map. Both studies showed the best spot separability for 0.7 and 1.1 mm thick lightguides. Additionally, measurements and simulations were acquired without using any lightguide, which still led to the previously observed spot pattern. In case of the intermediate and top crystal layer, the underlying crystal layers act as lightguides themselves. But for the bottom crystal layer, this finding supported the assumption that the 3M ESR vikuiti between the bonded crystals was partly transparent.

In the final detector modules, a lightguide of 1.1 mm thickness will most likely be used, due to the following reasons: the measurements showed slightly better results for the 1.1 mm lightguide in the top and intermediate layers, the glass sheets of 1.1 mm thickness are more easily available and to handle than those of 0.7 mm thickness.

Chapter 8

Evaluation of Reference Response for Positioning Algorithms

In the previous chapter, the focus has been on the acquisition of a preferably clean flood map, where all single crystal signals can be well separated. This is crucial for the calibration of a detector block. The raw data is separated with a nearest cluster net (Voronoi net), or it is used as a training data set, when algorithms like Maximum Likelihood (ML) or neural networks are involved during event positioning. A short overview over these methods has been given in section 2.2.6. For this project, Maximum Likelihood-based event positioning has been chosen. Therefore, the general workflow of calibration and event positioning with ML will be given in this chapter. Moreover, different potential reference responses will be evaluated.

8.1 Maximum Likelihood Event Positioning: Workflow

Calibration

Each detector module undergoes the same calibration procedure:

1. The detector modules in a scanner are irradiated with a source of 511 keV photons, so that a flood histogram can be calculated from the gain and saturation-corrected data. Figure 8.1a shows an exemplary (simulated) flood map of a detector module.
2. The flood map is divided into 16 regions, which correspond to the 16 dSiPM signals. Afterwards, a clustering algorithm is applied to separate the spots, which represent the signals from individual crystals. Figure 8.1b shows the clustering result for an accumulation point of Anger positions, that is marked in figure 8.1a (blue arrow).
3. Once all the clusters are found, the reference response matrix can be calculated for each spot. It contains the normalised mean light distribution $c_{i,k}$ that is measured by dSiPM k after a scintillation event occurred in crystal i . In other words, the mean light distribution (i.e. the mean number of counted photons on the dSiPM pixels) is calculated for all events contained in a cluster.

With 204 spots in the flood map and 16 dSiPM pixels, $c_{i,k}$ is a matrix of dimension 204×16 . Some of the 162 crystals give rise to two spots in the flood map, when they are located above a bond gap. In these cases, two separate light distributions are calculated, but assigned later on to the same crystal. An example for the reference response matrix can be seen in figure 8.1c, which shows the reference response matrix for the crystal indicated in figure 8.1a.

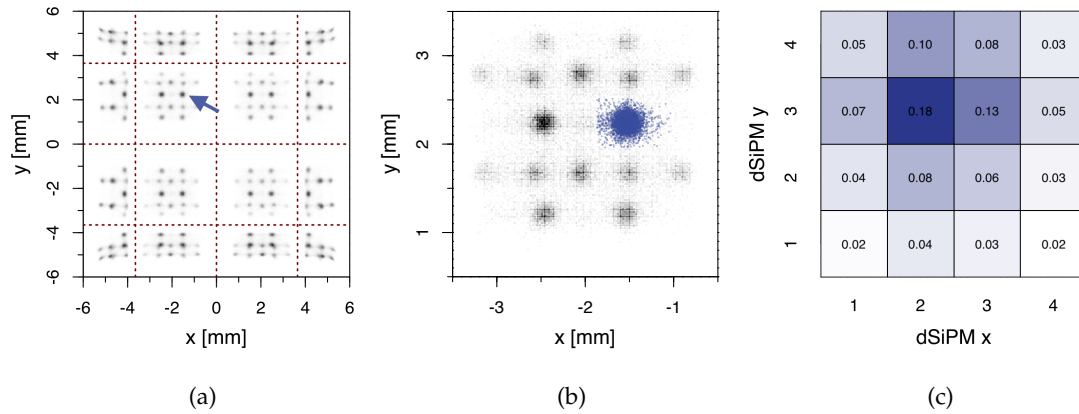


FIGURE 8.1: Steps of the calibration process: the flood histogram is divided into 4×4 regions, indicated by red lines in (a), to simplify the clustering of single crystal signals. The clustering result for the crystal signal marked in (a) is shown in figure (b). Once the crystal signals are separated, the reference response matrix can be calculated for each crystal. Figure (c) shows the reference response matrix for the crystal signal from figure (a).

Event Positioning

The basis for event positioning is the measured light distribution \mathbf{q} on the DPC tile. Since the normalised mean light distribution (reference response matrix) $c_{i,k}$ is known for each spot from the calibration step, the likelihood \mathcal{L}_i can be calculated with equation 2.10:

$$\mathcal{L}_i(i, S|\mathbf{q}) = \sum_k q_k \ln [c_{i,k} S] - c_{i,k} S \quad , \quad (2.10)$$

where S is the sum of all detected photons: $S = \sum_k q_k$. The best matching reference light distribution can be selected by determining the maximum likelihood value:

$$i^{\text{ML}} = \arg \max_i \mathcal{L}_i(i, S|\mathbf{q}) \quad . \quad (2.11)$$

The crystal which was most likely hit can be determined by means of a lookup table, which assigns all spot indices i to the corresponding crystals. Double spots, which originate from the same crystal, are assigned accordingly in this lookup table. Figure 8.2 shows the single steps of event positioning. The full iterative procedure including the update of energy and re-calculation of likelihood has been described in section 2.2.6 (based on reference [58]).




FIGURE 8.2: Steps of the event positioning: the likelihood is calculated from the acquired light distribution of a scintillation event (a) and all reference response matrices. The highest likelihood value is achieved for the best matching light distribution from the reference response matrix, shown in (b), and the corresponding spot index in the flood histogram (c). A lookup table assigns crystal layer and crystal location to the selected spot.

8.2 Simulations: Using the Ground Truth Information

For the used detector geometry, the clustering of all spots in the flood histogram is a crucial step. Especially overlapping spots at the outer edges require advanced clustering techniques, which make use of more parameters than only the 2D coordinates, e.g. energy or other light distribution parameters. In this chapter, the influence of different clustering approaches (reference responses) will be investigated by means of simulated data. Before the clustering approaches itself are discussed, the simulation and methods to make use of the ground truth information will be presented.

8.2.1 Simulation Setup

The simulation setup was similar to the one described in section 6.3 (figure 6.3): the three crystal layers of 9 mm height had a reflective coating, which was 70 % reflective between the single crystals, and 96 % reflective at the interfaces of crystal and air. A line source of 1 MBq irradiated the stack from the top, while the emission angle was confined to only emit 511 keV photons in the direction of the crystal stack. The acquisition time of 600 s yielded 2,706,998 events with application of the DPC digitiser from section 6.4. Generation of noise or crosstalk was omitted during the simulation, since the purpose of this simulation was the generation of scintillation events, their corresponding light distributions on the DPC tile, "measured" signals and the ground truth.

The ground truth was acquired with a GATE actor: the coordinates and energy deposit of a gamma photon's first interaction in the crystal block was saved for each event. Therefore, for each light distribution or simulated DPC signal, the corresponding crystal and crystal layer were known. This allowed one to compare

the result of a positioning algorithm to the true crystal and crystal layer ID. In this way, the influence of different reference responses was investigated. But before this could be done, methods to separate overlapping crystals were developed by means of the ground truth information, which was acquired during the simulation.

8.2.2 Distinguishing Signals from Different Crystal Layers

Both measured and simulated flood maps continuously showed overlapping spots at the edges, where the Centre of Gravity (CoG) positions were not sufficient to separate the crystal signals from each other. Apart from the CoG positions, further statistical moments of the distribution were investigated, since they contain characteristic information on the distribution. In theory, the light distribution could even be reconstructed by means of its moments– if certain conditions are met (existence of the moments, class of possible distribution functions is restricted, etc.)[157]. However, it was not intended to reconstruct the light distribution from its moments, but to reveal differences in the moments in order to separate the signals originating from different crystal layers¹.

Apart from the CoG positions, which correspond to the first moment, the energy (zeroth moment) and the standard deviation (derived from the second central moment) were investigated. The basic differences will be explained for crystals in the centre of the flood map, and transferred later on to overlapping spots at the edges.

Energy

The energy of a scintillation event was calculated as the sum of the counted photons². The aforementioned equation

$$Q = \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} \quad (2.4)$$

was applied to the hot pixel and its neighbouring pixels, i.e. in total 4, 6, or 9 dSiPMs. In general it could be observed that events from different layers have distinctive energies. Events from the top layer occurred most often, but had the lowest energy due to the large lightspread in the underlying crystal layers. The opposite case applied for the bottom layer, where most of the scintillation light was directly detected by the underlying dSiPM pixel, instead of being spread in all three layers. This lead

¹The analysis of moments is of special interest in detector designs with monolithic crystals, because they contain valuable information on the shape of the light distribution, so that this information can be used in the positioning algorithm [158], [159].

²To convert the number of counted photons to an energy, a Gaussian distribution is used to approximate the full energy peak of the count spectrum. The peak position is relating the number of counted photons to the measured energy (511 keV) for each crystal. In this chapter, different energy levels will be mentioned for events from different crystal layers. Strictly speaking, only the number of detected photons per amount of deposited energy varies for each layer. For better readability, the terms "energy" and "number of detected photons" will be used synonymously.

to fewer events with a higher measured energy. Events from the intermediate layer could be found in-between, both regarding frequency and energy level.

This effect is visualised in figure 8.3, which contains the energy spectra of events from different spots in the flood map: one spot was selected from each crystal layer and is marked in figure 8.3a. The energy spectra of events within these rectangles are shown in figure 8.3b, but the events are filtered by means of the ground truth information: only those events, which originated from the respective crystal layers, are shown. In case of the blue rectangle, only those events were selected, where the high-energy photon first interacted in the top crystal layer. Events originating from the intermediate and bottom layer were selected accordingly for the green and red rectangle.

The energy spectra clearly show three peaks, which could be assigned to the three crystal layers. The described observations were in good agreement with findings from *Chung et al.*, who observe decreasing energy peaks from bottom to top in a four-layer PET detector [55].

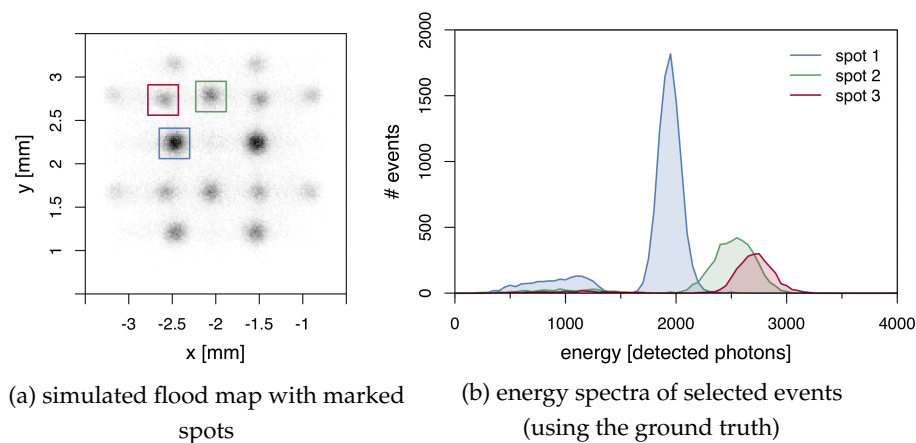


FIGURE 8.3: Energy spectra of single spots from the three crystal layers: when the DOI information was used to filter events, the energy spectra showed a distinct energy level for each crystal layer. Events from the bottom layer (red) had the highest energy, while those from the top layer (blue) had the lowest energy.

If the ground truth information was not used, and all recorded events inside the rectangles were evaluated, this resulted in the energy spectra shown in figure 8.4a. In contrast to figure 8.3b, the curves of spot 2 and 3 had an additional peak around 1950 photons: these events originated from the top layer, but their Anger positions matched a spot from the intermediate or bottom layer. This can occur after a far-Compton event, or due to a missing dSiPM signal, which can be caused by dead time of single detector pixels. In these cases, the Anger positions are shifted away from the expected spot in the flood map.

However, these events could be identified by the differing energy level. Figure 8.4b shows energy spectra from a measurement (same spots selected, 1.1 mm lightguide).

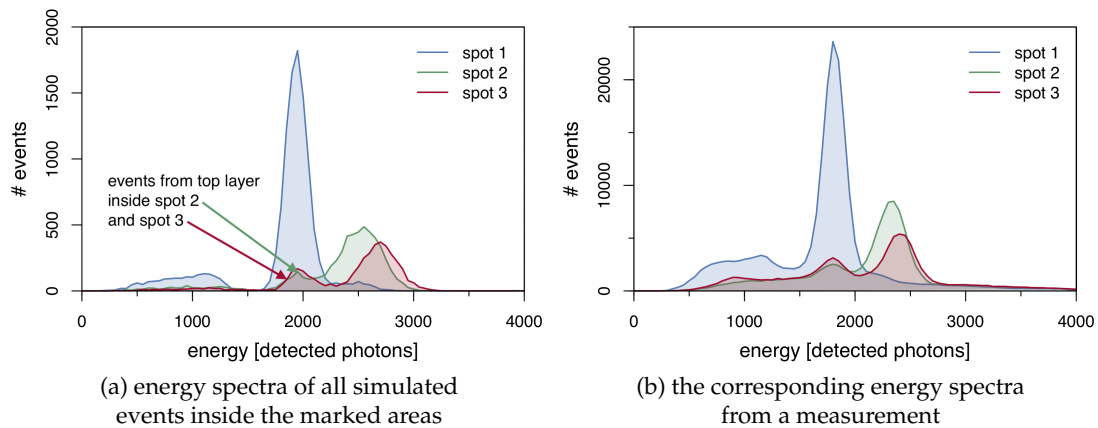


FIGURE 8.4: All events inside the marked areas from figure 8.3a are analysed: the energy spectra from simulation (a) and measurement (b) show the same formation of peaks. These peaks can be used to identify events from different crystal layers.

The same peaks could be observed in the measurement, so that energy thresholding appeared to be an option for separating overlapping spots.

While this distinction was straight-forward in the centre of the flood map, the case became more challenging at the edges, where clusters overlapped: due to the cut-off light distribution at the edges of the sensor, the energy levels were not as well separated as in the centre of the flood map. Apart from the Anger position and the energy of an event, further parameters should be considered to separate the clusters.

Light Distribution

An additional parameter for separation of overlapping spots could be the light distribution, that is detected by the dSiPM pixels for each event. Figure 8.5 shows the light distributions for the three spots, which have been introduced in figure 8.3. Again, only those events from the respective crystal layers were selected. To highlight the differences, a one-dimensional representation of the measured light fractions was chosen. The numbering scheme for the dSiPM pixels is shown in figure 8.5a.

The detected light fractions in figure 8.5b differed for each crystal layer. Therefore, the ratio of these light fractions could be used to distinguish the events. Especially the fractions detected by the direct neighbours of the hot pixel (in this case dSiPMs 2, 5, 7 and 10) could be used to calculate these ratios (e.g. ratio of dSiPM 6 and 10 to distinguish bottom and intermediate layers, or ratio of dSiPM 6 and 7 for top and bottom layers).

Figure 8.6 contains the light distributions of events from the three spots, but this time taken from measured data. The curves were similar to the simulated ones, apart from slight deviations in the detected light fractions on the hot pixel (dSiPM 6) and pixel 9 and 10. The concept for distinction of the events could still be applied, since the ratios of the detected light fractions appeared to be specific for each crystal layer. However, suitable dSiPM pairs have to be selected for this approach.

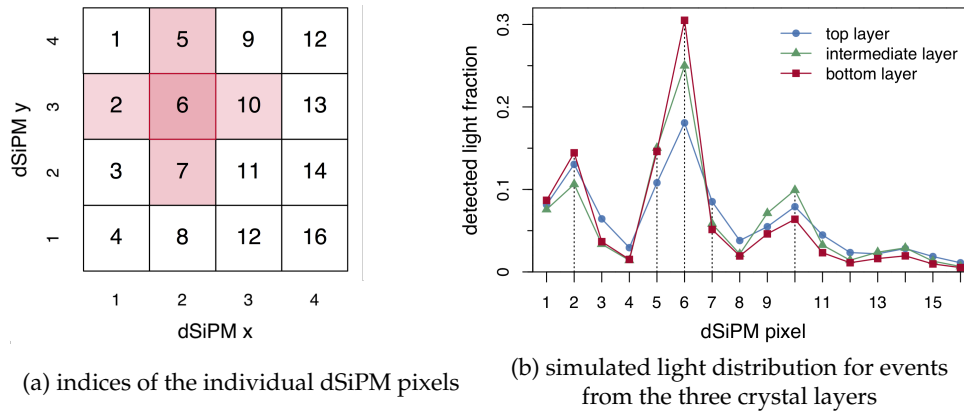


FIGURE 8.5: Simulated light distribution for events from the three crystal layers in the marked spots from figure 8.3. The highest fraction was detected on dSiPM 6 and its direct neighbours, which are indicated by dotted lines. The different ratios could be used to distinguish events from the three crystal layers.




FIGURE 8.6: Measured light distributions for events from the above defined spots. Compared to the simulation, the light fractions measured by dSiPM pixels 6, 9 and 10 differed slightly (marked with circles). However, the concept for event distinction could still be applied.

Standard Deviation

Another parameter that might be useful for distinction of crystal signals is the standard deviation of the light distribution, which is the square root of the second central moment:

$$\sigma_x = \sqrt{\frac{1}{Q} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} (x_i - a_x)^2} \quad \text{and} \quad \sigma_y = \sqrt{\frac{1}{Q} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} q_{i,j} (y_j - a_y)^2} \quad . \quad (8.1)$$

Analogous to energy and Anger positions (a_x, a_y) , only the hot pixel and its neighbouring pixels are considered. The standard deviation is relative to the Anger position of the event, and therefore a measure for the spatial extent of a spot. Since the light spread and consequently the detected light distribution on the DPC tile differed for the three crystal layers, a change in the standard deviation was expected.

The standard deviations of simulated data are presented in figure 8.7. In figure 8.7b, all events from the marked areas were used: in case of spot 1 and spot 2, only one distinct accumulation point was visible, whereas spot 3 had a second (weaker) accumulation point that overlapped with spot 1. When comparing to the ground truth information, which has been used in figure 8.7c, it becomes clear that the second point cluster contained events, which originated from the top layer, but whose Anger positions matched spot 1 (this has already been observed in the energy spectra).

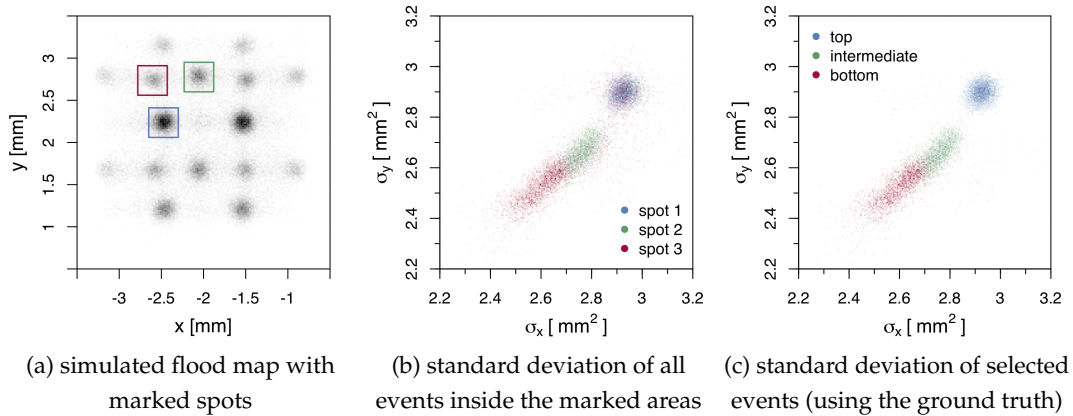


FIGURE 8.7: Standard deviation of the Anger positions, taken from the simulation. Figure (b) shows the standard deviation of all events inside the marked areas of figure (a), whereas figure (c) presents the standard deviation of events from the respective layers (selected with ground truth).

Similar patterns could be observed in the measured data, as shown in figure 8.8. Analysing all events from the marked areas in figure 8.8a lead to two point clusters for spot 3, and one cluster for spot 1 and 2 (figure 8.8b). After energy filtering (figure 8.8c) only three clusters remained, as expected from the ground truth in the simulated data. The standard deviation could therefore be an additional parameter for clustering.

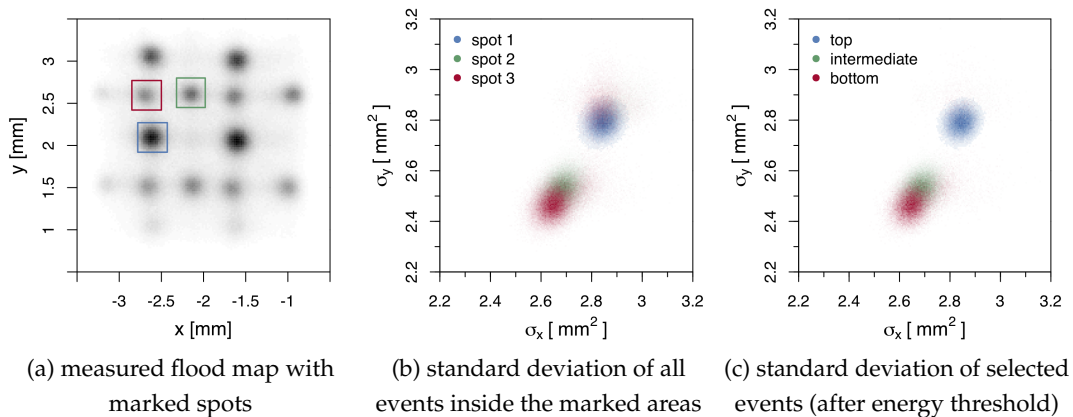


FIGURE 8.8: Standard deviations of the Anger positions from the measured data. Figure (b) shows the standard deviation of all events inside the marked areas of figure (a), whereas figure (c) presents the standard deviation for energy-filtered events (presumably from the respective layers).

8.2.3 Application to Overlapping Spots

It has been shown that parameters like energy, light distribution and standard deviation were specific to crystals from the three layers. In this section, their use for separation of overlapping spots will be demonstrated. Three cases could be distinguished, as shown in figure 8.9: at the top and bottom edge of the flood histogram, signals from the top and bottom crystal layers overlapped (blue marked areas). At the left and right edge of the flood map, signals from the bottom and intermediate layers overlapped (red marked areas). Moreover, the four outermost crystals of the top layer overlapped with the intermediate layer (green marked areas).




FIGURE 8.9: The coloured areas in the flood histogram indicate those spots, which contained two crystal signals: bottom and intermediate layers are marked in red, top and bottom layers are indicated in blue, and top and intermediate layers are marked in green.

Top and Bottom Layers Figure 8.10a shows a region of the simulated flood map, where the signals of top and bottom layers overlapped. The two signals inside the marked spot were separated by means of the previously described parameters. As can be seen in figure 8.10b, the energy spectrum of the simulated data showed two distinctive peaks, which could be assigned to the two crystal layers. The ground truth signal of both layers is shown in red and blue. Moreover, the standard deviation contained two well separated point clouds (figure 8.10c), so that a separation of both signals was feasible, if energy and/or standard deviation were included as parameters in the clustering process.

Figures 8.10d-8.10f confirm these findings for measured data: the two overlaying spots could be easily separated by considering energy and/or standard deviation as parameters in the clustering process. The light distribution, i.e. the measured number of photons on each dSiPM, did not have to be included in the clustering process, since energy and standard deviation already showed distinctive signals and were easier to consider in the clustering process (a suitable pair of dSiPMs would have to be chosen at first to integrate the light distribution into the clustering).

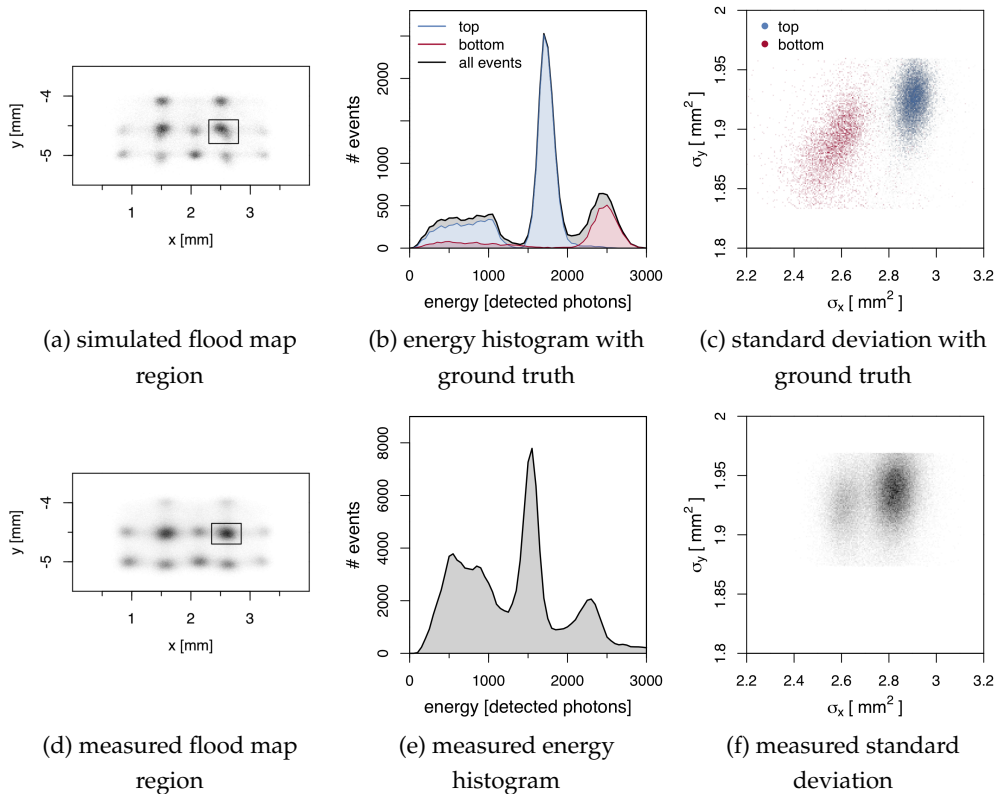


FIGURE 8.10: A single spot, which contained overlaying crystal signals of the bottom and the top layers, is shown in the simulated (a) and the measured flood map (d). The energy spectrum (b,e) and standard deviation (c,e) were used to separate the signals.

Top and Intermediate Layers Figure 8.11 shows the same analysis for a flood map region, that contains overlaying spots of the top and intermediate layers. The signals can easily be distinguished by the strongly pronounced energy peaks in figures 8.11b and 8.11e. In contrast to the previous case, the standard deviation is not clearly separated in figures 8.11c and 8.11f, and therefore not useful for separation of the crystal signals. When examining the measured flood map (figure 8.11d) closely, the two spots do not exactly overlap, in contrast to the simulated flood map. By including the energy into the clustering process, both crystal layer signals can already be distinguished.

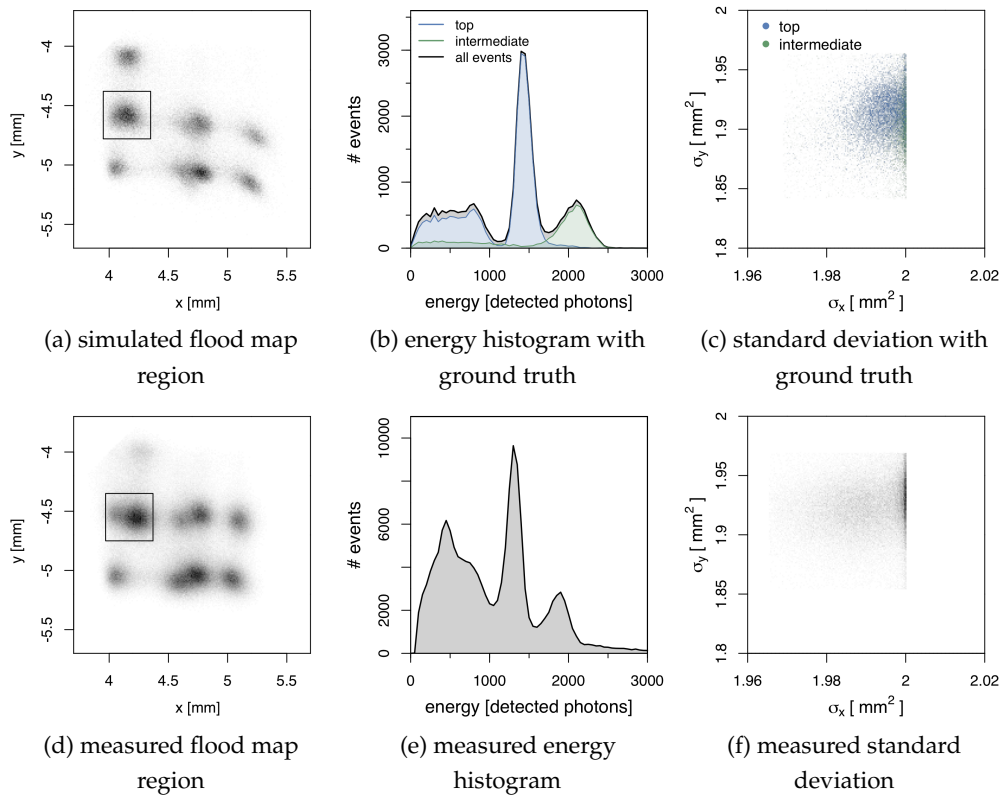


FIGURE 8.11: A single spot, which contained overlaying crystal signals of the top and intermediate layers, is shown in the simulated (a) and the measured flood map (d). The energy spectrum (b,e) could be used to separate the signals, while the standard deviation turned out to be not well separable.

Bottom and Intermediate Layers In case of the bottom and intermediate crystal layers, the signal separation appeared to be the most challenging. Figure 8.12 summarises the analysis results for both simulation (a-c) and measurement (d-f). The ground truth of the simulation revealed two energy peaks, that completely overlapped (figure 8.12b), and accordingly only one single energy peak could be observed in the measurement (figure 8.12e). While the standard deviation was well separable in the simulation (figure 8.12c), the separation was worse in the measured data, as can be seen in figure 8.12f: the two point clouds of the standard deviation overlapped and were hardly separable. At the same time the standard deviation was the only parameter, which could be used for separation, since also the light distribution did not reveal any distinct differences (figure 8.13a). Figure 8.13b shows exemplarily the ratio of two dSiPM light fractions (number 15 and 16), since these dSiPMs measured different light fractions for the two crystal layers. Nevertheless, the difference in the light fraction ratio was too small to be detectable in the measurement, so that the separation had to be based on the standard deviation.

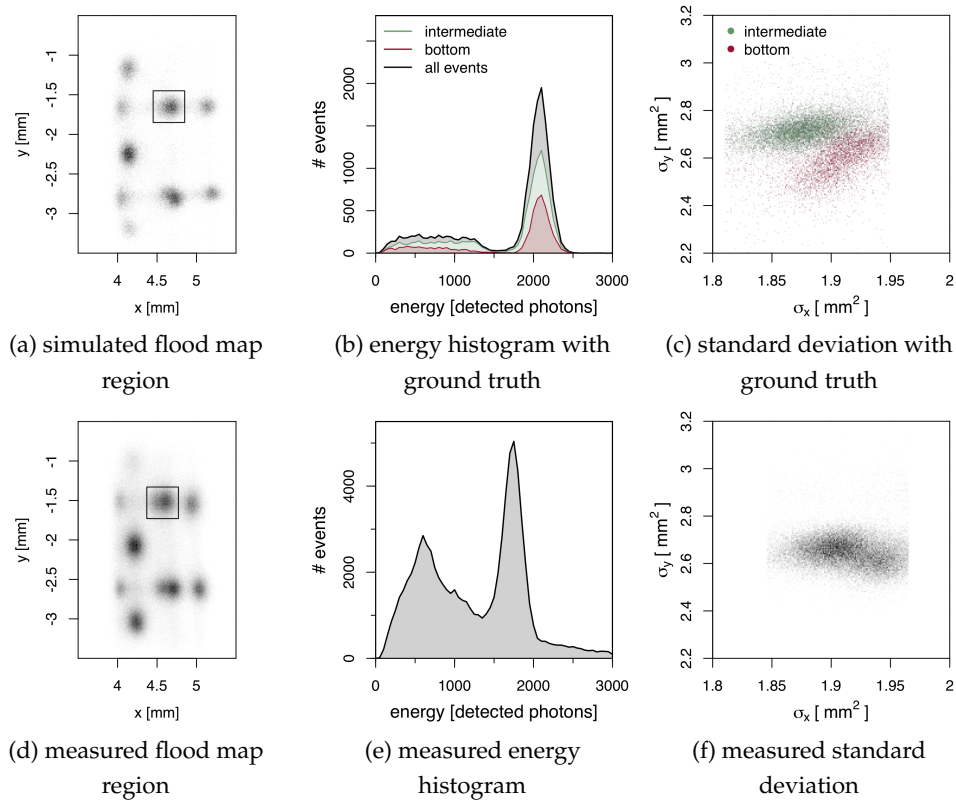


FIGURE 8.12: A single spot, which contained overlaying crystal signals of the bottom and intermediate layers is shown in the simulated (a) and the measured flood map (d). The energy peaks of both layers completely overlay (b,e). The standard deviation could be used to separate the signals in the simulated data (c), while this was more difficult for the measured data (f).

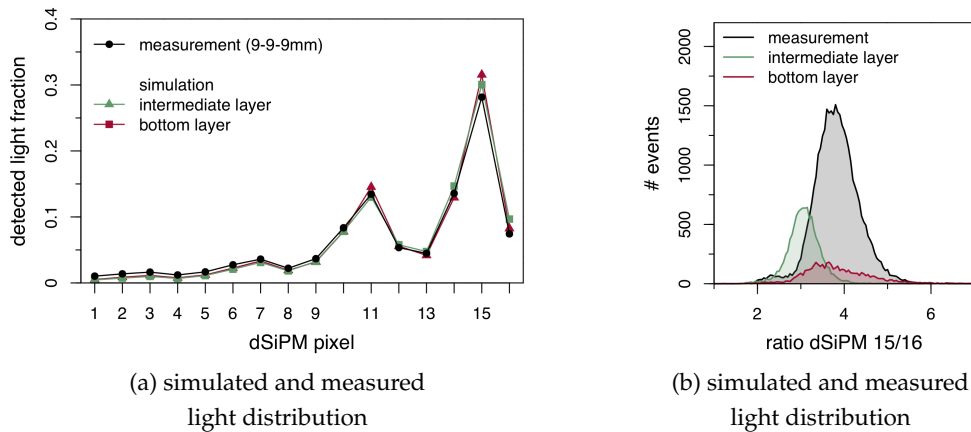


FIGURE 8.13: The light distribution for events from the marked spot in figures 8.12a and 8.12d: the ground truth information from the simulation showed extremely similar distributions for both layers, which made it impossible to separate the spot by making use of this information only. The measured light distribution was in good agreement with the simulated ones. Part (b) shows the ratio of dSiPM 15 and 16: while small differences were visible in the simulation, there was no usable information in the measurement.

In the simulated data, all overlaying spots could be separated. This was not the case for the measured data, where overlaying spots of the bottom and intermediate layers yielded too close signals to be separable at the outer edges.

Discussion

In general, the signals of the bottom and intermediate crystal layers have shown very similar energy peaks, which could already be observed in figure 8.4. The difference of the top layer signal compared to the other two layers was striking in these analyses. This could be related to the different surface finishes of the three layers: while the lateral surfaces of the top and bottom layers were polished, the intermediate layer laterals were "as cut" (see figure 6.1). It is difficult to judge how far the surface finishes influenced the spot energies, and if the different finishes caused the similarity of bottom and intermediate layers. It is commonly known that both crystal dimension and surface finish have a strong influence on the light output, and therefore the measured number of photons [92], [160]. An interesting observation was made by *St. James et al.*, who found a "directionality" in flood maps, when using the 3M ESR vikuiti: only in one direction of the flood map, all the crystal signals could be separated (independent from the alignment on the sensor). In the other direction, crystals at the edges could not be identified. *St. James et al.* make the manufacturing process responsible for this behaviour, since the crystals are covered with long reflector stripes in one direction, and shorter ones inside the slices in the other direction. Light crosstalk in the corners could lead to the observed directionality [161].

In section 8.5, the overlapping spots will be investigated again with the final geometry, which has crystal layers of 10, 9, and 8 mm thickness instead of 9 mm thickness for all three layers. Moreover, the same surface finish will be used for the three crystal layers. In the following section, the simulated calibration data will be separated by means of different clustering methods.

8.3 Clustering Methods

For separation of the spots in the calibration data, three clustering methods were evaluated, mainly differing in complexity. Each method was applied after a plain pre-segmentation, which divided the flood map into 4×4 regions. The methods will be described briefly, being followed by a comparison of the clustering results at the edges of the DPC tile.

1. Histogram-based segmentation

The first method was not per se a clustering method, but resembles more a segmentation. The general idea was to separate the spots in the flood map by analysing the projections in x and y direction.

For each region, the number of visible (i.e. not overlapping) clusters in x and y directions was provided, so that the specified number of minima and maxima was determined in density estimates of the 1D histograms. The regions were accordingly divided into rectangles, which contained either 0, 1 or 2 clusters. Figure 8.14a shows the segmented areas for the 5×3 clusters, that were to be determined. In a second step, the main energy peak was identified in each spectrum, and events above or below the peak were discarded by finding local minima in the energy spectrum above and below the main energy peak. The number of clusters inside each section was stored in a LUT, so that either one or two peaks were preserved from the energy spectra.

In figure 8.14a, the result of event separation is displayed. The cluster position was calculated as the mean of these energy-filtered events inside the rectangular area, while the energy was determined as the median value, to be less sensitive to outliers. Moreover, the mean light distribution was saved for each spot. In case of overlapping clusters, the events were separated by the two energy peaks, before cluster position and energy were calculated. Since this strategy was not feasible for overlapping spots of the bottom and intermediate layers, the same cluster centres and mean light distributions were assumed for the affected crystals.

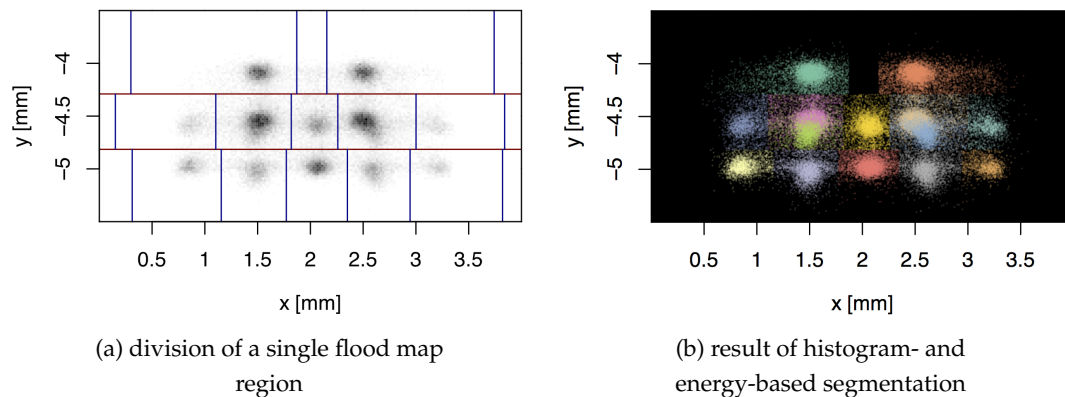


FIGURE 8.14: A single region of the flood map is shown in figure (a). According to the density estimates of the 1D histograms, the region was further divided into areas, which contained 0, 1 or 2 spots. By means of energy thresholding the clusters could be defined (b).

2. Fuzzy C-Means

The second method that was used for clustering is fuzzy c-means. This algorithm has been developed as a modification of k-means clustering [162]: while k-means clustering divides the data set into k discrete clusters, fuzzy c-means assigns each data point probabilities to belong to each of the "fuzzy" clusters, which results in overlapping clusters. How much the clusters fade into each other (*fuzziness* of the clusters) depends on the fuzzy factor, which is usually determined empirically. In

this case, a fuzzy factor of 2 has been chosen; the generally recommended range lies between 1.5 and 2.5 [163].

It is generally known that the outcome of fuzzy c-means (and also k-means clustering) strongly depends on the initialisation of the start parameters, such as the number of clusters or the cluster centres [164]. For this reason, the cluster positions from the histogram-based method were used as start parameters. Multi-dimensional clustering has been attempted, e.g. by using the event positions and energies, but did not yield reasonable results (even after scaling the parameters to similar ranges). Therefore, only the 2D positions were used for clustering, and energy thresholds were used afterwards to remove events below the energy peaks (far Compton events and noise). In contrast to the aforementioned method, all clusters within a flood map region were determined at the same time, and not sequentially.

Figure 8.15 contains a plot of the start parameters and the final result after fuzzy c-means clustering and energy thresholding.

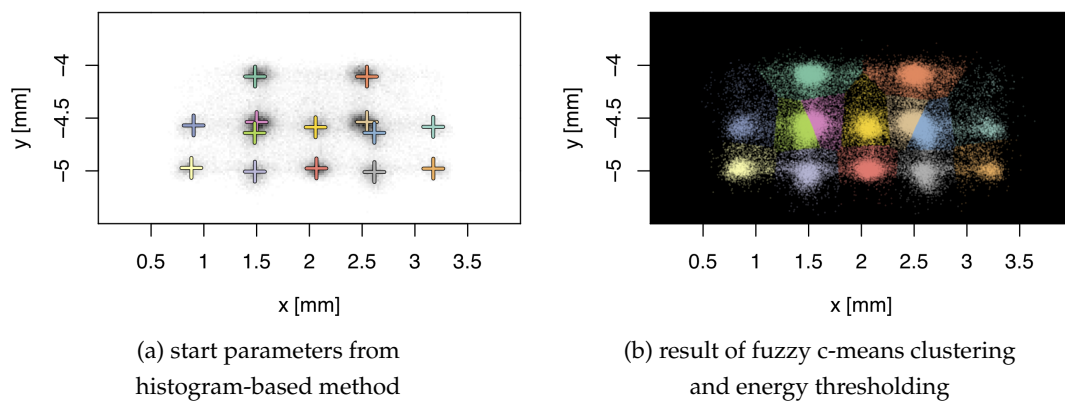


FIGURE 8.15: The cluster centres, which have been determined with the histogram-based method, were used as start parameters (a). Figure (b) shows the result of fuzzy c-means clustering and energy thresholding.

3. Model-based Clustering with Mclust

The third method for clustering was model-based and made use of the *mclust* package for *R* [134], [165]: *mclust* assumes a mixture of Gaussian distributions to have generated the data. The best fitting set of parameters for the Gaussian distributions is determined by expectation maximisation.

For each cluster, one multivariate Gaussian distribution is assumed: the cluster is centred at the mean, and its geometry (shape, volume, orientation) is determined by the covariance matrix [166]. Within the *mclust* package, 14 different models are implemented, which assume different shapes, volumes or orientations of the clusters, e.g. model "VEE": variable volume, equal shape and orientation, and ellipsoidal distribution of the clusters.

Since there were up to 16 clusters in each flood map region, the clustering was only applied inside the histogram-based sections, which contained either 0, 1 or 2

clusters— trying to determine all clusters at once would not always yield reasonable results. The underlying model (i.e. parametrisation of the covariance matrices) was selected individually (but automatically) by using the Bayesian Information Criterion (BIC), which can be automatically calculated for the data subsets [165].

For single spots, the Anger positions and event energies were passed to the `mclust` function. In case of overlapping spots, the standard deviations and energies were used in combination with the Anger positions. The light distribution was only used in case the clustering with the previously mentioned parameters failed. All parameters were provided along with the look-up table for the number of clusters. Additional to the specified number of expected clusters, a noise cluster was determined. Figure 8.16 shows the result of model-based clustering with noise clusters.

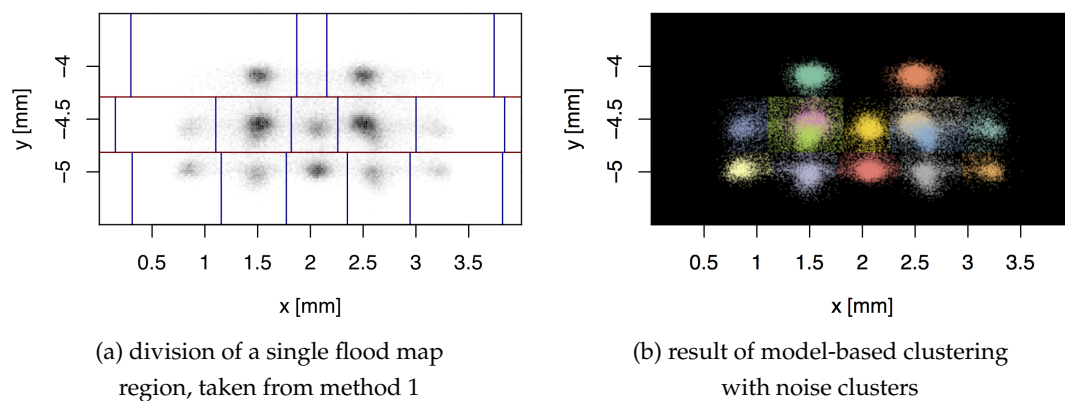


FIGURE 8.16: A single region of the flood map was divided into subsections, which contained 0, 1 or 2 spots (as in method 1, figure (a)). Model-based clustering was applied to identify the clusters and one noise cluster per subsection.

Results from Clustering Methods

In the centre of the flood histogram, the results of the three clustering approaches were comparable. Differences in the outcome became clear when regions at the outer edges of the flood map were investigated, especially at overlapping spots. The histogram-based method and model-based clustering seemed to separate overlapping clusters, while fuzzy c-means assigned one half of the combined spot to both underlying clusters (see figure 8.15b). This could be traced back to the fact, that simultaneous clustering of energy and positions was not possible with fuzzy c-means, so that the main separation was based on the Anger positions. Including energy into the clustering did not yield reasonable results, even after scaling the energy to a similar parameter range as the Anger positions. It is assumed, that the start parameters have not been defined to be already close enough to yield a usable clustering result, as this is a known problem in fuzzy c-means [164].

Since energy was the major separation criterion in the histogram-based method, two clusters were defined with the same centres, but the events were separated by using the energy peaks. Therefore, the clusters had different energies and mean light

distributions. In case of overlapping spots from the bottom and intermediate layers, where only one energy peak was visible, the same energy and mean light distribution was used. The model-based clustering approach determined two overlaying clusters with different centres, energies and mean light distributions by means of the standard deviation. In the next section, the influence of the clustering method (and therefore the calculated mean light distributions) on the positioning results will be investigated.

8.4 Influence of Response Functions on Positioning Results

After the mean light distributions had been calculated using the three clustering methods, the impact on event positioning could be determined. For this purpose, the crystal where the scintillation took place, had to be determined for each event. The resulting crystal could be compared to the ground truth from the simulation, i.e. the crystal, where the first interaction of the high-energy photon occurred. This crystal would yield the correct LOR for image reconstruction.

Methods

The simulated data set was split into two parts: 1.0×10^6 events had been used for the reference response (clustering and mean light distribution), and 1.4×10^6 events were used as a test data set. While the data for the reference response only made use of events, where the full neighbourhood of the hot pixel measured a signal, the test data set kept all events, including those with missing dSiPM signals. Acquisition of the reference response would correspond to a calibration scan, which allows for longer scan times and selection of these events. The test data set should be similar to measured data during a measurement, therefore all acquired events were analysed. Two methods for event positioning were compared:

1. Nearest Cluster

The event position was calculated by means of Anger logic from the light distribution on the DPC tile. Only the hot pixel and its 3, 5, or 8 neighbouring pixels were considered, and also events with incomplete signals were allowed. To identify the scintillation crystal, the nearest cluster centre (taken from the reference response) was selected.

2. Maximum Likelihood

The log-likelihood was calculated for the detected light distribution and each of the previously determined mean light distributions:

- a) S was the sum of all detected photons: $S^0 = \sum_k q_k$. Being proportional to the deposited energy, S was used as the energy prior.
- b) The log-likelihood was calculated with the formula

$$\mathcal{L}_i(i, S|\mathbf{q}) = \sum_k q_k \ln [c_{i,k} S] - c_{i,k} S \quad ,$$

where q_k were the detected charges on each dSiPM, and $c_{i,k}$ was the matrix of normalised light distributions (reference response matrix), which had been determined with the three different segmentation approaches. The largest likelihood value indicated the crystal, where the scintillation event had most likely taken place.

- c) The energy was updated: $S^{\text{ML}} = \sum_k q_k / \sum_k c_{i^{\text{ML}},k}$.
- d) The steps b) and c) were repeated, until $\mathcal{L}_i(i, S|\mathbf{q})$, i^{ML} and S^{ML} did not change anymore, or until a maximum number of iterations was reached.

When the same cluster centre, energy and light distribution had to be assumed for two overlapping spots (histogram-based method), and nearest cluster search or maximum likelihood positioning identified one of these crystals, either of the corresponding two crystals was selected randomly. As mentioned in section 8.2.1, the positioning results were compared to the ground truth from the simulation, so that the amount of correctly identified events could be assessed for each method.

Results

Figure 8.17 compares the amount of correctly identified crystals for the three clustering and two positioning methods. When all events were investigated, the difference between nearest cluster search and maximum likelihood positioning was around 10 % (51.7 % vs. 61.5 %). The superior performance of ML became evident when events with incomplete signal were investigated: nearest cluster search only identified 9.2 % of these events correctly, while ML positioning identified 38.2 % correctly. This finding was in line with the results of *Lerche et al.* [58], who state higher robustness of ML positioning to missing data. Moreover, they observed considerably smaller positioning errors with ML (compared to centre of gravity positioning), when less than 5 neighbouring dSiPM signals were missing in an event.




FIGURE 8.17: Comparison of different methods for clustering of the reference flood map and for event positioning.

In general, the ratio of 61.5 % correctly identified events did not appear to be very high, but an analysis of the ground truth data made this number comprehensible: the ground truth showed, that –in the ideal case– only 77 % of the events could be correctly identified. To correctly identify an event, the crystal of the *first* interaction of a high-energy photon has to be determined. But in a measurement, it is only possible to determine the crystal of the energy-weighted mean of all interactions, since the event position is calculated from those dSiPMs, which count the highest numbers of scintillation photons. For this reason, an event can only be correctly identified if the crystal of first photon impact coincides with the crystal of the energy-weighted mean of all the interactions (i.e. no far Compton scattering). This leads to a theoretical limit of 77 % for the used detector configuration.

The three clustering methods did not lead to significant differences in performance, as apparent in figure 8.17: the model-based clustering yielded the highest number of correctly identified events in all categories, but the differences to fuzzy c-means or histogram-based segmentation were less than 3 %. Differences were revealed when the number of correctly identified events was analysed for each crystal separately, as shown in figure 8.18.

All crystals, which were represented as a single, not-overlapping spot in the flood histogram, are marked in black. In this case, the ratio of correctly identified events was comparable for the three clustering methods. Of special interest are those crystal signals, which had to be separated from another crystal's signal in the flood histogram: these crystals are marked in blue, green and red, depending on the related crystal layers.

Figure 8.19 summarises the mean amount of correctly identified events for the three clustering methods, depending on the related crystal layers of the overlapping spots. The result for the bottom and intermediate layer (red) was the most diverging one. The best result could be achieved with model-based clustering (51.2 %). With fuzzy c-means, 49.8 % of the events from the bottom and intermediate layers could be identified correctly. During the histogram-based segmentation, the same light distribution and cluster centres had to be assumed for some of the overlapping spots of the bottom and intermediate layers. Since either of both crystals was selected randomly during the event positioning, the resulting ratio was around 36.0 %. For crystals of the top and intermediate layers (green) and top and bottom layers (blue), the results of the histogram- and model-based method were comparable, while fuzzy c-means yielded less correctly identified events. In case of the single crystal signals, all three methods performed equally (63.0 %).

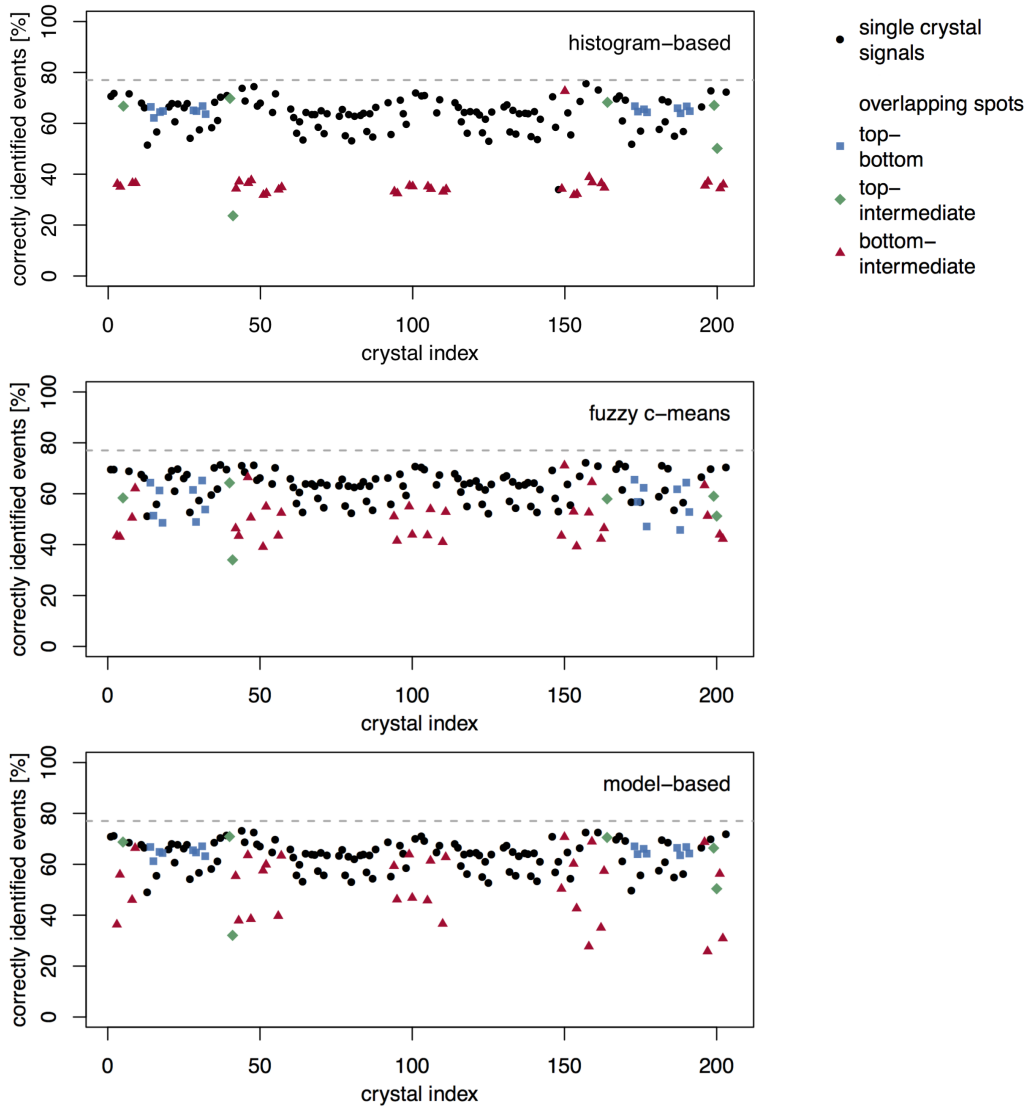


FIGURE 8.18: Ratio of correctly identified events for the individual crystals in the flood histogram. The crystals which gave rise to overlapping spots are marked in blue, green and red, and the theoretical limit of 77 % is marked with a grey line.




FIGURE 8.19: Mean amount of correctly identified crystals for the three types of overlapping spots and the single crystal spots.

Discussion

A positioning algorithm, that is robust to incomplete data, is of high importance for measurements. In the simulated data set, 22.5 % of the events have led to incomplete signals on the DPC tile. While these events can be discarded in a calibration scan, the case is different for a measurement with patients or volunteers, where the loss of sensitivity would be too large (the loss would be $1 - (1 - 0.225)^2 = 0.399 = 39.9\%$).

In total, model-based clustering yielded the highest ratio of correctly identified events. This result meets the expectations, because the model-based approach has been tailored to consider additional parameters such as the standard deviation, in order to improve the separation of overlapping clusters. In contrast, fuzzy c-means based the crystal identification on the euclidean distance of the single events, and an energy threshold was applied afterwards. This method divided the spots into two arbitrary halves, but still identified many scintillation events correctly.

The results of the histogram-based method were promising, but lacked precision at overlapping spots of the bottom and intermediate layer. The mentioned spots had been assigned the same cluster centres and light distributions, so that either of two crystals was randomly selected, which lead to low ratios of correctly identified events.

The model-based clustering approach was therefore the favoured method, based on the simulations. Certainly, the clustering method will have to be applied to measured data of the final detector geometry to make a decision. Lastly, it should be mentioned that the crystal layer has been identified correctly for 75.4 % (histogram-based), 75.9 % (fuzzy c-means) and 77.0 % (model-based) of the simulated events, when Maximum Likelihood positioning was applied to all events, including those with an incomplete signal distribution.

Other studies, which try to improve the crystal identification in pixelated PET detectors, mostly evaluate the performance by means of measurements. Since there is no ground truth available, a spatial resolution or hot rod phantom is measured and the same data set is analysed two times, using the standard and improved identification method. The influence can then be investigated by comparing line profiles inside the reconstructed images [76], [167]. This approach has the advantage, that all image degrading factors (noise, detector dead time, gain differences, positron range) are already included in the measurement, and that the overall influence on the reconstructed image can be highlighted.

On the other hand, the simulation method presented in this chapter allows one to study the influence of the clustering method in a very detailed way, e.g. for each crystal separately. It further allows one to optimise the signal separation at the edges, by investigating parameters like the standard deviation or the energy spectra separately for each crystal layer. But even more important are those events, which were not correctly identified: although 61.5 % of the events can be correctly identified, this corresponds to only 37.8 % correctly identified LORs. The ground truth information allows one to investigate the wrongly identified events in detail, and to develop new

strategies to improve the event positioning. Lastly, detailed information on the energy levels and standard deviation has been necessary to fully understand the measured flood maps, and to tailor the crystal identification method to a detector module.

The disadvantage of the simulation is the long computation time (use of super-computer required), which also limited the resources to match the simulation with measured data. However, the simulation time is related to the number of crystals, and is therefore significantly lower for double or single layer detectors. Since overlapping spots at the edges can occur in both single layer and multilayer detectors [108], [161], [168], the presented simulation study could be an interesting approach for evaluating crystal identification schemes for these detectors.

8.5 Measurements with a Detector Block of Final Size

So far, a small crystal stack with 8×8 crystals in the bottom layer has been used for all measurements. Having found an optimum crystal array size for the new BrainPET, first crystal arrays with the final dimensions (24×24 crystals in the bottom layer) have been provided for testing. In the final detector block, a customised dSiPM array will be used to match the large crystal array size: while the standard DPC tile has 4×4 dies, the customised tile will contain 6×6 dies. First analyses had to be conducted with the large crystal array and a standard DPC tile, because the customised tiles had not been available at that time.

As it has been mentioned in the previous section, the crystal surface finish was expected to have a strong influence on the scintillation light output and therefore the measured energy spectra. In the following, measurements with two different crystal arrays will be compared, one with a polished surface finish and one with a rough surface finish of the crystals' lateral faces.

8.5.1 Flood Maps with Ground and Polished Surface Finish

Methods

Flood histograms were recorded in the same measurement setup as before. The large crystal arrays were glued to a 1 mm thick lightguide (Herasil)³ and a standard DPC tile (4×4 dies). For event acquisition, the trigger was set to scheme 2 and the validation scheme was set to `0x55:0R`. To further minimise the recording of dark noise, the DPC tile was cooled to 11 °C. Flood histograms were calculated from the raw data after saturation and gain correction.

³Due to its availability in different layer thicknesses, Borofloat has been used during the initial measurements to find an optimum lightguide thickness. In these measurements, a 1 mm thick quartz glass lightguide (Herasil) was tested for comparison, due to its higher transmission in the UV range [122], [143].

Results

Figure 8.20 compares the resulting flood maps after 40 min irradiation with a ^{68}Ge line source. Similar to previous measurements, the flood histograms showed single regions, which corresponded to the 64 single dSiPM signals. In the centre, the single crystal signals formed a regular spot pattern in both measurements, while the spots partly overlapped at the outer edges. The blurred signals at the bottom and right edge were caused by the different dimensions of crystal arrays and DPC tile (the crystal arrays protrude over the DPC tiles' edges) and were not considered in the following analysis.

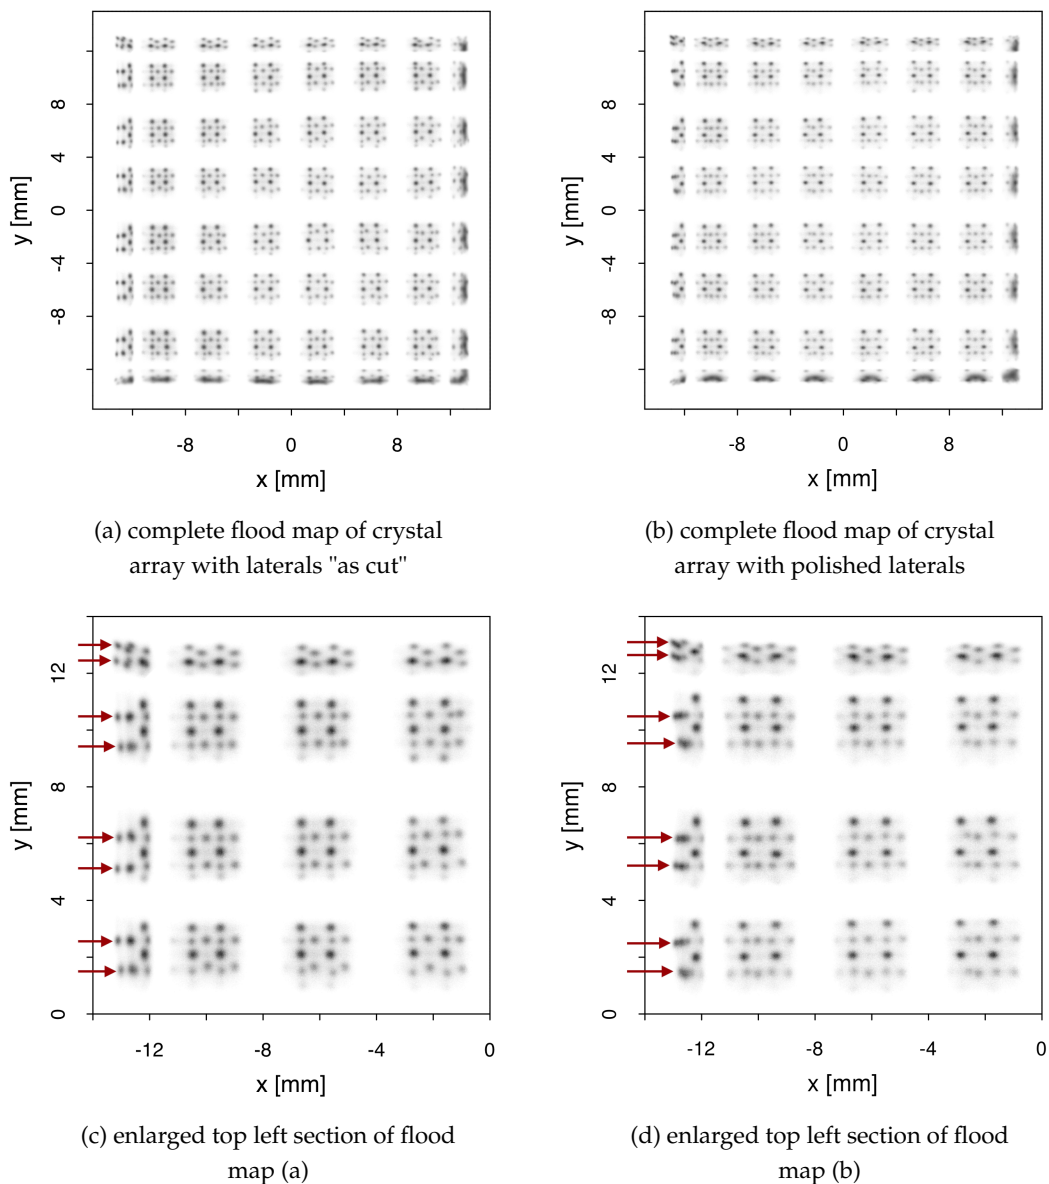


FIGURE 8.20: Recorded flood histograms of the crystal array with ground and polished surface finish are shown in figures (a) and (b). Figure (c) and (d) contain an enlarged view of the top left 4×4 dies. The spots at the edges were better separated in the ground finish crystal array (see red arrows).

The enlarged views in figures 8.20c and 8.20d show the top left section of the flood histograms, where the overlapping spots at the top and left edge were clearly visible. The crystal array with rough lateral faces yielded overlapping spots of the bottom and intermediate layer at the left edge, whereas the array with polished lateral faces additionally overlapped these spots with the outermost crystal row from the bottom layer (see red arrows in figures 8.20c and 8.20d).

At first sight, this observation spoke in favour of the ground surface finish, because the flood map indicated presumably easier separability of the crystal signals at the edges. On the other hand, polished crystals can yield a higher light output, depending on the crystal's aspect ratio [160]. Therefore, the impact of the crystals' surface finishes on energy resolution and separability of crystal signals was studied in detail in the following sections.

8.5.2 Energy Resolution of Both Crystal Arrays

Methods

To investigate the impact of crystal surface finish on a detector's energy resolution, the 6×6 centre regions of the flood histograms from figure 8.20 were analysed. Due to the cut-off light distribution, the dSiPM signals at the edges were excluded from the analysis, so that there were 144 spots to be analysed for each crystal layer. Since the spots in each flood map region formed a geometric pattern, the regions were separated into sections, which contained only one spot. The energy spectrum of each spot was calculated from the hot pixel's signal and its eight neighbouring dSiPMs⁴. Afterwards, it was approximated with either a single Gaussian distribution, or a sum of two Gaussians in case two energy peaks occurred. The energy resolution was then calculated as the ratio of the peak's FWHM and the peak position.

Results

Figure 8.21 compares the energy resolutions and the energy peak positions of the three crystal layers for both surface finishes. The energy resolutions of the three crystal layers lay between 12.19 % and 13.78 % for the rough crystal array, and between 12.51 % and 13.91 % for the polished crystal array.

⁴The energy resolution can also be analysed for the full light spot, which means that the signal of all triggered dSiPMs are used. Due to the trigger and validation thresholds, the number of fired dSiPMs can vary and lead to fluctuations in the measured light spot size. This has a direct impact on the energy resolution: depending on the thresholds, it can be higher or lower compared to the described energy resolution.

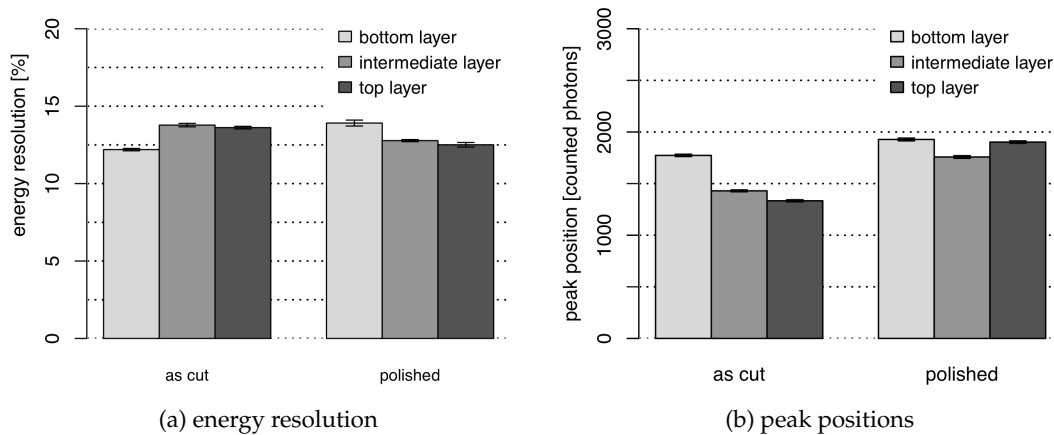


FIGURE 8.21: Energy resolutions and peak positions of the two crystal arrays with rough and polished lateral faces. Both parameters were calculated separately for the three crystal layers. The error bars show the standard error of the mean.

It was striking that the three crystal layers yielded different energy resolutions, especially in the polished array: the bottom layer performed worst, and the energy resolution decreased towards the top layer. This was not the case in the rough crystal array, where the bottom layer had the lowest energy resolution and where the intermediate layer yielded the highest value. Different energy resolutions had been expected for the three crystal layers due to the different crystal heights (different aspect ratio and therefore light output) and the scintillation light propagation inside the crystal array.

In total, the rough crystal array achieved a mean energy resolution of $(13.20 \pm 0.06)\%$, while the polished array achieved $(13.06 \pm 0.09)\%$. Under consideration of the standard errors, the two crystal arrays yielded the same overall energy resolution.

A difference between the two crystal arrays could be observed in the peak positions: in the rough crystal array, the highest number of photons was detected in events from the bottom layer and decreased towards the top layer. This effect had also been observed in the previously described simulations and measurements with the small crystal stack. In case of the polished crystal array, the top and bottom layers had very similar energy peak positions, whereas less photons were detected in events from the intermediate layer.

Discussion

The fact that such a high amount of photons was detected for events from the top layer could not be explained so far and contradicted the expectations from the simulations. At least, this observation showed that the scintillation light propagated in a very different way compared to the rough crystal array. The similar energy levels could become a disadvantage, when it comes to the separation of overlapping spots at the edges, because the energy levels of the three crystal layers have proven to be important for the separation. Judging by the absolute amount of detected photons,

the polished crystal layer had a higher light output. This observation meets the expectation from literature to find higher light output in the polished crystal array [160].

However, it remained difficult to compare the results to literature, because most fundamental studies were based on measurements and simulations of single crystals, and did not investigate the influence of the surface finish of crystals combined to an array. Especially with the 3M ESR vikuiti it was challenging to understand or even predict the light propagation in detail, because the reflector was partly transparent for certain incidence angles [150].

One study that had investigated a similar detector geometry with the 3M ESR vikuiti was related to the MINDView project: the detector consisted of three staggered crystal layers with polished LYSO crystals of $1.5 \times 1.5 \times 6 \text{ mm}^3$, which were shifted relatively to each other. Improved energy resolutions of 12.2 % (top), 9.6 % (intermediate) and 10.3 % (bottom) layers with a variation of 1 % could be achieved [110], and the energy peak positions increased from the top to the bottom layers. In this study, a different crystal aspect ratio, crystal material (i.e. light yield) and photodetector had been used, so that a comparison of further details was not possible.

8.5.3 Spot Separability in the Final Detector Block

Apart from the energy resolution, the separability of overlapping spots is an important feature that had to be investigated in detail. In the following, the separation of overlapping spots will be shown for three representative cases.

Top and Bottom Layers

The signals of the top and bottom layers were well separable in the rough crystal array, as visible in figure 8.22. Two distinct energy peaks were visible in the energy spectrum, and also the standard deviation clearly differed for both layers. The bottom row of figure 8.22 visualises the result of model-based clustering.

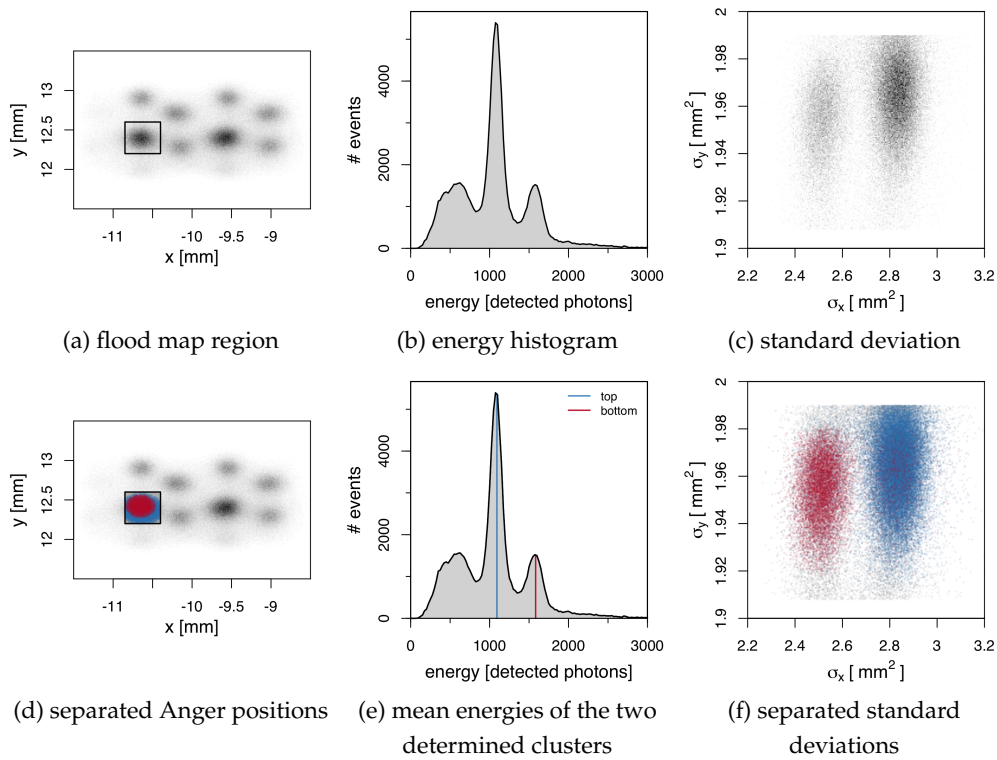


FIGURE 8.22: Rough crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the top and bottom crystal layers. Due to the distinct energy peaks and the clearly separated standard deviations, both spots could be easily identified, as shown in figures (d)-(f).

In contrast to this measurement, the separation was more complicated in the polished crystal array. As shown in figure 8.23, the energy peaks of both layers completely overlapped, so that the separation had to be based on the Anger positions and the standard deviations. This was still possible, because the spots did not completely overlap, but were slightly shifted in x direction. Moreover, the standard deviations were still separable, so that the results from model-based clustering (figures 8.23d-8.23f) appeared to be successful.

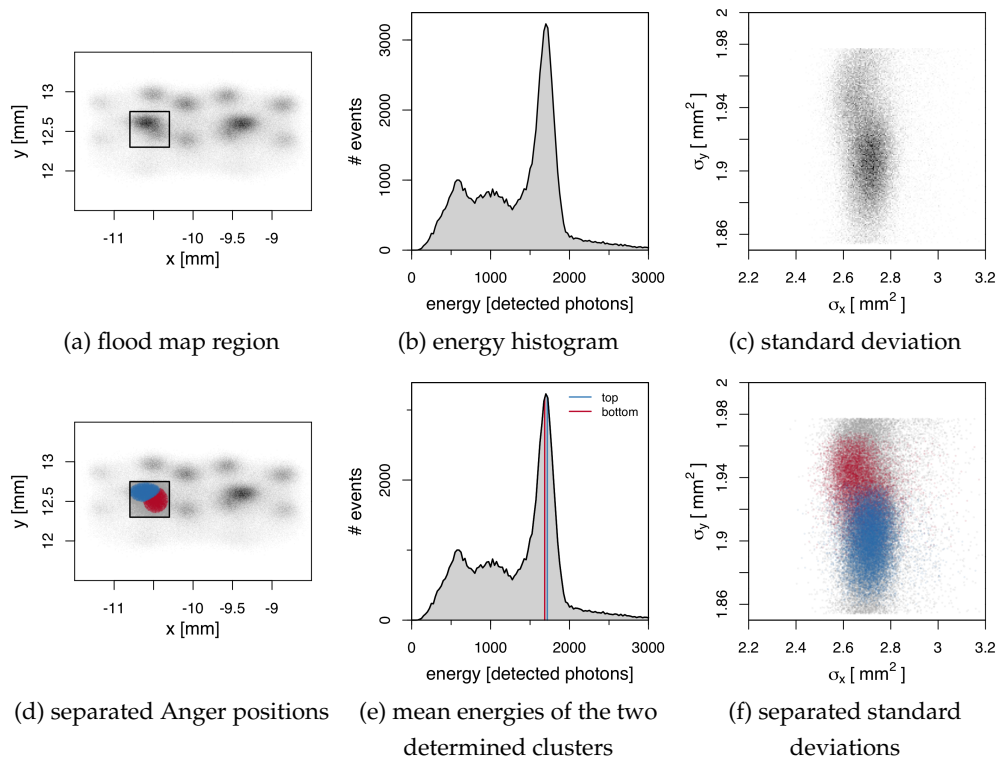


FIGURE 8.23: Polished crystal array: measured energy spectra and standard deviations for a point, which contained partly overlapping signals from the top and bottom crystal layers. Due to the congruent energy peaks, the separation was complicated compared to the rough crystal array, but appeared to be successful (figures (d)-(f)).

Bottom and Intermediate Layers

The separation of signals from the bottom and intermediate layers had been challenging in the previous measurement with the small crystal array. In the rough crystal array, the identification of both crystal signals was possible, as depicted in figure 8.24. Although the energy peaks were close to each other, and the standard deviations were not as clearly separated as in the previous case, model-based clustering could be applied to separate the crystal signals, as shown in the bottom row of figure 8.24.

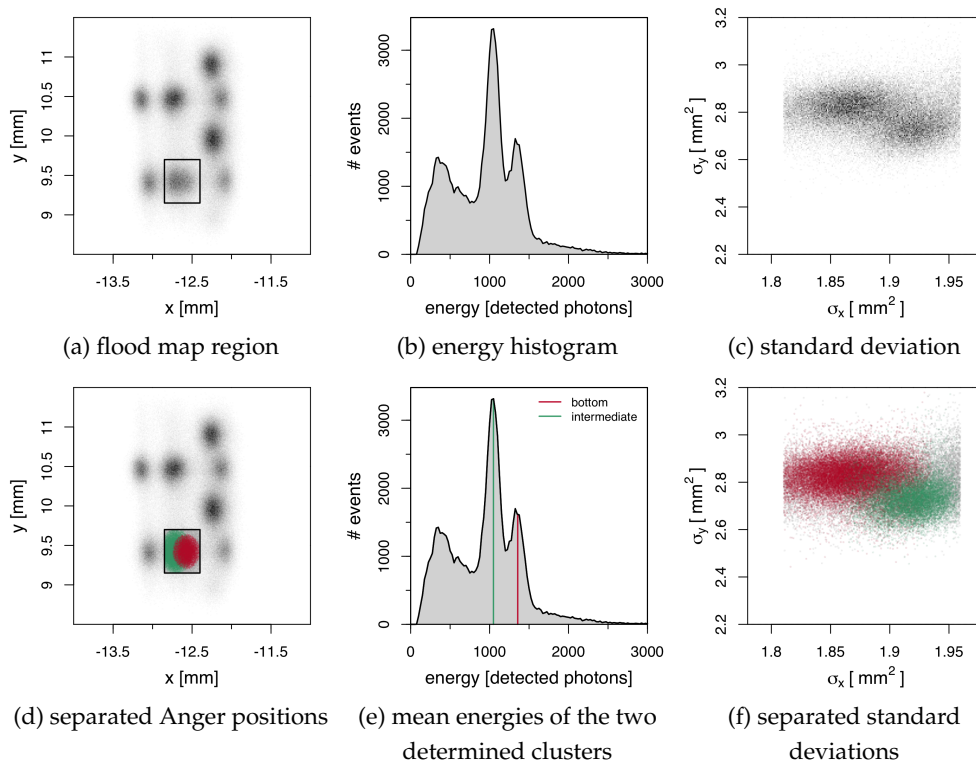


FIGURE 8.24: Rough crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the intermediate and bottom crystal layers. Although the energy peaks and standard deviations were closer to each other, compared to the previous case, model-based clustering could be used to identify both crystal signals. The results are shown in figures (d)-(f).

Figure 8.25 depicts the same situation for the polished crystal array. It has already been mentioned that three spots (instead of two in the rough array) overlapped at this edge. Since the energy levels of the three spots diffused in the same energy peak, the separation had to be focused on the Anger positions and standard deviations. For the case visualised in figure 8.25d and 8.25f, model-based clustering could only determine three sharply separated spots in a row, but no overlapping spots. These spots were regarded as not-separable.

Other spots from the same crystal rows seemed to be distinguishable by their different Anger positions, as indicated by the black arrows in figure 8.25a. But even in these cases, model-based clustering only divided the area into three rectangular

fields, instead of determining three spots of circular or oval shape, because there were no noticeable differences in energy or standard deviation (i.e. spot size). Compared to the rough crystal array, these crystal rows could not be considered to be separable.

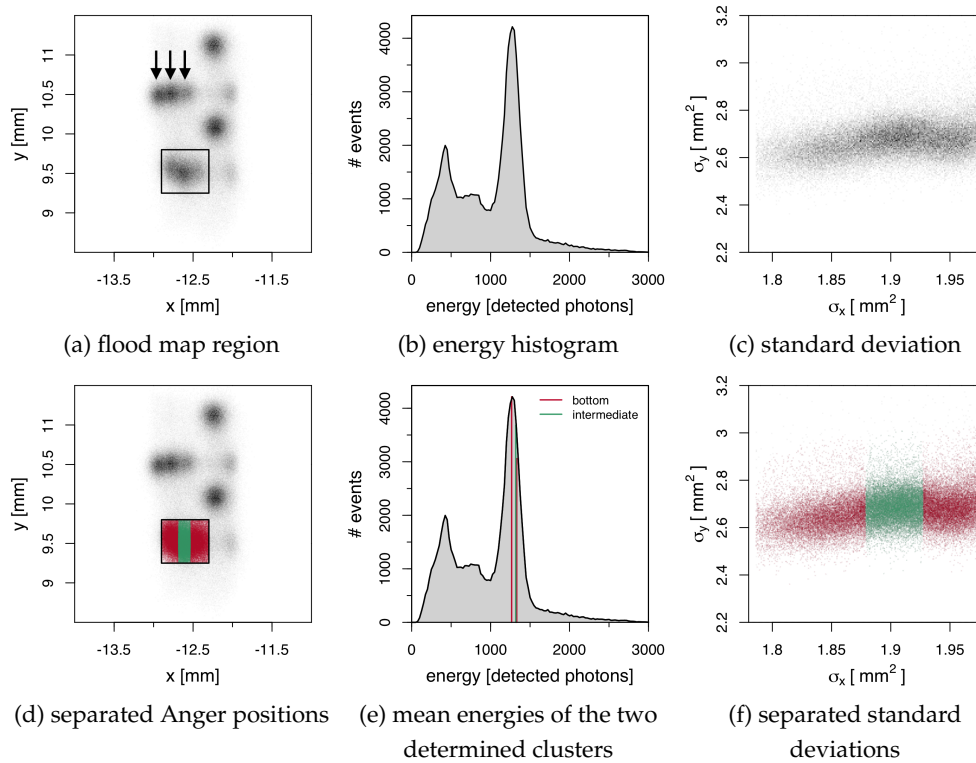


FIGURE 8.25: Polished crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the intermediate and bottom crystal layers. The single energy peaks and standard deviations overlapped, making it impossible to separate the spots with model-based clustering. The results are shown in figures (d)-(f).

Top and Intermediate Layers

The four outermost crystals at the corners of the top crystal layer gave rise to overlapping signals with the intermediate crystal layer. At least this held true for the rough crystal array: the resulting overlapping spots in the flood map were the most challenging ones, since the energy levels were very similar, and the standard deviations not separable. Still, model-based clustering could determine two clusters of different energy levels and sizes as can be seen in the bottom row of figure 8.26. It is noteworthy that only four spots of this kind occurred in the whole flood map, which consisted of 1634 crystal signals in total.

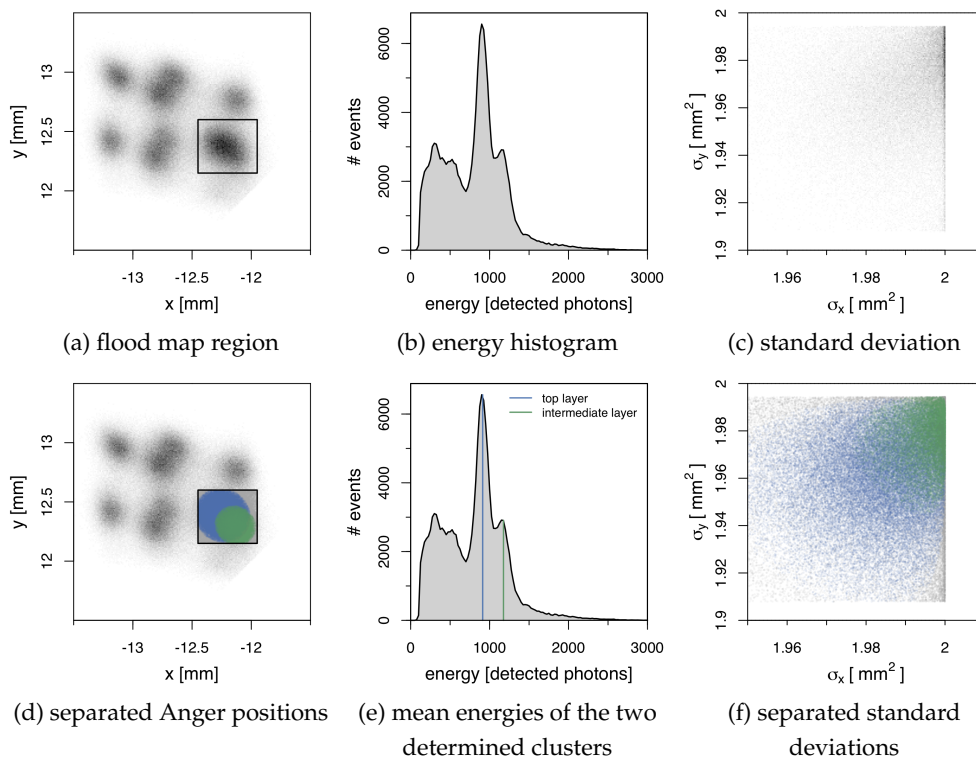


FIGURE 8.26: Rough crystal array: measured energy spectra and standard deviations for overlapping signals from the top and intermediate crystal layers. Despite the very close energy peaks, model-based clustering could identify both crystal signals (figures (d)-(f)).

In case of the polished crystal array, the spots of the top crystal layer were not overlapping and therefore all the crystals could be identified, as shown in figure 8.27.

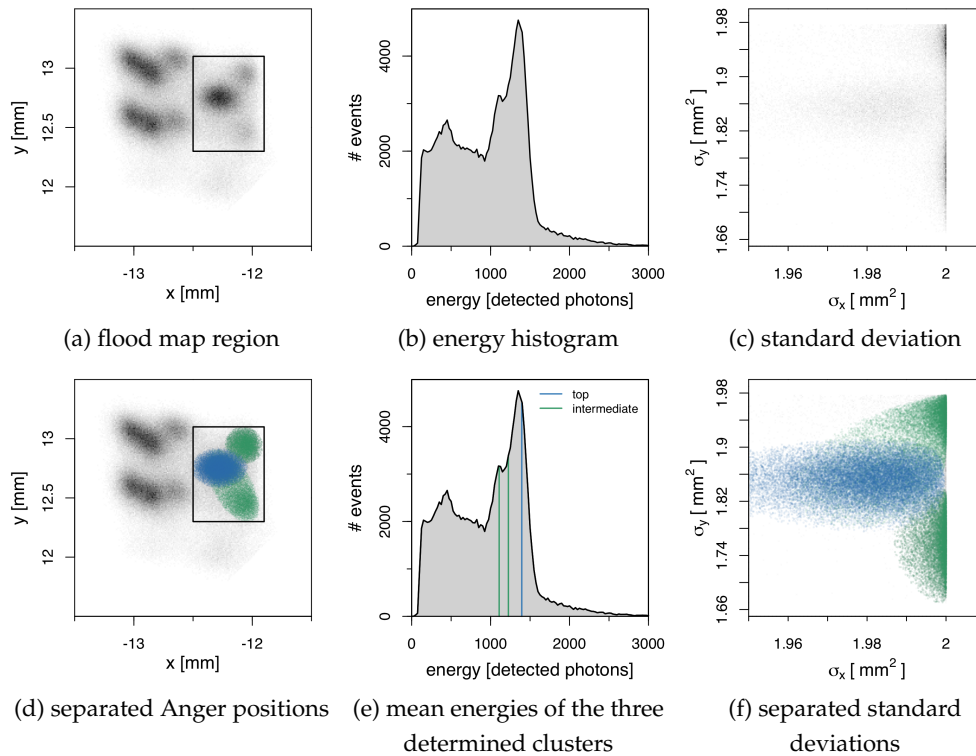


FIGURE 8.27: Polished crystal array: measured energy spectra and standard deviations for three signals from the top and intermediate crystal layers. There was no overlap of these points, so that model-based clustering could identify all three crystal signals (figures (d)-(f)).

Conclusion for Rough and Polished Crystal Arrays

Two aspects had been studied with the rough and the polished crystal array: energy resolution and separability of overlapping spots. The overall energy resolutions of the rough and the polished crystal array were very similar ($(13.20 \pm 0.06)\%$ and $(13.06 \pm 0.09)\%$), but differed for the single crystal layers. The polished array yielded the best energy resolution for the top layer, while it was the bottom layer in the rough crystal array.

The energy peak positions decreased from bottom to top layer in the rough array, while they were similar for the polished array. This observation was expected to influence the spot separability at the edges, because the different energy levels had been shown to be a valuable parameter for spot separation. This assumption proved to be true: in case of overlapping spots from the top and bottom layers and from the bottom and intermediate layers, only a single energy peak could be determined. The single energy peak made the separation more complicated– and even impossible for spots from the bottom and intermediate layer. Inside the whole crystal array, this would affect 5.8% of the crystals. In contrast, the rough crystal array yielded two energy peaks, which simplified the separation of the crystal signals and appeared to be more promising for automated separation. For this reason, crystal arrays with rough lateral faces will be used in the final detector.

Summary

In chapter 8, the simulated data were used to study the influence of reference responses for positioning algorithms.

The signals of all single crystals were identified in the flood map by means of three different segmentation and clustering algorithms. Since the ground truth of the simulation revealed specific energy levels and standard deviations (i.e. spot sizes) for signals from the three crystal layers, these parameters were considered during the separation of overlapping crystal signals.

Moreover, the influence of these three reference responses on the accuracy of event positioning was investigated. Two approaches were used for event positioning, nearest cluster search and maximum likelihood (ML) positioning. Maximum likelihood positioning showed superior performance in event positioning compared to nearest cluster search, especially for events with incomplete signal. Using model-based clustering for creation of the reference response lead to the highest amount of correctly identified events (61.5 % of 77 %). When the influence of the three reference responses was evaluated, similar ratios of correctly identified events (63 %) were determined for single (not overlapping) spots in the flood map. For overlapping spots the ratio differed significantly, but in all cases, model-based clustering yielded the highest amount of correctly identified events.

Two crystal stacks of the dimensions determined in chapter 4 were provided for further testing, which focused on the crystal surface finish.

In one crystal array, the lateral faces of the crystals were polished, while they were left "as cut" (rough) in the second crystal array. Since a higher light output was initially expected for the polished crystal array, the full energy peak positions and energy resolutions of both arrays were determined separately for each of the three crystal layers.

The rough and polished crystal arrays led to very similar energy resolutions ($(13.20 \pm 0.06)\%$ and $(13.06 \pm 0.09)\%$). As expected, the polished array had a higher light output, but the similar energy levels of the top and bottom crystal layer complicated the spot separation at the edges. Spots of the bottom and intermediate crystal layer could not be separated in the polished crystal array, which would affect 5.8 % of the crystals. Overlapping spots in the flood map acquired with the rough crystal array could be separated by considering the different energy levels and standard deviations in the clustering process. This makes it the preferred option for the new brain PET scanner.

Chapter 9

Conclusions and Outlook

Conclusions

In this thesis, a novel detector module for an MRI compatible human brain PET scanner was developed. Compared to currently available scanners, this device will comprise several new technologies, which are not yet state-of-the-art for PET/MRI scanners. Digital Silicon photomultipliers (dSiPMs) will be used for the detection of scintillation photons, crystals with a smaller pitch compared to whole body PET scanners, and a multilayer detector concept to achieve depth of interaction (DOI) information. Furthermore, all components will have to be insensitive to a magnetic field of the MRI (7 T), in order to be capable of hybrid imaging. At the beginning of this thesis, the general detector concept had already been determined: three staggered crystal layers should be used for DOI encoding. However, the actual dimensions of the detector modules were still to be optimised under consideration of the costs. Moreover, the combination of three crystal layers and digital SiPMs had not yet been investigated, so that the decoding of the flood map and optimisation of the lightguide thickness were still necessary. These aspects were investigated by means of GATE simulations and experimental studies.

The optimum detector module configuration for the new brain PET scanner was determined by means of GATE simulations. These simulations showed that a scanner option with three crystal layers and the largest detector modules (24×24 crystals in the bottom layer) would yield the highest sensitivity, since the solid angle coverage of the scanner would be maximised, and reduced gaps between the detector blocks would be achieved. Moreover, the three crystal layers yielded a higher depth of interaction resolution than a two crystal layer configuration, which made the spatial resolution more homogeneous within the whole field of view. This effect was especially visible at radial offsets, which cover the human cortex area in the acquired images [15]. Since these areas are of special interest for neuroscientific applications, the three crystal layer detector with the largest detector modules was the preferred design, despite higher costs compared to the other considered designs. Advanced simulations including the positron range, attenuation and photon acollinearity are currently work in progress, in order to estimate the spatial resolution more precisely.

While the sensitivity of different detector configurations was evaluated with standard GATE simulations, the simulations on spatial resolution required a sophisticated and customised hot rod phantom. The high counting statistics necessitated the use of the supercomputer JURECA [129]. A simulation of this extent exceeds commonly used GATE simulations in complexity, and provided valuable insights on the possible benefit from the DOI information, that would be available from detectors with two or three crystal layers. These insights were crucial for the decision making on the final detector block configuration.

Having determined an optimum detector module size, first measurements were conducted with a test detector. The concept of employing multiple crystal layers for DOI encoding has been proven feasible several times in literature [9], however, the combination of three crystal layers and digital SiPMs has not yet been investigated. The theoretically expected spot pattern was confirmed in the centre of flood map, whereas the identification of overlapping crystal signals at the flood map edges was not straightforward. Due to the propagation of scintillation photons in all three crystal layers, the overlaying pattern was not expected to be derivable from measurements with single and double crystal layer geometries. Furthermore, a direct comparison to measurements from literature was not possible, caused by the new combination of components. To gain deeper understanding of the flood map, the creation and propagation of scintillation photons was simulated with GATE.

Although the implementation of surfaces was challenging due to the "UNIFIED" surface model, simulated and measured flood maps were brought to good agreement. To obtain a ground truth information from the simulation, a GATE actor was implemented that stored information on the high-energy photon's interactions inside the detector module. This additional information allowed us to associate recorded scintillation events to the specific crystal where the scintillation had taken place– and therefore to decode the flood map. A second challenging aspect of this simulation was the long computation time, which required usage of the supercomputer JURECA again. With an average consumption of 10,000 core hours for a simulation of high statistics, the technical possibilities to match simulation and measurement in more detail were limited. GATE simulations of single PET detector modules have been described in literature before, but these focused on monolithic geometries [93]. The three crystal layers (and therefore the high amount of scintillator pixels and their surfaces) made the simulated geometry complex, which increased the computation time accordingly. Nevertheless, the simulations have been carried out on the supercomputer JURECA, since the simultaneously recorded ground truth information was the key to a detailed understanding of the flood map pattern.

To further refine the simulation, a new GATE digitiser was implemented that modelled the electronic readout of the DPC tiles. Being implemented on single photon level, effects like pile-up, dead time and the trigger and validation schemes of the

DPC tile were considered. This digitiser reproduced the measured energy spectra and spot sizes in the flood map more accurately compared to the standard digitiser. Moreover, the simulation became more realistic, since events with missing die signals (due to dead time or no trigger on the die) could be acquired, as observed in the measurements.

To the author's best knowledge, the DPC readout has not yet been implemented for GATE. In contrast to similar work for analogue SiPMs (e.g. G4SiPM [153] or GosSiP [152]), this digitiser is fully integrated into the GATE simulation. It is applied online and does not require any post-processing to gain the dSiPM output signal. So far, the novel GATE digitiser for the DPC readout has been used in simulations with pixelated crystal arrays. Future work should certainly test (and if applicable adapt) the digitiser for use with monolithic crystals, without causing memory overloads. In addition to the comparison with Tabacchini's analytical model [69], further comparisons of simulated and measured data are envisaged to validate the digitiser. For this purpose, less complex crystal geometries (single crystal arrays instead of the three staggered layers) should be preferred. To make the simulation more realistic, an extension of a sophisticated surface model for bonded crystal arrays and dedicated models for cross talk and afterpulsing would be required. So far, only a simple cross talk model has been implemented, that was used during the validation phase with Tabacchini's analytical model.

The developed GATE simulation, that allowed us to study the light transport inside a detector module, was employed to find an optimum lightguide thickness for the new brain PET modules. Being a connecting layer between the crystal stack and the DPC tile, the lightguide thickness has a crucial influence on the scintillation light distribution on the sensor [123], so that a variation of its thickness can be used to optimise the light spread on the sensor. As a figure of merit, the separability of crystal signals in the flood map was investigated as a function of different lightguide thicknesses. The best separability was achieved with 0.7 and 1.1 mm thick lightguides, which was confirmed by corresponding measurements with a test detector. It was decided to use a lightguide of 1.1 mm thickness in the final detector modules, since slightly better results were achieved for the top and intermediate crystal layers, and since it was easier to handle during the glueing process. An interesting finding was the fact, that a flood map can also be acquired without using any lightguide. The fact that all three crystal layer signals could be recorded without any lightguide, was in line with observations from literature, where the 3M ESR vikuiti between the bonded crystals was reported to be partly transparent [149], [150]. However, leaving out the lightguide would not have any benefit for the PET detector modules.

Having found an optimum detector module size and lightguide thickness, the last aspect addressed in the design phase was the accuracy of event positioning. To fully exploit the potential of the novel detector concept, the crystal where a

scintillation event took place has to be correctly identified during a measurement. Now that realistic data sets could be acquired from the simulation along with a ground truth information, different segmentation approaches were tested for the calibration step, and two event positioning schemes were compared. The number of correctly identified scintillation events was determined by means of the ground truth information: the crystal ID of the gamma photon's first interaction was compared to the one obtained from the event positioning methods. In that way, the number of correctly identified events could be determined for each combination of segmentation and event positioning approaches.

Using model-based clustering for creation of the reference response in combination with maximum likelihood positioning led to the highest amount of correctly identified events (61.5%). This supposedly low number becomes comprehensible, when it is considered that only 77% of all events can correctly identified in the ideal case, as revealed by the ground truth information. A correct identification is only possible, when the crystal of the gamma photon's first interaction coincides with the one containing the energy-weighted mean of all interactions.

Even though 61.5% of 77% of all events were already correctly identified in a single detector block, this would only amount to 37.8% correctly identified lines of response (LORs). This finding emphasises the need for improvement of event positioning approaches. The simulated data gives us the unique opportunity to select those events, where the scintillating crystal (of the gamma photon's first impact) could not be correctly identified. These events should be studied in detail, in order to develop new concepts for event positioning, and therefore to increase the number of correctly identified LORs.

Furthermore, the simulation studies of chapter 8 highlighted the importance of correctly identifying events with incomplete signal. If these events were discarded, the sensitivity would have been reduced by 39.9% in this simulation. Since deeper insights of the measured spot pattern with overlapping spots of different crystal layers could be gained, the separation of overlapping spots at the edges of the flood map was optimised. The presented simulation approach differs from those described in literature, where the impact of new crystal identification schemes usually has to be investigated in reconstructed images, since there is no ground truth information available [76], [167]. Especially for less complex detector geometries, e.g. with one or two crystal layers, the described simulations could be a promising approach to evaluate and optimise event positioning schemes for the specifically used geometry.

The optimum detector module configuration and a suitable event positioning approach have been determined for the new brain PET scanner. Two crystal stacks of the final dimensions were provided for further testing, which should decide on the crystal surface finish. In one crystal array, the lateral faces of the crystals were polished, while they were left "as cut" (rough) in the second crystal array. Due to the limitations of the UNIFIED model, the influence of the surface finish was not investigated with

simulations, but with measurements instead. The rough and polished crystal arrays led to very similar energy resolutions ($(13.20 \pm 0.06)\%$ and $(13.06 \pm 0.09)\%$). Despite higher output of scintillation light, the polished crystal array was not the preferred option. Different energy peak positions were observed for the three crystal layers in the rough crystal array. This resulted in excellent spot separability of the flood map, which means that all crystal signals could be identified. In contrast, 5.8% of the crystals could not be identified when using the polished crystal array. The latter was the decisive aspect, that made the rough crystal array the preferred option for the new brain PET scanner.

Outlook

In this work, it was shown that the combination of a crystal array with 24×24 crystals in the bottom layer, a $48 \times 48 \times 1$ mm lightguide and a customised DPC tile with 6×6 dies would yield the best performance of all considered detector configurations. So far, the experimental studies were conducted with components, that did not have the same dimensions or specifications as the final detectors will have. As soon as the customised DPC tiles will be available, all envisaged components will be evaluated in a first prototype of a detector module.

Future research should consider the influence of bias voltage and different trigger and validation schemes on the flood maps and energy spectra. First tests with different trigger and validation logic had already been acquired with the standard DPC tiles (not shown in this thesis for the sake of brevity), but extended tests had been postponed to the time, when the final DPC tiles would be available. A second important step will be the operation of a prototype module in the magnetic field of the MRI device.

In contrast to measurements with single detector modules, two modules will be used for measurements in coincidence. This will allow us to characterise energy and timing resolution of the novel detectors. A large part of future work will consist of fine-tuning of the data acquisition and processing unit.


The assembly of the dedicated brain PET scanner is scheduled for 2020. Having finished the design of all single components, the production and assembly of those is now in progress. At recent conferences, the high amount of contributions on brain PET detectors confirmed the demand for such a dedicated imaging device in neuroscience. The new brain PET scanner is therefore envisaged to make an important contribution to the field, by providing a commercial scanner combined with ultra-high field MRI for neuroscientific applications.

Appendix A


Additional Material

A.1 Schematic Drawing of Scanner Cassettes


BrainPET




Option I




Option II



Option III



Option IV



A.2 DPC Data Processing

An exemplary dark count map, recorded at 14.5 °C sensor temperature is shown in figure A.1. By disabling 10% of the micro-cells with the highest count rates, the dark count rate was reduced from 13.22 MHz to 3.46 MHz, as can be seen in the histograms below.

Dark Count Map

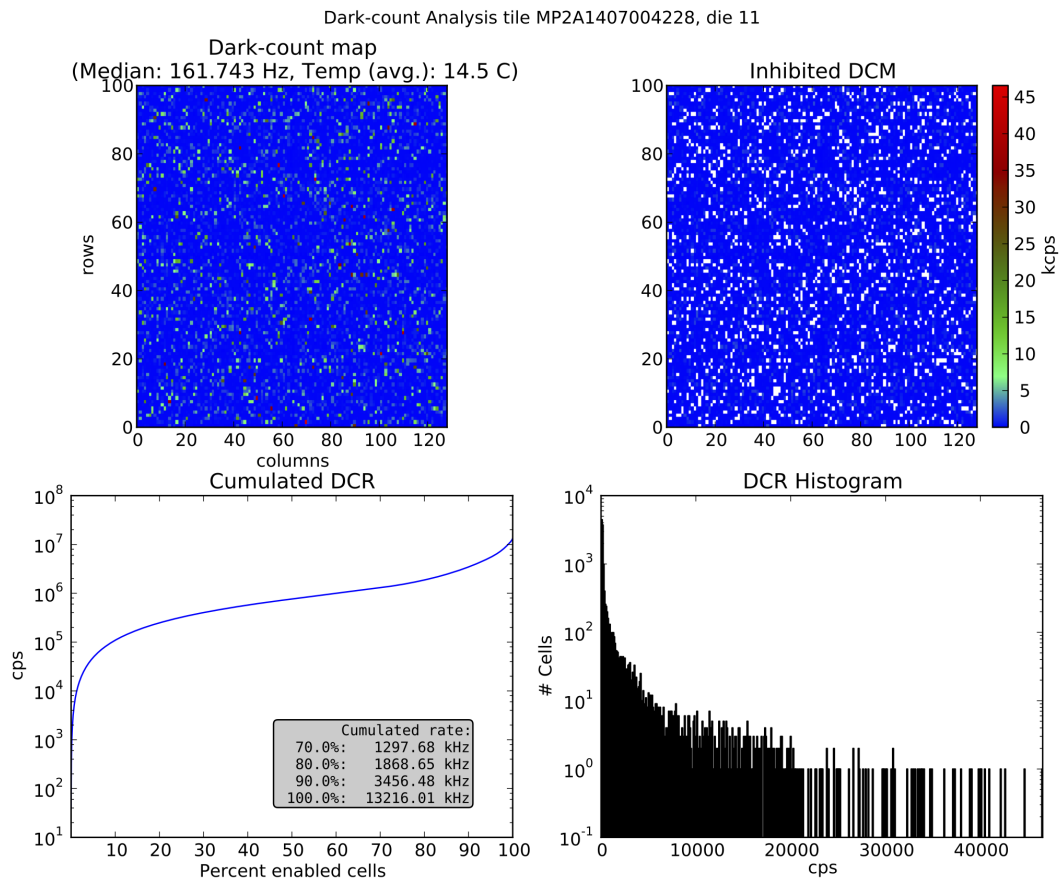


FIGURE A.1: The dark count map and inhibit map, which were recorded at 14.5 °C sensor temperature, are shown together with the corresponding histogram and cumulated dark count rate. The figure was created with the Philips DPC Tile-TEK software [66].

A.3 Validation of the GATE Digitiser

Trigger Probabilities

The probability that the trigger is set after a certain number of discharged cells is shown in figure A.2 for the three trigger schemes. Both *Tabacchini et al.* [69] and the GATE simulation assume 16 % crosstalk.

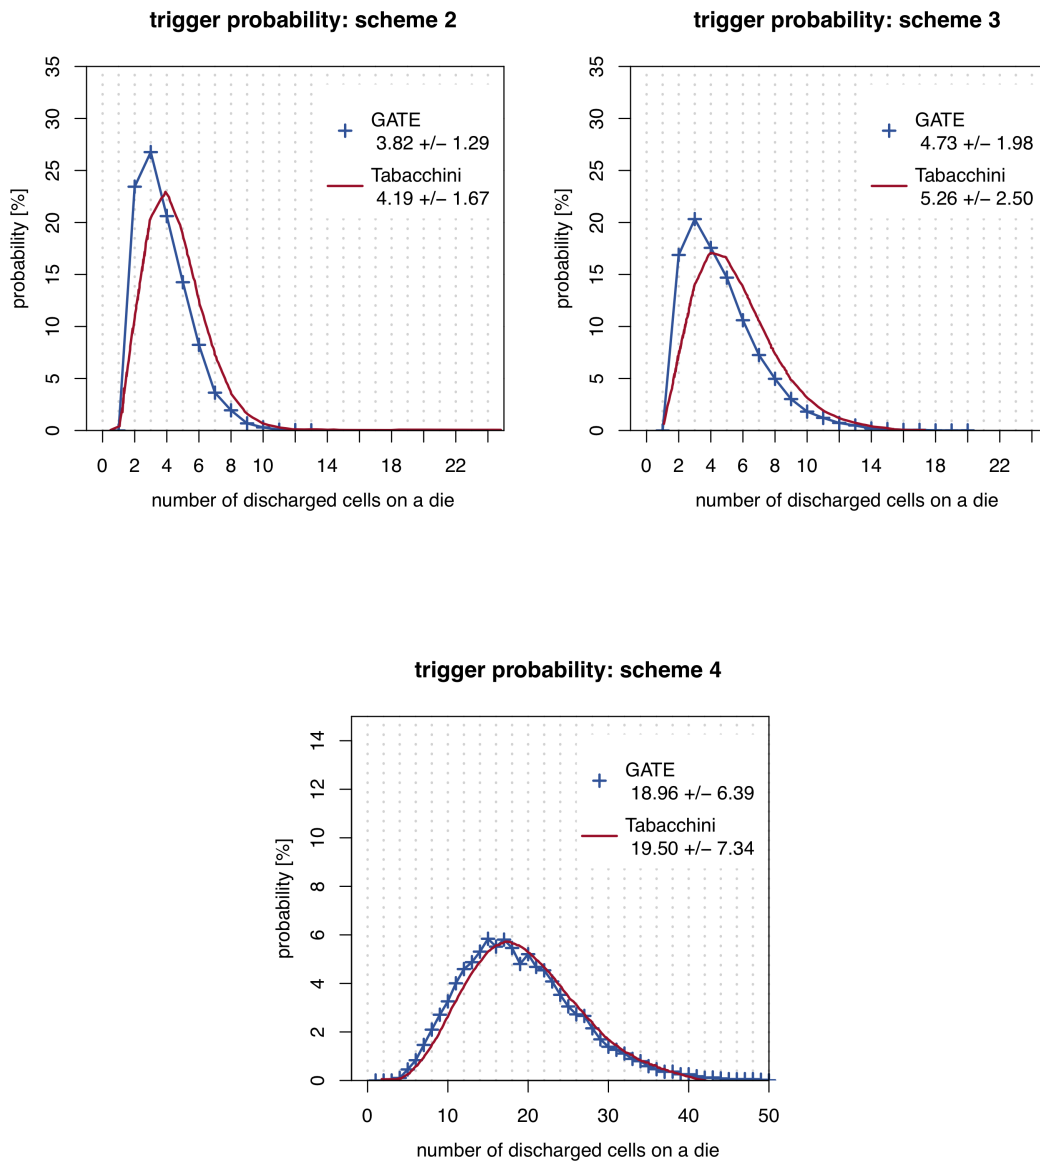


FIGURE A.2: Probability that the trigger is set after a certain number of cell discharges, displayed for trigger scheme 2 (top left), scheme 3 (top right) and scheme 4 (bottom row).

Discharged Cells When Trigger Condition Reached

The number of discharged cells at the exact moment of trigger generation is shown in figure A.3.

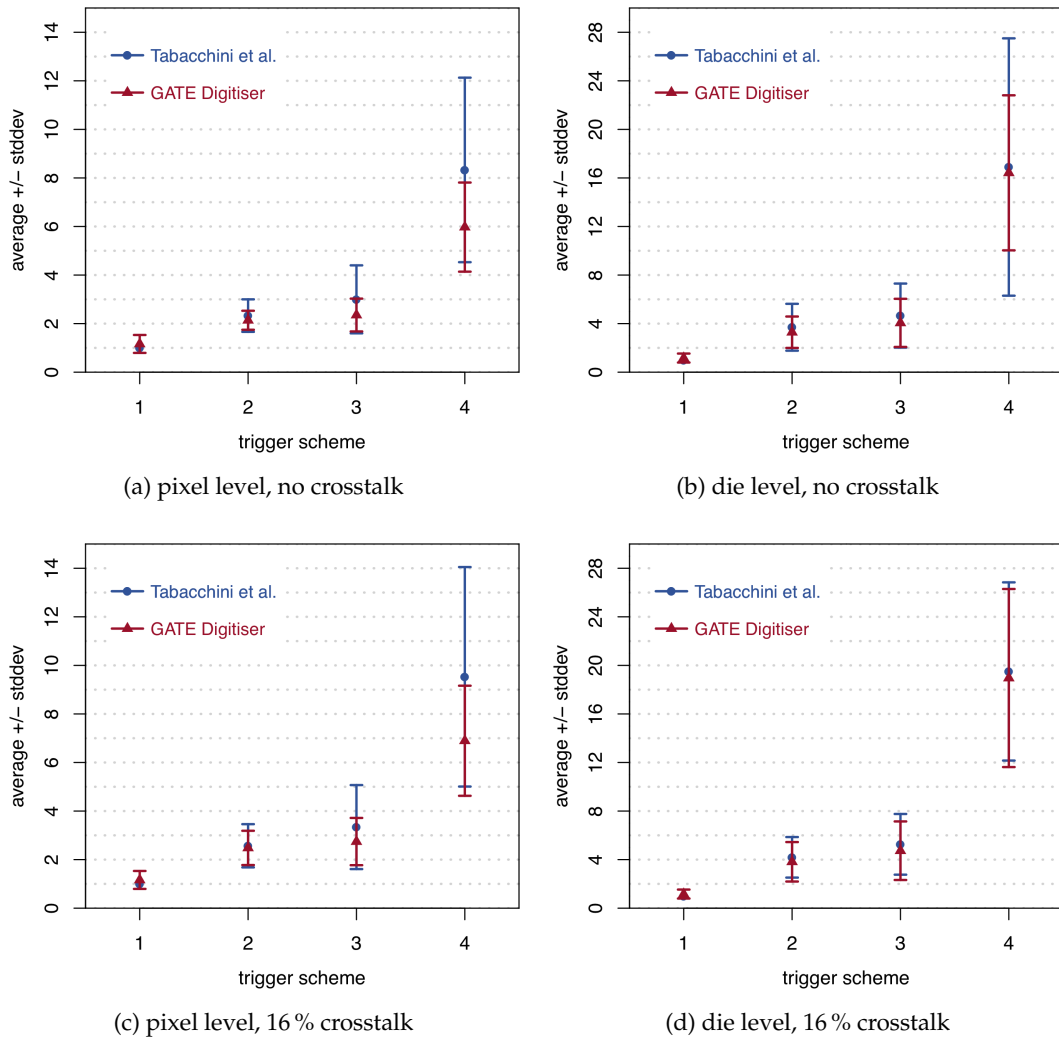


FIGURE A.3: Number of discharged cells at the exact time of trigger generation. The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d).

Figure A.4 shows the number of discharged cells at the exact moment of trigger generation, but this time the artificial scintillator with shorter rise and decay times was used.

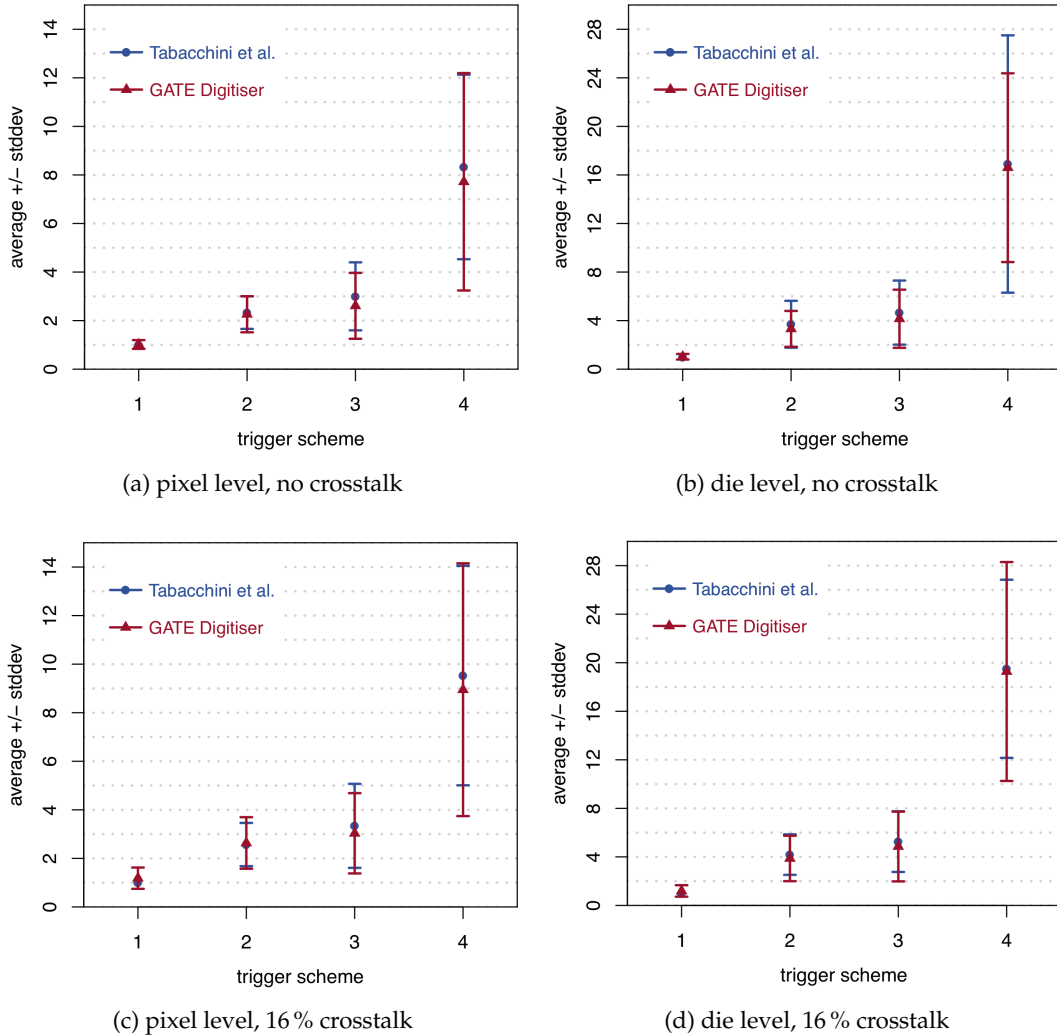


FIGURE A.4: Number of discharged cells at the exact time of trigger generation, using an artificial scintillator in the simulation (shorter rise and decay times). The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d).

Discharged Cells When Validation Condition Reached

Figure A.5 shows the number of discharged cells when the validation condition was fulfilled.

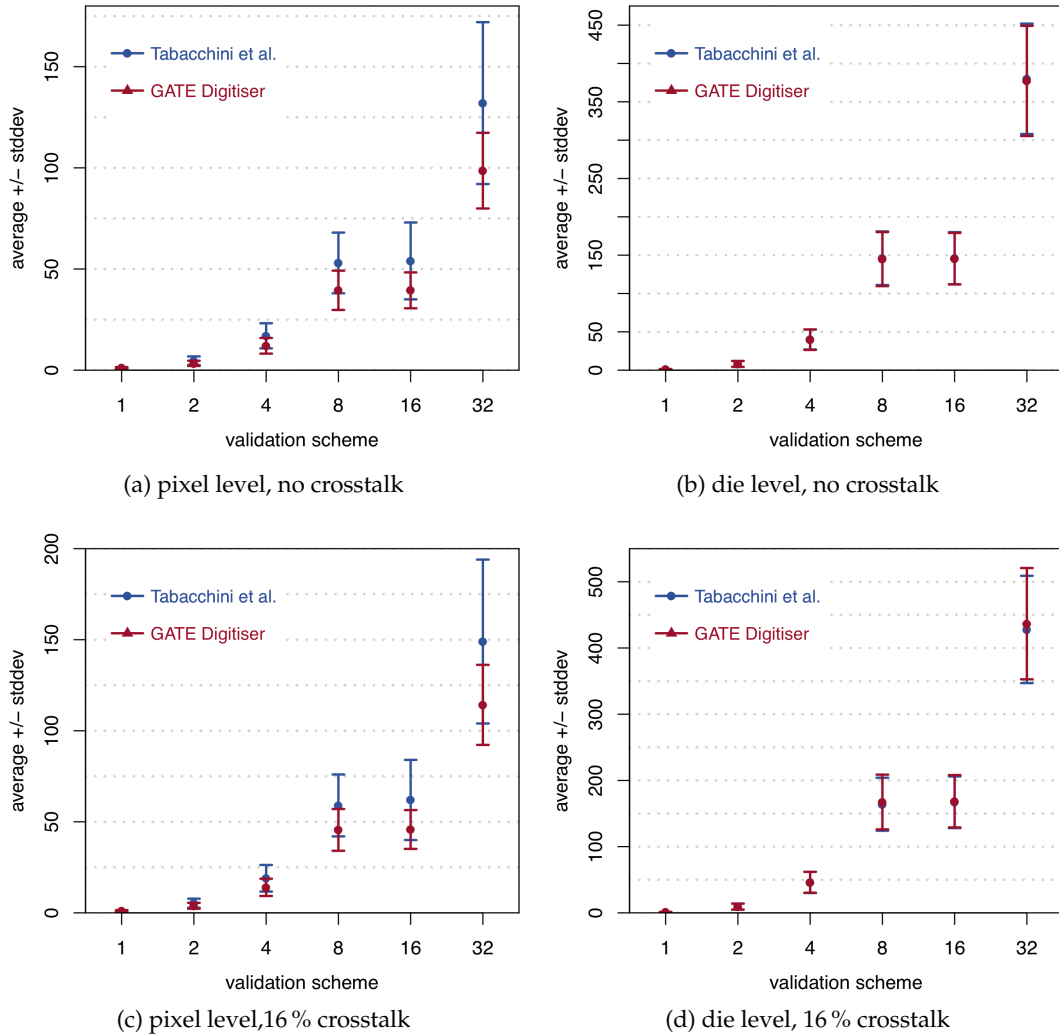


FIGURE A.5: Number of discharged cells at the time of fulfillment of the validation condition. The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d).

The number of discharged cells when the validation condition was fulfilled is shown in figure A.6, but this time the artificial scintillator with shorter rise and decay times was used.

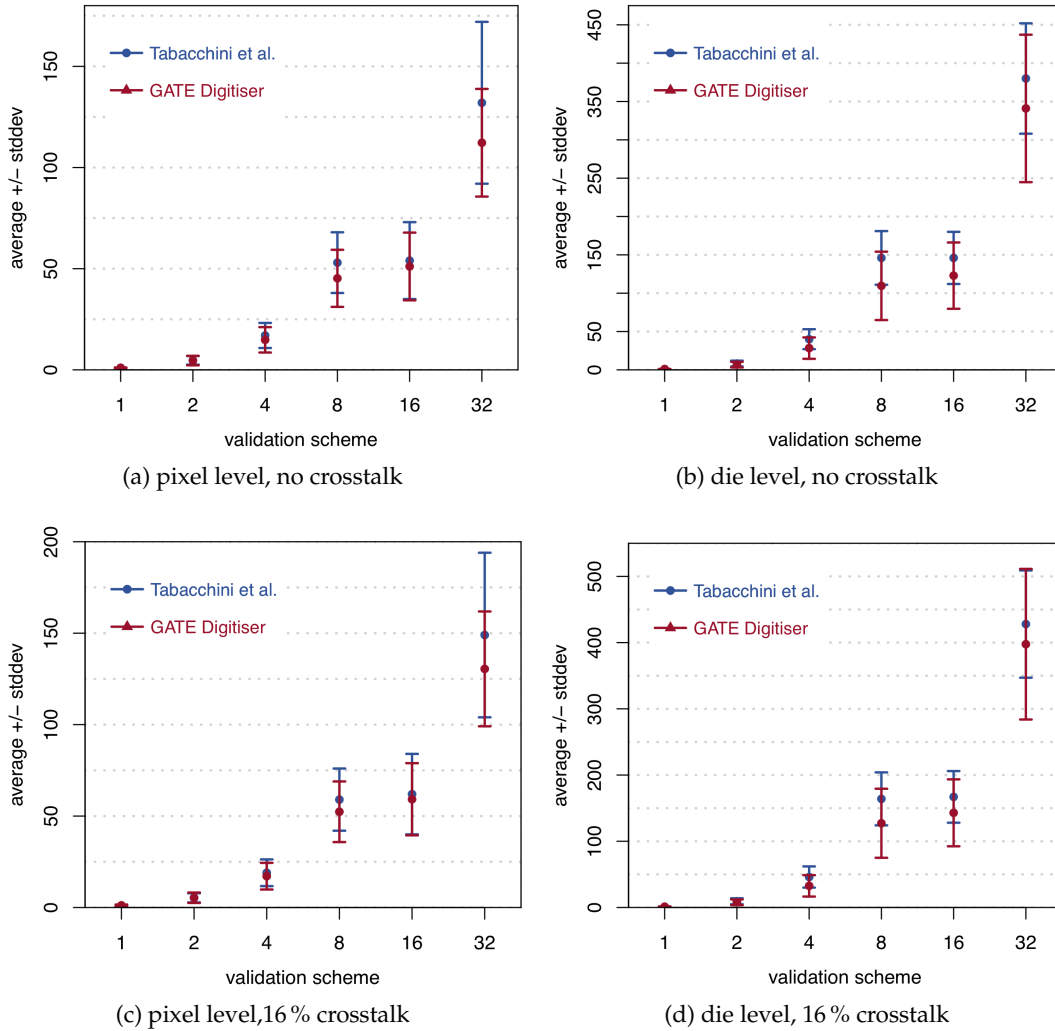


FIGURE A.6: Number of discharged cells at the time of fulfillment of the validation condition, using an artificial scintillator in the simulation (shorter rise and decay times). The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16% crosstalk (c,d).

Cumulative Validation Probabilities

The probability for validating an event after a certain number of cell discharges is shown in figures A.7 to A.9 for the different validation patterns.




FIGURE A.7: Cumulative probability of fulfilling the validation schemes 2-OR, 4-OR, and 8-OR after a certain number of cell discharges. 16 % crosstalk was assumed.




FIGURE A.8: Cumulative probability of fulfilling the validation schemes 4-AND and 8-AND after a certain number of cell discharges. 16 % crosstalk was assumed.




FIGURE A.9: Cumulative probability of fulfilling validation scheme 8-AND after a certain number of cell discharges. The results are shown with and without the influence of crosstalk.

A.4 Simulations Concerning the Reflectivity of 3M ESR Vikuiti

The measurements of chapter 7 have shown that a flood map can also be recorded without using any lightguide between the crystal stack and the DPC tile. It was assumed that the partial transparency of the 3M ESR vikuiti allowed enough horizontal light spread in the bottom crystal layer to yield the clearly resolvable signals in the centre of the flood map.

Figure A.10 shows the resulting flood maps of simulations, that did not include any lightguide and assumed different reflectivity values for the 3M ESR vikuiti. The cover glass of the DPC tile was included, though.

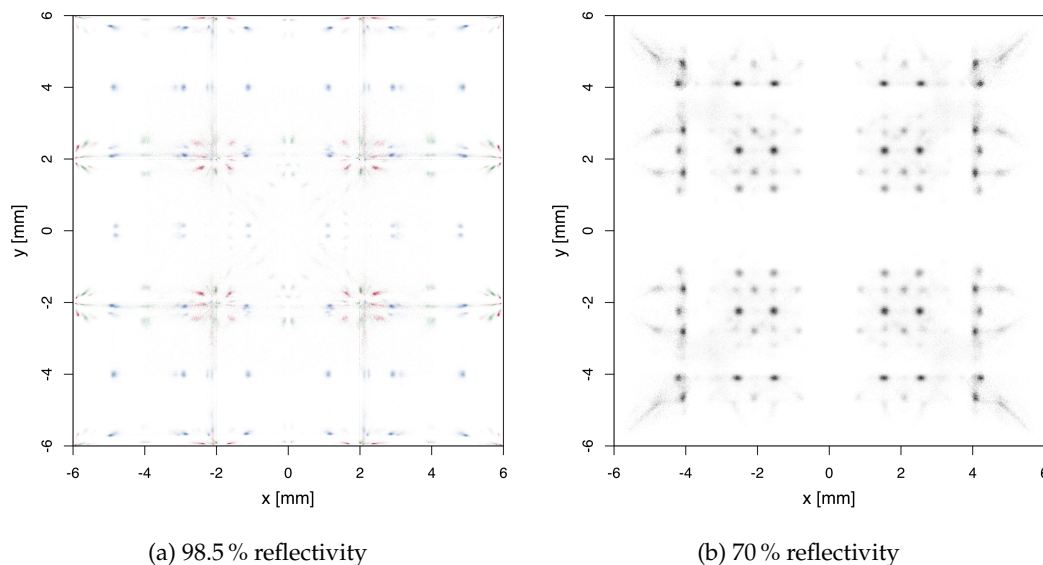


FIGURE A.10: Simulated flood maps without any lightguide, using different reflectivity values for the 3M ESR vikuiti. In (a), the signals of the three crystal layers are drawn in different colours: red– bottom layer, green– intermediate layer, blue– top layer.

So far, the Anger positions were always calculated from the hot pixel signal and all neighbouring pixels, in case all of them had detected scintillation photons. Due to the reduced scintillation light spread in the simulation from figure A.10a, the Anger positions had to be calculated from all available scintillation events, independent from the number of available DPC pixel signals. It is clearly visible that only the signals from the top crystal layer were resolved. Single signals of the intermediate crystal layer were visible, but only very few spots belonging to the bottom crystal layer were resolved. In contrast, the simulation with 70 % reflectivity of the vikuiti (figure A.10b) showed the expected spot pattern in the centre of the flood map. This result confirmed, that the cover glass of the DPC tile did not have any significant impact on the scintillation light spread, and that the vikuiti reflectivity was the parameter of highest impact in these simulations.

List of Figures

2.1	The coincident events comprise true (a), scattered (b), random (c) and multiple (d) coincidences.	8
2.2	The parallax error in one transaxial direction in different crystal blocks. While it is quasi not existing in the centre field of view (FOV) (limited to crystal size), the parallax error becomes larger towards the periphery of the FOV. The LORs in the areas marked in grey are assigned a fixed LOR (black lines), which can differ a lot from the true LORs (red lines).	10
2.3	The scintillation process described in the band model. a) Possible excitations of an electron, leaving a hole in the valence band. b) Electron caught in a trap, the remaining energy is transferred to the crystal lattice (phonon). c) De-excitation of an electron under emission of light γ_g with E_g - self absorption is probable. d) Recombination via activator states, emission of scintillation light γ_{scint}	12
2.4	Two possible detector layouts: either a monolithic crystal is directly glued to a photodetector array, or a crystal array is used in combination with a lightguide, which distributes the scintillation light over multiple sensor pixels.	14
2.5	Energy spectrum of a single LSO crystal, which was coupled to a light diffuser and a photomultiplier. The spectrum was recorded using a ^{22}Na source. Data redrawn from reference [37].	15
2.6	Different DOI capable detector geometries.	18
2.7	Flood map of a crystal array with 30×30 crystals. An accumulation point can be identified for each crystal. Data redrawn from reference [58].	20
2.8	The sensor architecture of a Philips dSiPM shown in a photograph of a tile: it consists of 16 dies with 4 pixels and 16 subpixels each. A subpixel is further divided into eight row-trigger-lines.	24
2.9	The event acquisition pipeline of a Philips dSiPM: when the trigger has been detected, a validation scheme is used to distinguish true scintillation events from dark counts. After integration and readout, a global recharge sets the die back to ready state. Redrawn from reference [66].	25

2.10	Each subpixel is divided into eight RTLs. To validate an event, the geometric distribution of photons in the RTLs is compared to a selectable pattern, which can be modified by setting gates 0-6. Redrawn from reference [66].	26
3.1	Coronal (a), sagittal (b) and transverse (c) view of the reconstructed images of a simulated phantom (BrainWeb [99]): while the grey matter structures are well resolved in the central transverse field of view (green arrows), off-centre structures appear blurred due to the degradation caused by the parallax error (red arrows). The simulated and reconstructed data was provided by Hancong Xu.	36
3.2	Crystal geometry for the next-generation brain PET: three staggered layers are supposed to yield DOI information, being composed of $2 \times 2 \times 8-10 \text{ mm}^3$ LSO crystal pixels. Finding the optimum crystal array size is one of the objectives of this thesis.	40
3.3	A crystal block prototype, consisting of 24×24 crystals in the bottom layer. The whole block is covered with a specular reflector, and a lightguide is glued to the bottom crystal array.	40
3.4	Top row: The three crystal arrays are shifted relatively to each other. The intermediate layer is only shifted by half a crystal pitch in x direction, while the top layer is shifted in both directions (one crystal pitch x and half crystal pitch y , relatively to the bottom layer). Bottom row: A schematic drawing of the spot pattern, which is expected to be found in the flood map (only a small section is shown).	41
4.1	A line source is aligned to the z axis of the scanner, emitting pairs of annihilation photons. As shown for scanner option III, the detector block signals of each cassette are calculated independently.	45
4.2	Axial sensitivity profiles for the three new scanner design options and the Siemens BrainPET for comparison. The increase in sensitivity due to less dead space and a longer axial FOV is clearly visible. For each scanner design, the peak sensitivity is given in the legend.	47
4.3	Axial sensitivity profiles for the three new scanner design options, and the two layer option (IV) for comparison. The sensitivity profile is equivalent to the one of option II with intermediate sized blocks. For each scanner design, the peak sensitivity is given in the legend.	49
4.4	The modified hot rod phantom inside scanner option III.	51
4.5	Reconstructed image of the hot rod phantom inside the three scanners (150 iterations, isotropic voxels of 1.25 mm^3 , one subset). Each image shows the sum of all planes in axial direction. The coloured lines indicate the position of the line profiles in figure 4.6.	52

4.6	Line profiles through two compartments of the phantom, as indicated in figure 4.5, for the BrainPET (top), the two layer scanner design (middle) and the three layer option (bottom). The decrease of peak to valley ratio is strongly visible for the BrainPET, while the DOI consideration in option IV and III homogenises the peak to valley ratio over the FOV.	53
4.7	Peak to valley ratios of the line profiles for the 2.4 mm rods (left side in figure 4.6). The ratio has been calculated for all peaks and their neighbouring valleys, which are located between the two outermost rods. The legend contains the minimum and maximum peak to valley ratios for each scanner option.	54
4.8	Line profiles for the 1.6 mm and 4.8 mm rods in the scanner with three crystal layers. Except for the two outermost rods, the 1.6 mm rods could be identified.	54
5.1	Lightguide and scintillator arrays were aligned on the DPC tile by means of 3D printed stencils.	58
5.2	Projected alignment of the three crystal layers on the DPC tile: the white semi-transparent squares represent the single crystals within a crystal layer, while the DPC tile is shown in the background photo. The section shows 5×5 pixels (2.5×2.5 dies) of the DPC tile.	58
5.3	Measurement setup with a line source: crystal stack and DPC tile were located in a light-tight black housing. The line source was positioned at a fixed distance by means of a foam cube.	60
5.4	Approximate, relative gain estimation for the single dSiPM pixels: the count spectra for two pixel signals are shown in (a), while the corresponding gain histograms are shown in (b). A threshold determines the relative gains, and the inverse values of the relative gains are used as the relative gain correction factors.	61
5.5	Flood histogram of the test detector. The three crystal layers can be clearly resolved in the centre part, while the spot pattern at the outer edges partly overlays.	62
5.6	Detailed view of a single dSiPM signal from the centre of the flood map. The spots could be assigned to the three crystal layers and were well distinguishable.	63
5.7	When a crystal is located above a bond gap or above two dSiPM pixels, two event positions are possible, depending on the maximum light detection on either of both sides.	64
5.8	Exposed view of the flood histogram from figure 5.5: the missing spots are marked in red, while not identifiable signals are marked with blue boxes.	65

6.1	The simulated detector geometry and the corresponding surfaces. The parameters of the surfaces 1-9 are defined in table 6.1.	68
6.2	Visualisation of a scintillation event inside the test detector. The interactions of the high-energy photon are represented by a red line and highlighted with numbers, while the scintillation photons are visualised with green lines. The light yield was lowered to 500 photons per MeV, so that the optical photons' trajectories can still be distinguished.	71
6.3	Setup of the simulation using a line source: the source is placed 65 mm from the top crystal layer, and emits 511 keV photons in the specified solid angles.	72
6.4	Simulated and measured flood map for the detector prototype with 1 mm lightguide.	73
6.5	Simulated and measured 3M ESR vikuiti reflectivity for different incidence angles at a wavelength of 420 nm. Data adapted from references [139], [150], [151].	75
6.6	Simulated flood map with reduced 3M ESR vikuiti reflectivity.	76
6.7	Exposed view of a single dSiPM signal, (a),(c) simulated and (b),(d) measured. The same number of events is displayed. For comparison, the two energy peak positions in the energy spectra are marked with dotted lines.	77
6.8	Outline of the simulation: the digitiser module is called after each event.	79
6.9	Structure of the digitiser module. The pulse lists of multiple events from the same time slice are evaluated together. The underlying read-out scheme is shown in figure 2.9.	80
6.10	Number of discharged cells at the exact time when the trigger condition is reached. The results are shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16 % crosstalk (c,d).	83
6.11	Probability, that the trigger is set after a certain number of cell discharges, displayed for trigger pattern 2.	84
6.12	Cumulative probability, that validation scheme 8-AND is fulfilled after a certain number of cell discharges. The results are shown with and without the influence of crosstalk.	84
6.13	Emission profiles of a 50 ps laser pulse and an LSO scintillation event.	85
6.14	Number of discharged cells at the exact time when the trigger condition is reached, shown for pixel level (a,c) and die level (b,c), without crosstalk (a,b) and with 16 % crosstalk (c,d). In this simulation, the artificial scintillation source was used for photon emission.	86
6.15	Exposed view of a single dSiPM signal, (a),(c) simulated and (b),(d) measured. The same number of events is displayed. For comparison, the two energy peak positions in the energy spectra are marked with dotted lines.	88

- 6.16 Simulated energy spectra of the 16 dSiPM pixels. The lutetium background radiation was not considered in this simulation. 89
- 6.17 Measured energy spectra of the 16 dSiPM pixels. The lutetium-related peaks (blue arrows) are clearly visible at the outer edges, while they are not visible in the centre, where always 9 dSiPM signals are required for an event. The red arrows indicate events presumably originating from the 1077 keV γ emission line of ^{68}Ga . The top row and left row (white) were not considered in this analysis, since the crystal arrays of the detector prototype were smaller than the lightguide and DPC tile (alignment shown in figure 5.2c). 90
- 6.18 A single energy spectrum of a dSiPM at the right edge of the sensor tile is shown. The lutetium-related peak was dominant in events with up to 8 triggered dSiPMs (a), but decreased when more neighbouring dSiPMs (12 (b) or 16 (c)) were considered. The smaller amount of scintillation photons was not sufficient to trigger a signal of all 16 dSiPMs. 91
- 6.19 Simulated flood histogram: the event positions are encoded with the ground truth information (red: bottom layer, green: intermediate layer, blue: top layer). The 16 dSiPM signals have been normalised to the respective maximum values. 91
- 6.20 Separate flood histograms of the three crystal layers. Crystals above a bond gap gave rise to two spots with lower intensities in the flood histogram. Two exemplary rows of spots, which belong to only one crystal row each, are marked with grey boxes. 92
- 6.21 Exposed view of a single dSiPM signal at the right edge of the sensor, (a) measured and (b) simulated. 93
- 6.22 Exposed view of a single dSiPM signal at the lower edge of the sensor, (a) measured and (b) simulated. 93
- 6.23 Exposed view of a single dSiPM signal at the corner of the sensor, (a) measured and (b) simulated. 94
- 7.1 Flood histograms of measurements with 1.1 mm and 2.0 mm lightguide thickness. A comparable number of events is displayed in both figures. 96
- 7.2 Simulated flood histogram for the detector prototype with a 1.1 mm lightguide. The peak to valley ratio was evaluated for three spots on each of the four centre dSiPM pixels (marked with black boxes), using flood maps simulated with different lightguide thicknesses. 97
- 7.3 Setup for simulations with different lightguide thicknesses. The point source emits 511 keV photons in confined angles to efficiently irradiate the crystals of interest. 98

7.4	Analysis of the flood histogram: (a) The flood histogram was divided into horizontal slices, the slice of interest is indicated by a red arrow. (b) The 1D histogram of the selected horizontal slice is calculated, and the kernel density estimate is determined in (c) (red line). The peaks of interest and the valley in-between are marked with red arrows.	99
7.5	Mean PVRs (\pm standard errors) of the four crystals from each crystal layer.	99
7.6	Ratio of the number of detected events and the number of all first interactions of the high-energy photons, shown separately for each crystal layer. The ratio increases with lightguide thickness.	100
7.7	Mean PVRs (\pm standard errors) of the four crystals from each crystal layer, taken from measured data.	101
7.8	Flood histogram for a measurement without any lightguide. Despite the missing lightguide, all crystal signals can be clearly identified in the centre of the DPC tile.	103
7.9	Flood histograms of a single centre dSiPM for measurements with different lightguides. Not using any lightguide influenced the spot positions of the bottom crystal layer, as indicated by blue arrows in (a). With increasing lightguide thickness, the spot size became larger. Since the crystal array had to be removed and glued to different lightguides for these measurements, a small shift in y direction caused different spot intensities of the top layer in figure (b) and (c), marked with green arrows.	103
8.1	Steps of the calibration process: the flood histogram is divided into 4×4 regions, indicated by red lines in (a), to simplify the clustering of single crystal signals. The clustering result for the crystal signal marked in (a) is shown in figure (b). Once the crystal signals are separated, the reference response matrix can be calculated for each crystal. Figure (c) shows the reference response matrix for the crystal signal from figure (a).	108
8.2	Steps of the event positioning: the likelihood is calculated from the acquired light distribution of a scintillation event (a) and all reference response matrices. The highest likelihood value is achieved for the best matching light distribution from the reference response matrix, shown in (b), and the corresponding spot index in the flood histogram (c). A lookup table assigns crystal layer and crystal location to the selected spot.	109

-
- 8.3 Energy spectra of single spots from the three crystal layers: when the DOI information was used to filter events, the energy spectra showed a distinct energy level for each crystal layer. Events from the bottom layer (red) had the highest energy, while those from the top layer (blue) had the lowest energy. 111
- 8.4 All events inside the marked areas from figure 8.3a are analysed: the energy spectra from simulation (a) and measurement (b) show the same formation of peaks. These peaks can be used to identify events from different crystal layers. 112
- 8.5 Simulated light distribution for events from the three crystal layers in the marked spots from figure 8.3. The highest fraction was detected on dSiPM 6 and its direct neighbours, which are indicated by dotted lines. The different ratios could be used to distinguish events from the three crystal layers. 113
- 8.6 Measured light distributions for events from the above defined spots. Compared to the simulation, the light fractions measured by dSiPM pixels 6, 9 and 10 differed slightly (marked with circles). However, the concept for event distinction could still be applied. 113
- 8.7 Standard deviation of the Anger positions, taken from the simulation. Figure (b) shows the standard deviation of all events inside the marked areas of figure (a), whereas figure (c) presents the standard deviation of events from the respective layers (selected with ground truth). . . . 114
- 8.8 Standard deviations of the Anger positions from the measured data. Figure (b) shows the standard deviation of all events inside the marked areas of figure (a), whereas figure (c) presents the standard deviation for energy-filtered events (presumably from the respective layers). . . . 114
- 8.9 The coloured areas in the flood histogram indicate those spots, which contained two crystal signals: bottom and intermediate layers are marked in red, top and bottom layers are indicated in blue, and top and intermediate layers are marked in green. 115
- 8.10 A single spot, which contained overlaying crystal signals of the bottom and the top layers, is shown in the simulated (a) and the measured flood map (d). The energy spectrum (b,e) and standard deviation (c,e) were used to separate the signals. 116
- 8.11 A single spot, which contained overlaying crystal signals of the top and intermediate layers, is shown in the simulated (a) and the measured flood map (d). The energy spectrum (b,e) could be used to separate the signals, while the standard deviation turned out to be not well separable. 117

- 8.12 A single spot, which contained overlaying crystal signals of the bottom and intermediate layers is shown in the simulated (a) and the measured flood map (d). The energy peaks of both layers completely overlay (b,e). The standard deviation could be used to separate the signals in the simulated data (c), while this was more difficult for the measured data (f). 118
- 8.13 The light distribution for events from the marked spot in figures 8.12a and 8.12d: the ground truth information from the simulation showed extremely similar distributions for both layers, which made it impossible to separate the spot by making use of this information only. The measured light distribution was in good agreement with the simulated ones. Part (b) shows the ratio of dSiPM 15 and 16: while small differences were visible in the simulation, there was no usable information in the measurement. 118
- 8.14 A single region of the flood map is shown in figure (a). According to the density estimates of the 1D histograms, the region was further divided into areas, which contained 0, 1 or 2 spots. By means of energy thresholding the clusters could be defined (b). 120
- 8.15 The cluster centres, which have been determined with the histogram-based method, were used as start parameters (a). Figure (b) shows the result of fuzzy c-means clustering and energy thresholding. 121
- 8.16 A single region of the flood map was divided into subsections, which contained 0, 1 or 2 spots (as in method 1, figure (a)). Model-based clustering was applied to identify the clusters and one noise cluster per subsection. 122
- 8.17 Comparison of different methods for clustering of the reference flood map and for event positioning. 124
- 8.18 Ratio of correctly identified events for the individual crystals in the flood histogram. The crystals which gave rise to overlapping spots are marked in blue, green and red, and the theoretical limit of 77 % is marked with a grey line. 126
- 8.19 Mean amount of correctly identified crystals for the three types of overlapping spots and the single crystal spots. 126
- 8.20 Recorded flood histograms of the crystal array with ground and polished surface finish are shown in figures (a) and (b). Figure (c) and (d) contain an enlarged view of the top left 4×4 dies. The spots at the edges were better separated in the ground finish crystal array (see red arrows). 129
- 8.21 Energy resolutions and peak positions of the two crystal arrays with rough and polished lateral faces. Both parameters were calculated separately for the three crystal layers. The error bars show the standard error of the mean. 131

-
- 8.22 Rough crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the top and bottom crystal layers. Due to the distinct energy peaks and the clearly separated standard deviations, both spots could be easily identified, as shown in figures (d)-(f). 133
- 8.23 Polished crystal array: measured energy spectra and standard deviations for a point, which contained partly overlapping signals from the top and bottom crystal layers. Due to the congruent energy peaks, the separation was complicated compared to the rough crystal array, but appeared to be successful (figures (d)-(f)). 134
- 8.24 Rough crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the intermediate and bottom crystal layers. Although the energy peaks and standard deviations were closer to each other, compared to the previous case, model-based clustering could be used to identify both crystal signals. The results are shown in figures (d)-(f). 135
- 8.25 Polished crystal array: measured energy spectra and standard deviations for a point, which contained overlapping signals from the intermediate and bottom crystal layers. The single energy peaks and standard deviations overlapped, making it impossible to separate the spots with model-based clustering. The results are shown in figures (d)-(f). 136
- 8.26 Rough crystal array: measured energy spectra and standard deviations for overlapping signals from the top and intermediate crystal layers. Despite the very close energy peaks, model-based clustering could identify both crystal signals (figures (d)-(f)). 137
- 8.27 Polished crystal array: measured energy spectra and standard deviations for three signals from the top and intermediate crystal layers. There was no overlap of these points, so that model-based clustering could identify all three crystal signals (figures (d)-(f)). 138

List of Tables

2.1	Overview of common PET tracers and their clinical applications. Republished and adapted with permission of Eureka Science (FZC), from "Clinical Applications of Positron Emission Tomography (PET) Imaging in Medicine: Oncology, Brain Diseases and Cardiology", S. Kitson et al, <i>Current Radiopharmaceuticals</i> , vol. 2, no. 4, 2009 [4]; permission conveyed through Copyright Clearance Center, Inc.	6
2.2	Comparison of properties of common scintillators used for nuclear imaging. This table was published in S. R. Cherry, J. A. Sorenson, and M. E. Phelps, "Radiation Detectors" in <i>Physics in Nuclear Medicine</i> , Fourth Edition, Oxford: Elsevier LTD, 2012, p. 101, Copyright Elsevier 2008 [23]. Data for LYSO:Ce was added according to reference [31]. . .	13
2.3	Overview of the selectable trigger schemes. Adapted and republished from reference [69]. ©IOP Publishing Ltd and Sissa Medialab srl. Reproduced by permission of IOP Publishing. All rights reserved. . . .	26
2.4	Overview of the selectable validation schemes (adapted from references [66], [69]).	27
3.1	Overview over different clinical and preclinical scanners. The spatial resolution is given as axial \times transverse resolution. Data compiled from the respective sources. *Measured at 1 cm radial distance from the centre FOV. **Tapered monolithic crystals with 40 – 50 mm width.	38
4.1	Overview of different possible crystal array sizes, which could be used for the new scanner. For comparison, the dimensions of the currently used Siemens BrainPET are given.	43
4.2	Ratio of detected true singles/coincidences and emitted photons for each scanner design. A change in geometry influences the sensitivity significantly.	47
4.3	The different scanner designs with the respective scintillator volumes and number of crystals. Option IV was additionally considered, being a presumably cheaper alternative.	48
6.1	Surface definitions of simulated optical interfaces and surfaces. R: Reflectivity, T: Transmittance, E: Efficiency.	70

6.2	Ground truth information on the three interactions of the high-energy photon shown in figure 6.2. The information was recorded with a GATE actor.	71
6.3	Updated definitions for the simulated optical interfaces and surfaces. The surfaces which are not mentioned in this table remain unchanged (see table 6.1). R: Reflectivity, T: Transmittance, E: Efficiency.	76
6.4	Overview of the selectable validation schemes, with corresponding names from the Philips DPC manual [66] and the publication of <i>Tabacchini et al.</i> [69].	82

Bibliography

- [1] D. A. Rich, "A Brief History of Positron Emission Tomography," *Journal of Nuclear Medicine Technology*, vol. 25, no. 1, pp. 4–11, 1997.
- [2] Z. Chen, S. D. Jamadar, S. Li, *et al.*, "From simultaneous to synergistic MR-PET brain imaging: A review of hybrid MR-PET imaging methodologies," *Human Brain Mapping*, vol. 39, no. 12, pp. 5126–5144, 2018.
- [3] J. Zheng, D. A. Jaffray, and C. Allen, "Nanosystems for Multimodality In vivo Imaging," in *Multifunctional Pharmaceutical Nanocarriers*, V. Torchilin, Ed., New York, NY: Springer, 2008, pp. 409–430, ISBN: 9780387765518.
- [4] S. Kitson, V. Cuccurullo, A. Ciarmiello, D. Salvo, and L. Mansi, "Clinical Applications of Positron Emission Tomography (PET) Imaging in Medicine: Oncology, Brain Diseases and Cardiology," *Current Radiopharmaceuticals*, vol. 2, no. 4, pp. 224–253, 2009.
- [5] S. Vandenberghe and P. K. Marsden, "PET-MRI: a review of challenges and solutions in the development of integrated multimodality imaging," *Physics in Medicine and Biology*, vol. 60, no. 4, R115–R154, 2015.
- [6] J. van Sluis, J. de Jong, J. Schaar, *et al.*, "Performance Characteristics of the Digital Biograph Vision PET/CT System," *Journal of Nuclear Medicine*, vol. 60, no. 7, pp. 1031–1036, 2019.
- [7] P. J. Slomka, T. Pan, and G. Germano, "Recent Advances and Future Progress in PET Instrumentation," *Seminars in Nuclear Medicine*, vol. 46, no. 1, pp. 5–19, 2016.
- [8] W. W. Moses, "Fundamental limits of spatial resolution in PET," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 648, S236–S240, 2011.
- [9] M. Ito, S. J. Hong, and J. S. Lee, "Positron Emission Tomography (PET) Detectors with Depth-Of-Interaction (DOI) Capability," *Biomedical Engineering Letters*, vol. 1, no. 2, pp. 70–81, 2011.
- [10] S. R. Cherry, T. Jones, J. S. Karp, *et al.*, "Total-Body PET: Maximizing Sensitivity to Create New Opportunities for Clinical Research and Patient Care," *Journal of Nuclear Medicine*, vol. 59, no. 1, pp. 3–12, 2017.
- [11] A. J. Gonzalez, F. Sanchez, and J. M. Benlloch, "Organ-Dedicated Molecular Imaging Systems," *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 2, no. 5, pp. 388–403, 2018.

- [12] E. S. Harmon, M. O. Thompson, C. R. Schmidlein, J. N. Turner, and A. Krol, "Towards 50 ps TOF-PET for brain imaging," in *Medical Imaging 2019: Biomedical Applications in Molecular, Structural, and Functional Imaging*, B. Gimi and A. Krol, Eds., International Society for Optics and Photonics, vol. 10953, SPIE, 2019, pp. 1095303-1 –1095303-11.
- [13] O. Kraff, A. Fischer, A. M. Nagel, C. Mönninghoff, and M. E. Ladd, "MRI at 7 Tesla and Above: Demonstrated and Potential Capabilities," *Journal of Magnetic Resonance Imaging*, vol. 41, no. 1, pp. 13–33, 2014.
- [14] K. Wienhard, M. Schmand, M. Casey, *et al.*, "The ECAT HRRT: Performance and First Clinical Application of the New High Resolution Research Tomograph," *IEEE Transactions on Nuclear Science*, vol. 49, no. 1, pp. 104–110, 2002.
- [15] H. Herzog, K.-J. Langen, C. Weirich, *et al.*, "High resolution BrainPET combined with simultaneous MRI," *Nuklearmedizin*, vol. 50, no. 2, pp. 74–82, 2011.
- [16] S. Jan, D. Benoit, E. Becheva, *et al.*, "GATE V6: a major enhancement of the GATE simulation platform enabling modelling of CT and radiotherapy," *Physics in Medicine and Biology*, vol. 56, no. 4, pp. 881–901, 2011.
- [17] P. Colombino, B. Fiscella, and L. Trossi, "Study of Positronium in Water and Ice from 22 to -144°C by Annihilation Quanta Measurements," *Il Nuovo Cimento*, vol. 38, no. 2, pp. 707–723, 1965.
- [18] G. Hevesy, "The Absorption and Translocation of Lead by Plants," *Biochemical Journal*, vol. 17, no. 4-5, pp. 439–445, 1923.
- [19] P. A. M. Dirac, "A Theory of Electrons and Protons," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 126, no. 801, pp. 360–365, 1930.
- [20] P. Lecoq, A. Annenkov, A. Gektin, M. Korzhik, and C. Pedrini, *Inorganic Scintillators for Detector Systems: Physical Principles and Crystal Engineering*, ser. Particle Acceleration and Detection. Berlin, Heidelberg, New York: Springer, 2006, ISBN: 9783540277668.
- [21] K. Kubota, "From tumor biology to clinical PET: A review of positron emission tomography (PET) in oncology," *Annals of Nuclear Medicine*, vol. 15, no. 6, pp. 471–486, 2001.
- [22] J. S. Fowler, N. D. Volkow, G. J. Wang, Y. S. Ding, and S. L. Dewey, "PET and Drug Research and Development," *Journal of Nuclear Medicine*, vol. 40, no. 7, pp. 1154–1163, 1999.
- [23] S. R. Cherry, J. A. Sorenson, and M. E. Phelps, *Physics in Nuclear Medicine*, Fourth edition. Oxford: Elsevier LTD, 2012, 544 pp.
- [24] E. Rota Kops and H. Herzog, "Alternative methods for attenuation correction for PET images in MR-PET scanners," in *2007 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE, 2007.

- [25] J. M. Ollinger, "Model-based scatter correction for fully 3D PET," *Physics in Medicine and Biology*, vol. 41, no. 1, pp. 153–176, 1996.
- [26] J. F. Oliver and M. Rafecas, "Improving the singles rate method for modeling accidental coincidences in high-resolution PET," *Physics in Medicine and Biology*, vol. 55, no. 22, pp. 6951–6971, 2010.
- [27] N. J. Shah, Ed., *Hybrid MR-PET Imaging: Systems, Methods and Applications*, ser. New Developments in NMR. Cambridge: The Royal Society of Chemistry, 2018, ISBN: 9781788010740.
- [28] M. Defrise, P. E. Kinahan, and C. J. Michel, "Image Reconstruction Algorithms in PET," in *Positron Emission Tomography: Basic Sciences*, London: Springer, 2005, pp. 63–91.
- [29] D. L. Bailey, D. W. Townsend, P. E. Valk, and M. N. Maisey, *Positron Emission Tomography*. London: Springer, 2005, ISBN: 1852337982.
- [30] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments*. Berlin, Heidelberg: Springer, 1994, ISBN: 9783540572800.
- [31] L. Pidol, A. Kahn-Harari, B. Viana, *et al.*, "High Efficiency of Lutetium Silicate Scintillators, Ce-doped LPS, and LYSO Crystals," *IEEE Transactions on Nuclear Science*, vol. 51, no. 3, pp. 1084–1087, 2004.
- [32] P. J. Slomka, T. Pan, D. S. Berman, and G. Germano, "Advances in SPECT and PET Hardware," *Progress in Cardiovascular Diseases*, vol. 57, no. 6, pp. 566–578, 2015.
- [33] M. Janecek and W. W. Moses, "Simulating Scintillator Light Collection Using Measured Optical Reflectance," *IEEE Transactions on Nuclear Science*, vol. 57, no. 3, pp. 964–970, 2010.
- [34] M. A. Miller. (2016). "Focusing on high performance." Manufacturer information, Koninklijke Philips: Advanced Molecular Imaging Physics, [Online]. Available: <https://philipsproductcontent.blob.core.windows.net/assets/20170523/360753349c5d4a6aa46ba77c015e75b4.pdf> (visited on 01/14/2020).
- [35] H. O. Anger, "Scintillation Camera," *Review of Scientific Instruments*, vol. 29, no. 1, pp. 27–33, 1958.
- [36] C. J. Marriott, J. E. Cadorette, R. Lecomte, *et al.*, "High-Resolution PET Imaging and Quantitation of Pharmaceutical Biodistributions in a Small Animal Using Avalanche Photodiode Detectors," *Journal of Nuclear Medicine*, vol. 35, no. 8, pp. 1390–1396, 1994.
- [37] M. Moszyński, M. Kapusta, A. Nassalski, *et al.*, "New Prospects for Time-of-Flight PET With LSO Scintillators," *IEEE Transactions on Nuclear Science*, vol. 53, no. 5, pp. 2484–2488, 2006.

- [38] M. M. Bé, V. Chisté, C. Dulieu, *et al.*, *Table of Radionuclides (Vol. 5–A= 22 to 244)*, ser. Monographie BIPM-5. Sèvres: Bureau International des Poids et Mesures, 2010, vol. 5.
- [39] G. F. Knoll, *Radiation Detection and Measurement*, Fourth edition. New York, NY: Wiley, 2010, 864 pp.
- [40] M. Moszyński, “Energy resolution of scintillation detectors,” in *Hard X-Ray and Gamma-Ray Detector Physics VII*, R. B. James, L. A. Franks, and A. Burger, Eds., vol. 5922, SPIE, 2005, pp. 16–29.
- [41] M. Balcerzyk, M. Moszyński, M. Kapusta, *et al.*, “YSO, LSO, GSO and LGSO. A Study of Energy Resolution and Nonproportionality,” *IEEE Transactions on Nuclear Science*, vol. 47, no. 4, pp. 1319–1323, 2000.
- [42] M. Grodzicka, M. Moszyński, T. Szcześniak, *et al.*, “Energy resolution of small scintillation detectors with SiPM light readout,” *Journal of Instrumentation*, vol. 8, no. 2, P02017, 2013.
- [43] G. Borghi, V. Tabacchini, and D. R. Schaart, “Towards monolithic scintillator based TOF-PET systems: practical methods for detector calibration and operation,” *Physics in Medicine and Biology*, vol. 61, no. 13, pp. 4904–4928, 2016.
- [44] R. Vinke and C. S. Levin, “A method to achieve spatial linearity and uniform resolution at the edges of monolithic scintillation crystal detectors,” *Physics in Medicine and Biology*, vol. 59, no. 12, pp. 2975–2995, 2014.
- [45] C. Levin, “Design of a High-Resolution and High-Sensitivity Scintillation Crystal Array for PET With Nearly Complete Light Collection,” *IEEE Transactions on Nuclear Science*, vol. 49, no. 5, pp. 2236–2243, 2002.
- [46] J. D. Thiessen, M. A. Koschan, C. L. Melcher, *et al.*, “A phoswich detector design for improved spatial sampling in PET,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 882, pp. 124–128, 2018.
- [47] R. Bugalho, B. Carriço, C. S. Ferreira, *et al.*, “Experimental characterization of the Clear-PEM scanner spectrometric performance,” *Journal of Instrumentation*, vol. 4, no. 10, P10011, 2009.
- [48] Y. Yang, Y. Wu, J. Qi, *et al.*, “A Prototype PET Scanner with DOI-Encoding Detectors,” *Journal of Nuclear Medicine*, vol. 49, no. 7, pp. 1132–1140, 2008.
- [49] M. Pizzichemi, G. Stringhini, T. Niknejad, *et al.*, “A new method for depth of interaction determination in PET detectors,” *Physics in Medicine and Biology*, vol. 61, no. 12, pp. 4679–4698, 2016.
- [50] M. Pizzichemi, “Positron Emission Tomography: state of the art and future developments,” *Journal of Instrumentation*, vol. 11, no. 8, p. C08004, 2016.

- [51] M. Ito, M. S. Lee, and J. S. Lee, "Continuous depth-of-interaction measurement in a single-layer pixelated crystal array using a single-ended readout," *Physics in Medicine and Biology*, vol. 58, no. 5, pp. 1269–1282, 2013.
- [52] A. Vandenbroucke, A. M. K. Foudray, P. D. Olcott, and C. S. Levin, "Performance characterization of a new high resolution PET scintillation detector," *Physics in Medicine and Biology*, vol. 55, no. 19, pp. 5895–5911, 2010.
- [53] S. J. Hong, S. I. Kwon, M. Ito, *et al.*, "Concept Verification of Three-Layer DOI Detectors for Small Animal PET," *IEEE Transactions on Nuclear Science*, vol. 55, no. 3, pp. 912–917, 2008.
- [54] T. Xu, Q. Wei, G. Gong, Y. Liu, and T. Ma, "Development of a MR-compatible DOI-TOF Detector Module for PET Imaging Systems," in *2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD)*, IEEE, 2016.
- [55] Y. H. Chung, J. Y. Hwang, C.-H. Baek, *et al.*, "Monte Carlo simulation of a four-layer DOI detector with relative offset in animal PET," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 626-627, pp. 43–50, 2011.
- [56] F. Nishikido, N. Inadama, E. Yoshida, H. Murayama, and T. Yamaya, "Four-layer DOI PET detectors using a multi-pixel photon counter array and the light sharing method," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 729, pp. 755–761, 2013.
- [57] N. Inadama, H. Murayama, M. Hamamoto, *et al.*, "8-Layer DOI Encoding of 3-Dimensional Crystal Array," *IEEE Transactions on Nuclear Science*, vol. 53, no. 5, pp. 2523–2528, 2006.
- [58] C. W. Lerche, A. Salomon, B. Goldschmidt, *et al.*, "Maximum likelihood positioning and energy correction for scintillation detectors," *Physics in Medicine and Biology*, vol. 61, no. 4, pp. 1650–1676, 2016.
- [59] D. Renker and E. Lorenz, "Advances in solid state photon detectors," *Journal of Instrumentation*, vol. 4, no. 4, P04004, 2009.
- [60] F. Zappa, S. Tisa, A. Tosi, and S. Cova, "Principles and features of single-photon avalanche diode arrays," *Sensors and Actuators A: Physical*, vol. 140, no. 1, pp. 103–112, 2007.
- [61] R. Lecomte, "Novel detector technology for clinical PET," *European Journal of Nuclear Medicine and Molecular Imaging*, vol. 36, no. S1, pp. 69–85, 2008.
- [62] A. G. Stewart, V. Saveliev, S. J. Bellis, *et al.*, "Performance of 1-mm² Silicon Photomultiplier," *IEEE Journal of Quantum Electronics*, vol. 44, no. 2, pp. 157–164, 2008.

- [63] S. Cova, A. Lacaita, and G. Ripamonti, "Trapping Phenomena in Avalanche Photodiodes on Nanosecond Scale," *IEEE Electron Device Letters*, vol. 12, no. 12, pp. 685–687, 1991.
- [64] A. Lacaita, F. Zappa, S. Bigliardi, and M. Manfredi, "On the Bremsstrahlung Origin of Hot-Carrier-Induced Photons in Silicon Devices," *IEEE Transactions on Electron Devices*, vol. 40, no. 3, pp. 577–582, 1993.
- [65] P. Eckert, H.-C. Schultz-Coulon, W. Shen, R. Stamen, and A. Tadday, "Characterisation studies of silicon photomultipliers," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 620, no. 2-3, pp. 217–226, 2010.
- [66] Philips Digital Photon Counting, "Tile-TEK – User Manual v1.02," Koninklijke Philips N.V., Apr. 2016, User Manual.
- [67] C. Degenhardt, G. Prescher, T. Frach, *et al.*, "The Digital Silicon Photomultiplier: A Novel Sensor for the Detection of Scintillation Light," in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE, 2009.
- [68] T. Frach, G. Prescher, C. Degenhardt, *et al.*, "The Digital Silicon Photomultiplier: Principle of Operation and Intrinsic Detector Performance," in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE, 2009.
- [69] V. Tabacchini, V. Westerwoudt, G. Borghi, S. Seifert, and D. R. Schaart, "Probabilities of triggering and validation in a digital silicon photomultiplier," *Journal of Instrumentation*, vol. 9, no. 6, P06016, 2014.
- [70] S. Gundacker, E. Auffray, K. Pauwels, and P. Lecoq, "Measurement of intrinsic rise times for various L(Y)SO and LuAG scintillators with a general study of prompt photons to achieve 10 ps in TOF-PET," *Physics in Medicine and Biology*, vol. 61, no. 7, pp. 2802–2837, 2016.
- [71] D. Schug, J. Wehner, B. Goldschmidt, *et al.*, "Data Processing for a High Resolution Preclinical PET Detector Based on Philips DPC Digital SiPMs," *IEEE Transactions on Nuclear Science*, vol. 62, no. 3, pp. 669–678, 2015.
- [72] D. Hu, B. E. Atkins, M. W. Lenox, B. Castleberry, and S. B. Siegel, "A Neural Network Based Algorithm for Building Crystal Look-up Table of PET Block Detector," in *2006 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE, 2006.
- [73] S. España, R. Marcinkowski, V. Keereman, S. Vandenberghe, and R. V. Holen, "DigiPET: sub-millimeter spatial resolution small-animal PET imaging using thin monolithic scintillators," *Physics in Medicine and Biology*, vol. 59, no. 13, pp. 3405–3420, 2014.
- [74] I. J. Myung, "Tutorial on maximum likelihood estimation," *Journal of Mathematical Psychology*, vol. 47, no. 1, pp. 90–100, 2003.

- [75] H. H. Barrett, W. C. J. Hunter, B. W. Miller, *et al.*, "Maximum-Likelihood Methods for Processing Signals From Gamma-Ray Detectors," *IEEE Transactions on Nuclear Science*, vol. 56, no. 3, pp. 725–735, 2009.
- [76] N. Gross-Weege, D. Schug, P. Hallen, and V. Schulz, "Maximum likelihood positioning algorithm for high-resolution PET scanners," *Medical Physics*, vol. 43, no. 6Part1, pp. 3049–3061, 2016.
- [77] T. Beyer and D. W. Townsend, "Putting 'clear' into nuclear medicine: a decade of PET/CT development," *European Journal of Nuclear Medicine and Molecular Imaging*, vol. 33, no. 8, pp. 857–861, 2006.
- [78] G. Delso, S. Furst, B. Jakoby, *et al.*, "Performance Measurements of the Siemens mMR Integrated Whole-Body PET/MR Scanner," *Journal of Nuclear Medicine*, vol. 52, no. 12, pp. 1914–1922, 2011.
- [79] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince, *MRI From Picture to Proton*. Cambridge: Cambridge University Press, 2006, ISBN: 9780521865272.
- [80] D. W. Harberts and M. van Helvoort, "MRI image distortion due to magnetic materials in medical implants," in *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, IEEE, 2015.
- [81] J. F. Schenck, "The role of magnetic susceptibility in magnetic resonance imaging: MRI magnetic compatibility of the first and second kinds," *Medical Physics*, vol. 23, no. 6, pp. 815–850, 1996.
- [82] D. Strul, D. Cash, S. Keevil, *et al.*, "Gamma Shielding Materials for MR-Compatible PET," *IEEE Transactions on Nuclear Science*, vol. 50, no. 1, pp. 60–69, 2003.
- [83] A. Berneking, R. Trincherro, Y. Ha, *et al.*, "Design and Characterization of a Gradient-Transparent RF Copper Shield for PET Detector Modules in Hybrid MR-PET Imaging," *IEEE Transactions on Nuclear Science*, vol. 64, no. 5, pp. 1118–1127, 2017.
- [84] B. J. Lee, R. D. Watkins, C.-M. Chang, and C. S. Levin, "Low Eddy Current RF Shielding Enclosure Designs for 3T MR Applications," *Magnetic Resonance in Medicine*, vol. 79, no. 3, pp. 1745–1752, 2017.
- [85] S. Agostinelli, J. Allison, K. Amako, *et al.*, "Geant4— a simulation toolkit," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [86] D. Sarrut, M. Bardiès, N. Boussion, *et al.*, "A review of the use and potential of the GATE Monte Carlo simulation code for radiation therapy and dosimetry applications," *Medical Physics*, vol. 41, no. 6, pp. 064301-1–064301-14, 2014.
- [87] N. Metropolis, "The Beginning of the Monte Carlo Method," *Los Alamos Science Special Issue*, vol. 15, pp. 125–130, 1987.

- [88] P. Kowalski, P. Moskal, W. Wiślicki, *et al.*, "Multiple scattering and accidental coincidences in the J-PET detector simulated using GATE package," *Acta Physica Polonica A*, vol. 127, no. 5, pp. 1505–1512, 2015.
- [89] S. Surti and J. S. Karp, "Impact of event positioning algorithm on performance of a whole-body PET scanner using one-to-one coupled detectors," *Physics in Medicine and Biology*, vol. 63, no. 5, p. 055 008, 2018.
- [90] I. Somlai-Schweiger and S. I. Ziegler, "CHERENCUBE: Concept definition and implementation challenges of a Cherenkov-based detector block for PET," *Medical Physics*, vol. 42, no. 4, pp. 1825–1835, 2015.
- [91] R. Ota, R. Yamada, T. Moriya, and T. Hasegawa, "Cherenkov radiation-based three-dimensional position-sensitive PET detector: A Monte Carlo study," *Medical Physics*, vol. 45, no. 5, pp. 1999–2008, 2018.
- [92] E. Berg, E. Roncali, and S. R. Cherry, "Optimizing light transport in scintillation crystals for time-of-flight PET: an experimental and optical Monte Carlo simulation study," *Biomedical Optics Express*, vol. 6, no. 6, pp. 2220–2230, 2015.
- [93] D. J. van der Laan, D. R. Schaart, M. C. Maas, *et al.*, "Optical simulation of monolithic scintillator detectors using GATE/GEANT4," *Physics in Medicine and Biology*, vol. 55, no. 6, pp. 1659–1675, 2010.
- [94] B. Aklan, B. W. Jakoby, C. C. Watson, *et al.*, "GATE Monte Carlo simulations for variations of an integrated PET/MR hybrid imaging system based on the Biograph mMR model," *Physics in Medicine and Biology*, vol. 60, no. 12, pp. 4731–4752, 2015.
- [95] A. Ghabrial, D. Franklin, and H. Zaidi, "A Monte Carlo simulation study of the impact of novel scintillation crystals on performance characteristics of PET scanners," *Physica Medica*, vol. 50, pp. 37–45, 2018.
- [96] S. Jan, G. Santin, D. Strul, *et al.*, "GATE: a simulation toolkit for PET and SPECT," *Physics in Medicine and Biology*, vol. 49, no. 19, pp. 4543–4561, 2004.
- [97] V. Cuplov, I. Buvat, F. Pain, and S. Jan, "Extension of the GATE Monte-Carlo simulation package to model bioluminescence and fluorescence imaging," *Journal of Biomedical Optics*, vol. 19, no. 2, pp. 026004-1–026004-12, 2014.
- [98] H. W. A. M. de Jong, F. H. P. van Velden, R. W. Kloet, *et al.*, "Performance evaluation of the ECAT HRRT: an LSO-LYSO double layer high resolution, high sensitivity scanner," *Physics in Medicine and Biology*, vol. 52, no. 5, pp. 1505–1526, 2007.
- [99] B. Aubert-Broche, M. Griffin, G. Pike, A. Evans, and D. Collins, "Twenty New Digital Brain Phantoms for Creation of Validation Image Data Bases," *IEEE Transactions on Medical Imaging*, vol. 25, no. 11, pp. 1410–1416, 2006.

- [100] N. Palomero-Gallagher and K. Zilles, "Cortical layers: Cyto-, myelo-, receptor- and synaptic architecture in human cortical areas," *NeuroImage*, vol. 197, pp. 716–741, 2019.
- [101] M. Hodolic, R. Topakian, and R. Pichler, "¹⁸F-fluorodeoxyglucose and ¹⁸F-flumazenil positron emission tomography in patients with refractory epilepsy," *Radiology and Oncology*, vol. 50, no. 3, pp. 247–253, 2016.
- [102] C. Catana, A. Drzezga, W.-D. Heiss, and B. R. Rosen, "PET/MRI for Neurologic Applications," *Journal of Nuclear Medicine*, vol. 53, no. 12, pp. 1916–1925, 2012.
- [103] M. Z. Pajak, D. Volgyes, S. L. Pimlott, *et al.*, "NEMA NU4-2008 Performance Evaluation of Albira: A Two-Ring Small-Animal PET System Using Continuous LYSO Crystals," *Open Medicine Journal*, vol. 3, no. 1, pp. 12–26, 2016.
- [104] D. McElroy, W. Pimpl, B. Pichler, *et al.*, "Characterization and Readout of MADPET-II Detector Modules: Validation of a Unique Design Concept for High Resolution Small Animal PET," *IEEE Transactions on Nuclear Science*, vol. 52, no. 1, pp. 199–204, 2005.
- [105] T. Yamaya, N. Hagiwara, T. Obi, *et al.*, "Preliminary Resolution Performance of the Prototype System for a 4-Layer DOI-PET Scanner: jPET-D4," *IEEE Transactions on Nuclear Science*, vol. 53, no. 3, pp. 1123–1128, 2006.
- [106] T. Yamaya, E. Yoshida, T. Obi, *et al.*, "First Human Brain Imaging by the jPET-D4 Prototype With a Pre-Computed System Matrix," *IEEE Transactions on Nuclear Science*, vol. 55, no. 5, pp. 2482–2492, 2008.
- [107] M. Watanabe, A. Saito, T. Isobe, *et al.*, "Performance evaluation of a high-resolution brain PET scanner using four-layer MPPC DOI detectors," *Physics in Medicine and Biology*, vol. 62, no. 17, pp. 7148–7166, 2017.
- [108] M. Ito, S. J. Hong, J. S. Lee, *et al.*, "Four-layer DOI Detector with a Relative Offset in Animal PET System," in *2007 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE, 2007.
- [109] L. Moliner, M. J. Rodriguez-Alvarez, J. V. Catret, *et al.*, "NEMA Performance Evaluation of CareMiBrain dedicated brain PET and Comparison with the whole-body and dedicated brain PET systems," *Scientific Reports*, vol. 9, no. 1, 2019.
- [110] A. J. González, S. Majewski, F. Sánchez, *et al.*, "The MINDView brain PET detector, feasibility study based on SiPM arrays," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 818, pp. 82–90, 2016.
- [111] J. M. Benlloch, A. J. González, R. Pani, *et al.*, "The MINDVIEW project: First results," *European Psychiatry*, vol. 50, pp. 21–27, 2018.

- [112] C. Catana, "Development of Dedicated Brain PET Imaging Devices: Recent Advances and Future Perspectives," *Journal of Nuclear Medicine*, vol. 60, no. 8, pp. 1044–1052, 2019.
- [113] A. M. Grant, T. W. Deller, M. M. Khalighi, *et al.*, "NEMA NU 2-2012 performance studies for the SiPM-based ToF-PET component of the GE SIGNA PET/MR system," *Medical Physics*, vol. 43, no. 5, pp. 2334–2343, 2016.
- [114] A. M. Karlberg, O. Sæther, L. Eikenes, and P. E. Goa, "Quantitative comparison of PET performance—Siemens Biograph mCT and mMR," *EJNMMI Physics*, vol. 3, no. 1, 2016.
- [115] C. C. Constantinescu and J. Mukherjee, "Performance evaluation of an Inveon PET preclinical scanner," *Physics in Medicine and Biology*, vol. 54, no. 9, pp. 2885–2899, 2009.
- [116] Mediso Medical Imaging Systems. (2017). "nanoScan[®] PET: Preclinical Imaging Systems." Manufacturer information, [Online]. Available: https://www.mediso.de/media/Broschuere%20nanoScan_PET-CT_%20PET-MR_2017%2008_web.pdf (visited on 09/18/2019).
- [117] Bruker BioSpin. (2016). "Albira Si." Manufacturer information, [Online]. Available: https://www.bruker.com/fileadmin/user_upload/8-PDF-Docs/PreclinicalImaging/Brochures/AlbiraSi_Brochure_T14907.pdf (visited on 08/27/2018).
- [118] S. H. Maramraju, S. D. Smith, S. S. Junnarkar, *et al.*, "Small animal simultaneous PET/MRI: initial experiences in a 9.4T microMRI," *Physics in Medicine and Biology*, vol. 56, no. 8, pp. 2459–2480, 2011.
- [119] N. J. Shah, H. Herzog, C. Weirich, *et al.*, "Effects of Magnetic Fields of up to 9.4T on Resolution and Contrast of PET Images as Measured with an MR-BrainPET," *PLoS ONE*, vol. 9, no. 4, e95250, 2014.
- [120] J. H. Duyn, "The future of ultra-high field MRI and fMRI for study of the human brain," *NeuroImage*, vol. 62, no. 2, pp. 1241–1248, 2012.
- [121] 3M Electronic Display Lighting Optical Systems Division. (2018). "Vikuiti[™] Enhanced Specular Reflector (ESR)." Manufacturer information, [Online]. Available: <http://multimedia.3m.com/mws/media/4967180/vikuiti-enhanced-specular-reflector-esr.pdf> (visited on 07/01/2018).
- [122] Schott Technical Glass Solutions GmbH, "BOROFLOAT[®]33 – Optische Eigenschaften," 2019, Manufacturer information. [Online]. Available: https://www.schott.com/d/borofloat/4569b380-94c1-4139-ae58-51b1fe35c70f/1.3/borofloat33_opt_de_web.pdf (visited on 09/11/2019).
- [123] S. Surti, J. Karp, R. Freifelder, and F. Liu, "Optimizing the Performance of a PET Detector using Discrete GSO Crystals on a Continuous Lightguide," *IEEE Transactions on Nuclear Science*, vol. 47, no. 3, pp. 1030–1036, 2000.

- [124] National Electrical Manufacturers Association, *Performance measurements of positron emission tomographs*, ser. NEMA Standard Publication NU 2-2007. Rosslyn, VA, 2007.
- [125] H. Herzog, L. Tellmann, C. Hocke, *et al.*, "NEMA NU2-2001 Guided Performance Evaluation of Four Siemens ECAT PET Scanners," *IEEE Transactions on Nuclear Science*, vol. 51, no. 5, pp. 2662–2669, 2004.
- [126] A. Kolb, H. F. Wehrl, M. Hofmann, *et al.*, "Technical performance evaluation of a human brain PET/MRI system," *European Radiology*, vol. 22, no. 8, pp. 1776–1788, 2012.
- [127] H. Thoen, V. Keereman, P. Mollet, R. V. Holen, and S. Vandenberghe, "Influence of detector pixel size, TOF resolution and DOI on image quality in MR-compatible whole-body PET," *Physics in Medicine and Biology*, vol. 58, no. 18, pp. 6459–6479, 2013.
- [128] T. F. Budinger, S. E. Derenzo, G. T. Gullberg, W. L. Greenberg, and R. H. Huesman, "Emission Computer Assisted Tomography with Single-Photon and Positron Annihilation Photon Emitters," *Journal of Computer Assisted Tomography*, vol. 1, no. 1, pp. 131–145, 1977.
- [129] D. Krause and P. Thörnig, "JURECA: Modular supercomputer at Jülich Supercomputing Centre," *Journal of large-scale research facilities JLSRF*, vol. 4, A132, 2018.
- [130] C. W. Lerche, T. Kaltsas, L. Caldeira, *et al.*, "PET attenuation correction for rigid MR Tx/Rx coils from ¹⁷⁶Lu background activity," *Physics in Medicine and Biology*, vol. 63, no. 3, p. 035 039, 2018.
- [131] M. M. Bé, V. Chisté, C. Dulieu, *et al.*, *Table of Radionuclides (Vol. 7–A= 14 to 245)*, ser. Monographie BIPM-5. Sèvres: Bureau International des Poids et Mesures, 2013, vol. 7.
- [132] R. Brun and F. Rademakers, "ROOT– An object oriented data analysis framework," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 389, no. 1-2, pp. 81–86, 1997.
- [133] RStudio Team, *Rstudio: Integrated Development Environment for R*, RStudio, Inc., Boston, MA, 2016. [Online]. Available: <http://www.rstudio.com/>.
- [134] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2017. [Online]. Available: <https://www.R-project.org/>.
- [135] C. W. Lerche, S. Lodomez, V. Schulz, and B. Weissler, *Apparatus and Method for the Evaluation of Gamma Radiation Events*, Patent application PCT/EP2014/058865, International publication number WO 2014/180734 A2, 2014.

- [136] J. Kang, Y. Choi, K. J. Hong, *et al.*, "PET Detector Configuration with Thick Light Guide and GAPD Array Having Large-Area Microcells," in *IEEE Nuclear Science Symposium and Medical Imaging Conference*, IEEE, 2010.
- [137] A. Levin and C. Moisan, "A More Physical Approach to Model the Surface Treatment of Scintillation Counters and its Implementation into DETECT," in *1996 IEEE Nuclear Science Symposium. Conference Record*, IEEE, 1996.
- [138] S. Nayar, K. Ikeuchi, and T. Kanade, "Surface Reflection: Physical and Geometrical Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 611–634, 1991.
- [139] E. Lorincz, G. Erdei, I. Peczei, *et al.*, "Modeling and Optimization of Scintillator Arrays for PET Detectors," *IEEE Transactions on Nuclear Science*, vol. 57, no. 1, pp. 48–54, 2010.
- [140] M. Janecek, "Reflectivity Spectra for Commonly Used Reflectors," *IEEE Transactions on Nuclear Science*, vol. 59, no. 3, pp. 490–497, 2012.
- [141] E. Leming, A. D. Santo, F. Salvatore, B. Camanzi, and A. Lohstroh, "A GEANT4 Monte Carlo simulation to describe the time response of a coupled SiPM and LYSO detection system," *Journal of Instrumentation*, vol. 9, no. 6, p. C06008, 2014.
- [142] R. Mao, L. Zhang, and R.-Y. Zhu, "Optical and Scintillation Properties of Inorganic Scintillators in High Energy Physics," *IEEE Transactions on Nuclear Science*, vol. 55, no. 4, pp. 2425–2431, 2008.
- [143] Heraeus Quarzglas GmbH & Co. KG. (2011). "Quarzglas für die Optik: Daten und Eigenschaften." Manufacturer information, [Online]. Available: https://www.heraeus.com/media/media/hca/doc_hca/products_and_solutions_8/optics/Daten_und_Eigenschaften_Quarzglas_fuer_die_Optik_DE.pdf (visited on 05/03/2018).
- [144] E. Roncali, M. Stockhoff, and S. R. Cherry, "An integrated model of scintillator-reflector properties for advanced simulations of optical transport," *Physics in Medicine and Biology*, vol. 62, no. 12, pp. 4811–4830, 2017.
- [145] M. Stockhoff, S. Jan, A. Dubois, S. R. Cherry, and E. Roncali, "Advanced optical simulation of scintillation detectors in GATE V8.0: first implementation of a reflectance model based on measured data," *Physics in Medicine and Biology*, vol. 62, no. 12, pp. L1–L8, 2017.
- [146] J. Nilsson, V. Cuplov, and M. Isaksson, "Identifying key surface parameters for optical photon transport in GEANT4/GATE simulations," *Applied Radiation and Isotopes*, vol. 103, pp. 15–24, 2015.
- [147] E. Dietz-Laursonn, *Peculiarities in the Simulation of Optical Physics with Geant4*, 2016. arXiv: 1612.05162v1 [physics.ins-det].

- [148] E. Auffray, B. Frisch, F. Geraci, *et al.*, "A Comprehensive & Systematic Study of Coincidence Time Resolution and Light Yield Using Scintillators of Different Size and Wrapping," *IEEE Transactions on Nuclear Science*, vol. 60, no. 5, pp. 3163–3171, 2013.
- [149] B. J. Pichler, F. Bernecker, G. Boning, *et al.*, "A 4×8 APD Array, Consisting of Two Monolithic Silicon Wafers, Coupled to a 32-Channel LSO Matrix for High-Resolution PET," *IEEE Transactions on Nuclear Science*, vol. 48, no. 4, pp. 1391–1396, 2001.
- [150] F. Loignon-Houle, C. M. Pepin, S. A. Charlebois, and R. Lecomte, "Reflectivity quenching of ESR multilayer polymer film reflector in optically bonded scintillator arrays," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 851, pp. 62–67, 2017.
- [151] F. Padera and C. Lynch. (2015). "Measurement of Enhanced Specular Reflector (ESR) Films Using a LAMBDA 1050 UV/Vis/NIR Spectrometer and URA Accessory." Field Application Report, PerkinElmer, Inc., [Online]. Available: https://www.perkinelmer.com/uk/lab-solutions/resources/docs/FAR_Measurement-of-Enhanced-Specular-Reflector-Films-Using-LAMBDA-1050-and-URA-Accessory-012190_01.pdf (visited on 03/01/2019).
- [152] 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," *Journal of Instrumentation*, vol. 7, no. 8, P08011, 2012.
- [153] T. Niggemann, E. Dietz-Laursonn, T. Hebbeker, *et al.*, "G4SiPM: A novel silicon photomultiplier simulation package for Geant4," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 787, pp. 344–347, 2015.
- [154] M. Basunia, "Nuclear Data Sheets for A = 176," *Nuclear Data Sheets*, vol. 107, no. 4, pp. 791–1026, 2006. [Online]. Available: <https://www.nndc.bnl.gov/nudat2/decaysearchdirect.jsp?nuc=176LU&unc=nds> (visited on 09/15/2019).
- [155] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*, ser. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC press, 1986, 176 pp., ISBN: 0412246201.
- [156] S. J. Sheather and M. C. Jones, "A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 53, no. 3, pp. 683–690, 1991.
- [157] V. John, I. Angelov, A. A. Öncül, and D. Thévenin, "Techniques for the reconstruction of a distribution from a finite number of its moments," *Chemical Engineering Science*, vol. 62, no. 11, pp. 2890–2904, 2007.

- [158] C. W. Lerche, V. Herrero-Bosch, M. Spaggiari, *et al.*, "Fast Circuit Topology for Spatial Signal Distribution Analysis," in *2010 17th IEEE-NPSS Real Time Conference*, IEEE, 2010.
- [159] C. W. Lerche, V. Herrero-Bosch, N. Ferrando-Jodar, *et al.*, "Efficient readout electronics for multi-anode photomultiplier," in *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics XII*, A. Burger, L. A. Franks, and R. B. James, Eds., SPIE, 2010.
- [160] V. C. Spanoudaki and C. S. Levin, "Investigating the temporal resolution limits of scintillation detection from pixellated elements: comparison between experiment and simulation," *Physics in Medicine and Biology*, vol. 56, no. 3, pp. 735–756, 2011.
- [161] S. S. James, Y. Yang, Y. Wu, *et al.*, "Experimental characterization and system simulations of depth of interaction PET detectors using 0.5 mm and 0.7 mm LSO arrays," *Physics in Medicine and Biology*, vol. 54, no. 14, pp. 4605–4619, 2009.
- [162] J. C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms (Advanced Applications in Pattern Recognition)*. New York: Plenum Press, 1981, ISBN: 0306406713.
- [163] K.-L. Wu, "Analysis of parameter selections for fuzzy c-means," *Pattern Recognition*, vol. 45, no. 1, pp. 407–415, 2012.
- [164] J. Nayak, B. Naik, and H. S. Behera, "Fuzzy C-Means (FCM) Clustering Algorithm: A Decade Review from 2000 to 2014," in *Computational Intelligence in Data Mining - Volume 2*, New Delhi: Springer, 2014, pp. 133–149.
- [165] L. Scrucca, M. Fop, T. B. Murphy, and A. E. Raftery, "mclust 5: Clustering, Classification and Density Estimation Using Gaussian Finite Mixture Models," *The R Journal*, vol. 8, no. 1, pp. 205–233, 2016.
- [166] M. Z. Rodriguez, C. H. Comin, D. Casanova, *et al.*, "Clustering algorithms: A comparative approach," *PLoS ONE*, vol. 14, no. 1, e0210236, 2019.
- [167] A. J. Chaudhari, A. A. Joshi, S. L. Bowen, *et al.*, "Crystal identification in positron emission tomography using nonrigid registration to a Fourier-based template," *Physics in Medicine and Biology*, vol. 53, no. 18, pp. 5011–5027, 2008.
- [168] Q. Wei, T. Ma, T. Xu, *et al.*, "Crystal identification for a dual-layer-offset LYSO based PET system via Lu-176 background radiation and mean shift algorithm," *Physics in Medicine and Biology*, vol. 63, no. 2, 02NT01, 2018.