

DISS. ETH NO. 30327

The SAFIR-II Positron Emission Tomography Insert

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by

Jan Friedrich Debus

M.SC., Ruprecht-Karls University of Heidelberg
born on 13.01.1995

accepted on the recommendation of

Prof. Dr. Günther Dissertori
Prof. Dr. Peter Fischer
Dr. Werner Lustermann

2024

Abstract

This thesis presents the design and individual components of the SAFIR-II (Small Animal Fast Insert for MRI detector II) high-performance PET (Positron Emission Tomography) scanner. Created to acquire preclinical PET images within five seconds, the scanner enables the study of fast kinematic processes within small rodents. This is made possible through the ability to process data at a high targeted measurement activity of up to 500 MBq. Designed as an MRI (Magnetic Resonance Imaging)-compatible PET insert, it can acquire images simultaneously with the Bruker BioSpec 70/30 MRI it operates in.

Additionally presented are evaluations of the scanner's performance using various measurements and tests. The detector can acquire data at measurement activities of up to 500 MBq while losing less than 5% of its data to detector dead time effects. It exhibits a coincidence time resolution of 221 ps and an energy resolution of 12.1%. These values are shown to depend on the measurement activity and on the timing threshold of the front-end ASIC (Application Specific Integrated Circuit) using a dedicated line-source measurement. A sensitivity profile was evaluated following the 2008 NEMA NU4 standard, reaching a peak sensitivity of 3.89% when using ICSR (Inter-Crystal-Scatter-Recovery). SAFIR-II's compatibility with the MRI system was demonstrated using quality assurance sequences provided by the MRI manufacturer. A thorough investigation of the scanner's image quality was also performed. A spatial resolution of down to 1.7 mm was observed using the measurement of a Derenzo-Hot rod phantom. An image quality investigation according to the NEMA NU4 standard yielded a Uniformity of 2.96% at 500 MBq, outperforming several commercial scanners currently available. Satisfactory results were observed for the SOR (Spill-over-Ratio) of the Air and Water chambers, with values of 0.057 and 0.075 being comparable to those of a commercial scanner. The RC (Recovery Coefficient) values were 0.92, 0.79, 0.58, 0.32, and 0.053 for the 5 mm, 4 mm, 3 mm, 2 mm, and 1 mm rods, respectively. The impact of the various image corrections on the scanner's image quality was examined, with the normalization correction exhibiting a higher significance compared to the predecessor system SAFIR-I. An in-vivo study featuring a Sprague-Dawley rat was performed at an injected activity of 283 MBq. Images reconstructed using as little as ten seconds of data from this study show SAFIR-II's capability to resolve small anatomical structures such as the myocardium.

Zusammenfassung

Diese Arbeit beschreibt den Aufbau und die individuellen Komponenten des SAFIR-II (Small Animal Fast Insert for MRI detector II) Hochleistungs-PET (Positronen-Emissions-Tomographie) Detektors. SAFIR-II wurde erbaut um präklinische PET-Bilder innerhalb von fünf Sekunden aufnehmen zu können, wodurch Studien zur Kinetik von schnell ablaufenden metabolischen Prozessen in Mäusen und Ratten ermöglicht werden. Dies wird dadurch erreicht, dass der Detektor Daten bei einer hohen Mess-Aktivität von bis zu 500 MBq verarbeiten kann. SAFIR-II wurde als MRT (Magnetresonanztomographie)-kompatibler PET-Einsatz konzipiert, und kann Bilder zeitgleich mit einem Bruker BioSpec 70/30 MRT aufnehmen.

Es werden ebenfalls Messungen und Tests zu SAFIR-II's Leistungsvermögen vorgestellt. Es wird gezeigt, dass SAFIR-II Daten bei Messaktivitäten von bis zu 500 MBq verarbeiten kann, wobei weniger als 5% der Daten durch Totzeit-Effekte verloren gehen. Der Detektor misst Koinzidenzen mit einer Zeitauflösung von 221 ps und einer Energieauflösung von 12.1%. Es wurde durch die Messung einer Linien-Quelle gezeigt, dass diese Werte sowohl von der Messaktivität als auch vom zeitlichen Schwellenwert des verwendeten ASIC's (application-specific integrated circuit) abhängig sind. Die Sensitivität des Detektors wurde unter Verwendung des nach dem NEMA-NU4 Standard von 2008 spezifizierten Messprotokolls evaluiert. Das Sensitivitätsprofil wurde gemessen, und es wurde gezeigt, dass der Detektor unter Verwendung der ICSR (Inter-Crystal-Scatter-Recovery)-Technik eine maximale Sensitivität von 3.89% aufweist. Zusätzlich wurde die Kompatibilität zwischen SAFIR-II und dem MR-System mit Hilfe von Testsequenzen des MR-Herstellers demonstriert. Mehrere Messungen zur Bildqualität des Detektors wurden ebenfalls durchgeführt. Die Messung eines Derenzo-Phantoms hat gezeigt, dass SAFIR-II Strukturen mit einer räumlichen Auflösung von bis zu 1.7 mm darstellen kann. Eine Untersuchung der Bildqualität anhand des NEMA NU4 Standards ergab, dass der Detektor bei 500 MBq eine Uniformität von 2.96% aufweist, und somit mehrere kommerziell verfügbare Systeme übertrifft. Für die SOR (Spill-over-Ratio) der Luft- und Wasser-Kammern wurden zufriedenstellende Werte beobachtet, welche mit 0.057 und 0.075 vergleichbar mit den Werten kommerzieller Systeme sind. Die RC (Recovery Coefficient) Parameter der 5 mm, 4 mm, 3 mm, 2 mm, und 1 mm Röhren waren jeweils 0.92, 0.79, 0.58, 0.32, und 0.053. Der Einfluss verschiedener Bildkorrekturmethode auf die Bildqualität wurde ebenfalls untersucht, wobei die Normalisierung des Detektors einen signifikant höheren Effekt auf Bilder hatte als noch beim Vorgängersystem SAFIR-I. Eine Studie mit einer Sprague-Dawley Ratte wurde durchgeführt, der 283 MBq FDG ($[^{18}\text{F}]$ Fluorodeoxyglucose) injiziert wurden. Bilder dieser Studie, welche über einen Zeitraum von zehn Sekunden aufgenommen wurden, waren detailliert genug, um darauf kleine anatomische Strukturen, wie zum Beispiel den Herzmuskel der Ratte, zu erkennen.

Acknowledgements

Working on a project such as SAFIR, especially during a global pandemic, was an arduous task I would never have managed to complete on my own. Thus, I would like to thank the following people for their support, be it while working on this thesis or at any point before that.

Prof. Günther Dissertori, without whom I would never have had the chance to work on this project at all. And who always made time to listen in case of a problem despite being in one of the highest offices imaginable.

Dr. Werner Lusterma, who supervised the design and assembly for all SAFIR detectors. Thank you for your guidance throughout four and a half continuous years, listening to my problems on a weekly basis.

Dr. Christian Ritzer, who designed the original `petaAnalysis`, FTR, and DAQ-control software. He assisted me numerous times and provided a fountain of knowledge. Dr. Agnieszka Zagoździńska-Bochenek, who designed the SDIP3 and several other PCBs used in SAFIR-II. Robert Becker, who engineered the mechanical support structure and installed the cooling system. Volker Commichau, who provided the MR-compatible DC-DC converters and the SBTV3. Ulf Röser and Jan Sörensen, who provided major support in the electrical cabling and other electrical engineering tasks. SAFIR-II would never have been possible without them.

Prof. Peter Fischer and Dr. Michael Ritzert, who designed the PETA-8 ASIC. They provided me with continuous assistance even after I finished my master's degree, be it by answering every single one of my E-Mails concerning PETA-8, or being present for my defense.

Dr. Somayeh Saghamanesh, who provided the MR-Fusion images in Chapter 4. Dr. Pascal Bebie, who assisted in adapting the random correction to SAFIR-II.

Dr. Matthias Wyss, Dr. Afroditi Eleftheriou, Anita Siebert, and Prof. Bruno Weber. They participated in the in-vivo rat study, provided insight and assistance concerning SAFIR-II's MRI compatibility, and provided comments for some of the more medical parts of this thesis.

Jennifer Zollinger, Bettina Lareida, and Gabriela Amstutz, for providing assistance in any administrative matter I faced. They always found a way to make me laugh every

time I entered their office.

Gianluca Vagli, Dr. Lukas Gerritzen, Vasilije Perovic, Kaustuv Datta, Dr. Diego Sanz Becerra, and everyone else at the IPA. They were my colleagues and friends through this time, facing similar challenges as I did.

Daniel Coquelin, Dr. Oskar Taubert, and Frank Schlegel, for tons and tons of Malarkey.

Dr. Charlie Debus and Dr. Markus Götz, for making sure this document was free of Malarkey. For caring, for normal group behaviour, and for everything else.

Sebastian Geiser, Sarah Bartha, Ingmar Middelhoff, Gesa Berthold, Rebecca Pons, Arvid Weyrauch, Dr. Lars Heyen, and many others, for making the years studying physics that much more enjoyable.

Dr. Vitali Karasenko, friend and mentor, who taught me the foundation of many of the skills I needed for this work.

Stefan Wild, who started me down the path of physics in the first place.

Calvin Schmäbler and Patrik Meyer, for more than a decade of friendship.

And, of course, my family. My sister, Charlie Debus, who deserves to be mentioned twice. My brother, Nils Debus, for being my best friend. My father, Christian Debus, who continued working after retirement just so he could see his son more often. And my mother, Cornelia Debus, who provides a home that I always can and want to return to.

Contents

Abstract	I
Zusammenfassung	III
Acknowledgements	V
Table of Contents	VII
1 Introduction	1
1.1 Positron Emission Tomography: A Theoretical Overview	2
1.2 Technological Advancement in the Field of Positron Emission Tomography	4
1.3 The SAFIR Project	5
1.4 Outline	6
2 System Design, Components, and Functionality	7
2.1 The SAFIR-II High-Rate Positron Emission Tomography Insert	7
2.1.1 Design Requirements	7
2.1.2 System Overview	8
2.2 Photon Detection	11
2.2.1 Crystal Size and Spatial Resolution	11
2.2.2 Scintillator Material and Energy Resolution	12
2.2.3 Crystal Geometry and System Sensitivity	14
2.2.4 SiPMs and Overvoltage	17
2.3 PETA8: Digitization of Analog Signals	18
2.3.1 The Front End: Energy and Timing	18
2.3.2 Data Readout Architecture	20
2.3.3 Detector Dead Time and Data Loss	21
2.3.4 SP8: The Detector Head Module	27
2.4 Supporting Electronics	28
2.4.1 Digital Control: SDIP3	28
2.4.2 Power Conversion: SSPD	30
2.4.3 Bias Voltage Control: SBTV3	30
2.4.4 Fast Control Signals: SFCM & SFCD	31
2.4.5 Power Distribution: SPPD & SBD	32
2.5 Mechanical Support	34
2.5.1 Air Cooling and Insulation	34
2.5.2 Shielding and MRI Compatibility	36

3	System Control, Data Acquisition, and Raw Data Analysis	43
3.1	Control Software	43
3.1.1	Power Supply Control	44
3.1.2	Ethernet Control	45
3.1.3	Bias Control	48
3.2	Data Analysis and Calibration	49
3.2.1	Energy Filtering	49
3.2.2	Coincidence Sorting	53
3.2.3	Timing Calibration	54
3.2.4	Threshold of the PETA-8 Front End	58
3.2.5	Inter Crystal Scatter Recovery	61
4	From Coincidences to Images	65
4.1	PET Image Reconstruction	66
4.1.1	Projection Data	66
4.1.2	Filtered Backprojection	67
4.1.3	Iterative Reconstruction	69
4.1.4	Ordered Subset Processing	72
4.2	Image Quality	73
4.2.1	Derenzo Phantom Measurement	73
4.2.2	Image Quality Phantom Measurement	78
4.3	Correcting for Artefacts	81
4.3.1	Attenuation Effects	81
4.3.2	Detector Normalization	82
4.3.3	Random & Scattered Coincidences	86
4.3.4	Results	89
4.4	SAFIR-II In-vivo: Zeus the Rat	95
5	Conclusion	98
5.1	SAFIR-II by Comparison	99
5.2	Outlook	101
	Glossary	103
	Bibliography	106

1 Introduction

At some point, everything's going to go south on you. You're going to say, 'This is it. This is how I end.' Now, you can either accept that, or you can get to work. You solve one problem, and then you solve the next problem, and then the next. And if you solve enough problems, you get to go home.

(Mark Watney)

One of the persistent cornerstones of modern science has been the need for ever-more advanced technology. While Albert Einstein's prediction of gravitational waves in 1916 was spot on, direct physical evidence of this prediction was only producible a century later, following the technological development of sufficiently sensitive laser interferometry [1]. The 7 TeV center-of-mass energy of the Large Hadron Collider's (LHC) first run enabled the discovery of the Higgs-Boson [2] but has since been surpassed in subsequent runs. Plans for an even higher-luminosity upgrade of the LHC are well underway to search for new frontiers in high-energy physics [3]. Since the launch of the first space telescope in 1968 [4], astronomers have attempted to peer further and further into the cosmos by pushing for ever-increasing resolution and sensitivity through devices such as the Hubble [5] and James-Webb [6] telescopes. Present in each of these instances were questions about the inner workings of our universe, resulting in the development of new technologies that provided answers beyond their initial purpose.

The field of medical technologies is no stranger to such advancements. Magnetic resonance imaging (MRI) has been able to peer into the human body at ever-increasing spatial resolutions through ever-increasing magnetic field strengths [7], while steady advances in neuroscience and robotics have enabled increasingly sophisticated prosthetics [8]. Artificial neural networks have been used in many new algorithms for disease detection and treatment [9], aided by the high computational performance provided by new silicon technologies.

While on a much smaller scale, this thesis, as well as the collaborative project that spawned it, answers the need for such technological advancement. Most positron emission tomography (PET) scanners center their design around imaging with maximal image quality and minimal radiation dosage, resulting in a slow but precise acquisition. This makes them incompatible with research that desires to investigate fast kinetic processes, biological processes that are completed within comparatively short periods of time. The attempt to design an MRI-compatible PET scanner capable of observing such processes is the foundation of the Small Animal Fast Insert for MRI (SAFIR) project.

1.1 Positron Emission Tomography: A Theoretical Overview

PET is based on the localization of radioactive β^+ emitters via the detection of photons created by the resulting positron-electron annihilations. Preceding each scan is the production of a radiotracer, a chemical compound created through the combination ('labeling') of a select metabolically active substance with a β^+ emitting radioisotope. Upon introducing the tracer into a living organism, it accumulates in regions of interest according to its metabolic behavior. Which compound is used is highly dependent on the focus of the scan, as the specific metabolic behavior of different substances can be used to investigate various things. As an example, the most commonly used tracer is [^{18}F]Fluorodeoxyglucose (FDG), a glucose analog labeled with Fluorine-18 (^{18}F). In metabolic processes it initially operates as regular glucose, being transported into cells and phosphorylated in the same way as glucose. However, it does not undergo further metabolism afterwards, causing it to remain trapped in cells. As it accumulates in regions with high glucose uptake, such as the brain or cancerous tissue, FDG-based PET scans are widely used as a tool for oncology and neuroscience. To provide a second relevant example, Oxygen-15 labeled water is metabolically inert, causing it to distribute freely via the bloodstream once injected. Thus, it is useful for measuring and quantifying blood flow in the brain, heart, or elsewhere [10].

While the metabolic properties of a radiotracer give rise to a specific tracer distribution, the decay products of the radioisotope enable the observation of this distribution. A β^+ decay results in the emission of a positron from the isotope's nucleus, following the decay equation



Although it possesses momentum, the positron travels only a short distance within tissue before annihilating with a surrounding electron. This distance varies depending on the utilized isotope and surrounding material but is usually less than 2 mm for most PET tracers in organic tissue. The result of the annihilation process is the basis of PET: two photons, each possessing an energy of 511 keV, being emitted in opposite directions of each other. By detecting both of these photons, it can be inferred that the positron-electron annihilation is located along a line between the two detection points. An image of a radiotracer's distribution can then be created by detecting a large number of such lines at various angles. Figure 1.1 visualizes the basic concepts of a PET scan.

Most PET scanners rely on scintillating crystals and photosensors to detect the emitted 511 keV photons. The crystals are responsible for absorbing the high-energy photons and converting them into visible light, after which the photosensors are utilized to generate an analog electric signal based on the amount of light detected. Early PET scanners utilized photomultiplier tubes (PMTs) as photosensors, but most modern designs have exchanged them for silicon photomultipliers (SiPMs) due to their lower cost, smaller size, and, in the case of PET-MRI scanners, necessary immunity to magnetic fields. Following the path of detected photons, properties of the analog signals generated by the photosensor (such as the arrival time and amplitude) are recorded, either via application-specific inte-

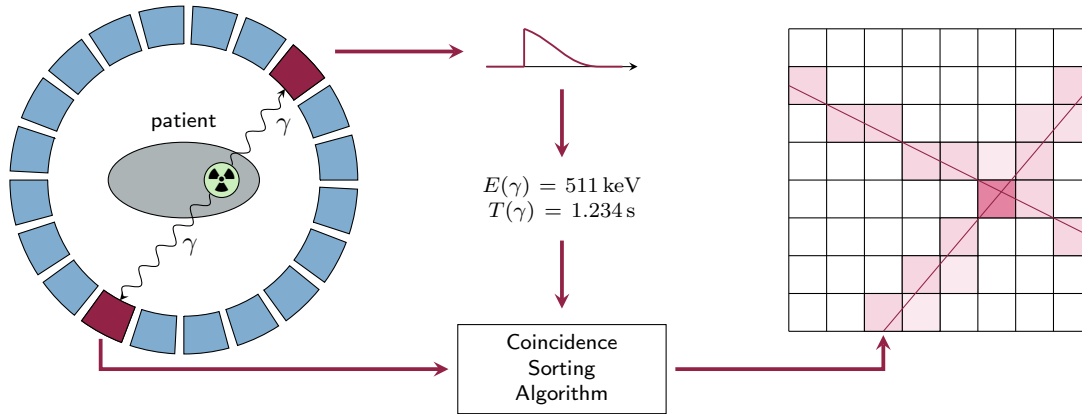


Figure 1.1: Basic functional principle behind PET image acquisition. **Left:** A β^+ emitting radiotracer results in two 511 keV photons being emitted at an angle of 180° , which interact with surrounding scintillator crystals. **Middle:** The scintillation light from the crystals creates a signal in the connected photo-sensor, which is recorded and matched with other signals in post processing to find coincidences. **Right:** Detected coincidences create LORs, which are used by a reconstruction algorithm to create a voxelized distribution of the tracer.

grated circuits (ASICs) featuring time-to-digital converters (TDCs) and analog-to-digital converters (ADCs), or similar methods. As such, the PET scanner now has information on a detected photon's energy, arrival time, and location. This information is then processed, either online during Data Acquisition (DAQ) or offline after a completed scan. If a detected photon, also called a hit, does not exhibit an energy matching the expected 511 keV, it might have been involved in a scattering process before detection, causing it to deviate from its original path. Another possibility is that the sensor was triggered by something other than an annihilation photon, such as the intrinsic radiation present for some crystal materials. Whatever the reason, those hits are commonly rejected and not processed further. Furthermore, as they are created in the same annihilation process, it makes sense that any two photons of interest should be detected more or less simultaneously, barring some measurement inaccuracies as well as time differences resulting from the photons' travel time between their creation and detection. Thus, a PET scanner searches for pairs of coincident hits while rejecting singular hits without a partner or groups of more than two. Considering the detected locations of such coincidences, as previously mentioned, a line can be drawn between them. After recording a large number of such lines of response (LORs), a reconstruction algorithm can then be utilized to infer an image of the tracer distribution.

1.2 Technological Advancement in the Field of Positron Emission Tomography

Since its inception in the 1950s and the construction of the first commercial scanner, ECAT II, in 1978 [11], PET has proven to be an invaluable tool in medical diagnostics, particularly in oncology. Though the necessity of an on-site cyclotron and the initially low photon detection sensitivity of early scanners hindered its widespread use, its usefulness in detecting small lesions for tumor diagnosis and staging was recognized early on. Since then, many steps have been taken to improve PET instrumentation as a whole. The initially low sensitivity and difficult-to-manufacture NaI(Tl) scintillation crystals were replaced first by the superior bismuth germanium oxide (BGO) and later on by lutetium oxyorthosilicate (LSO) and lutetium-yttrium oxyorthosilicate (LYSO) crystals. These crystals were utilized in larger numbers as well as smaller sizes to feature higher spatial resolutions, sensitivities and larger fields of view (FOVs). Recently, the EXPLORER scanner [12] developed at UC Davis has shown the capabilities and usefulness of total-body PET scanners. Aiding this push for higher sensitivity was the improved timing precision exhibited by LSO crystals, which enabled measuring the difference in arrival time for annihilation photons. This provided additional positioning information for the decay, allowing the use of time-of-flight (TOF) based techniques to improve a scanner's signal-to-noise ratio. Nowadays, most clinical PET scanners have integrated TOF information within their reconstruction algorithm with timing precisions of down to 214 ps [13, 14].

All of these technological advancements have ultimately helped to improve image quality for both clinical and preclinical applications. In tandem with these developments, efforts have been made to address a second issue with PET: the lack of detailed anatomical information. In contrast to the widely utilized modalities of computed tomography (CT) and MRI, which are able to provide detailed three-dimensional information of a patient's morphology, PET's nature as a functional imaging technique limits it to only ever provide information on the injected tracer's distribution and concentration. An obvious solution to this problem consists of combining PET with one of these other modalities. For PET-CT, this has the added benefit of providing the PET scanner with an up-to-date attenuation map to utilize in its image reconstruction. However, both PET-CT and PET-MRI face additional technical challenges. In the case of PET-CT, any X-ray created by the CT can obviously trigger a PET scanner's scintillation crystals, making it impossible to operate both modalities simultaneously and in the same space. It is possible to place a CT scanner and a PET scanner back-to-back in a single device instead and acquire data by slowly moving the patient through the scanner. A prototype employing this technique became operational at the University of Pittsburg in 1998 [15]. Combined PET-CT has been commercially available since 2001 with the "Discovery LS" and has since found widespread use in hospitals worldwide, to the point that PET-CT has replaced regular PET in most medical centers around the world. Unfortunately, a patient undergoing both a PET and a CT scan is exposed to the radiation dosage from both, which is why the alternative combination of PET and MRI was a subject of interest even

before the success of PET-CT. MRI has the added benefit of an increased soft-tissue contrast compared to CT, but due to the high magnetic field strengths employed, any PMTs would cease to function in proximity to an MRI. Operating two separate devices and moving a patient between the MRI and PET scan is the most straightforward solution to this and has been implemented in scanners such as the Philips Ingenuity PET-MRI [16]. This sequential imaging procedure is not without drawbacks, as it can lead to inaccuracies between the images due to patient movement [17]. Thus, 'true' PET-MRI was only possible following another technological leap. SiPMs capable of operating within an MRI's field would replace PMTs, allowing for the simultaneous acquisition of PET and MRI in 2010 with the Siemens Biograph mMR [18].

Nevertheless, as with any technology, a number of challenges and technical limitations remain, two of which are of particular interest for this thesis. Firstly, the increasing demand for diagnosis and staging in oncology using tracers such as ^{18}F has been restricted by the required processing time of about 30 minutes per patient [13]. Reducing the time required to perform a scan would be a considerable benefit, but this collides with a second limitation. The push for higher system sensitivity has so far manifested in increasing image quality, but a high sensitivity could naturally also be used to decrease scan time while maintaining the existing image quality due to higher count rates. However, higher count rates result in a higher level of noise through so-called "random" coincidences as well as decreased signal due to increased detector dead times. These effects ultimately lower the scanner's signal-to-noise ratio (SNR) and, thus, its effective sensitivity [13]. Reducing these effects, be it via improved data collection or timing precision, is of high interest not only for clinical applications.

1.3 The SAFIR Project

In research, PET possesses a vast potential for various studies beyond the field of oncology, facilitated by the availability of numerous radiotracers with alternative uses. This is true for clinical studies but arguably even more so for preclinical investigations. Nowadays, several different preclinical PET, PET-CT, and PET-MRI scanners are available, both for purchase from companies as well as developed by academic research groups. The field of preclinical PET instrumentation as a whole focuses on many of the same topics as in clinical scanners: increased system sensitivity, larger scannable volumes, and improved image quality. SAFIR detector II (SAFIR-II), the focus of this thesis, was designed, developed, and constructed with a different goal in mind.

The long image acquisition times commonly found in PET often force researchers to work with singular, static images of a tracer's distribution. While this is sufficient in many cases, the examination of certain biological processes would benefit from acquiring several images within much shorter time frames, opening up the possibility of studying a tracer's distribution over time. For example, one of the main topics of interest for the SAFIR collaboration is an investigation into the Oxygen-Glucose-Index (OGI). The OGI denotes a useful measure when studying brain metabolism, as it details a cell's preferred process for adenosine triphosphate (ATP) production. ATP, being the primary compound that

1 Introduction

provides energy to cell processes, can be generated through glycolysis, a pathway in which glucose ($C_6H_{12}O_6$) is converted into pyruvate. During this process, which does not require oxygen, ATP is first consumed but later produced in higher quantities. Following this, pyruvate might be converted into lactate if not enough oxygen is available or undergo oxidative phosphorylation (OxPhos), which requires oxygen to generate additional ATP at a significantly higher yield compared to glycolysis. If one observes the rate at which glucose and oxygen are metabolized in relation to each other in a region of interest, one can use this information to infer the relative rate between glycolysis and OxPhos. The OGI measures precisely this, and a study conducted in 1988 has shown a link between the OGI and the brain's activation status by observing the change in metabolic rate for glucose and oxygen in humans after providing visual stimulation [19]. This study observed two groups of five volunteers each, measuring the metabolic rate for glucose in one group and for oxygen in the other using PET imaging and comparing these measurements. Reproducing this investigation while observing both the metabolic rate for glucose and oxygen in the same subjects, as well as their temporal distribution, is of considerable interest. To achieve this, one would need a scanner capable of acquiring a series of PET images with a scan time of only a few seconds per image. This scanner would also have to provide accurate quantitative information on the activity concentration in a region of interest determined via anatomical information from an MRI scan. These requirements form the basis for SAFIR-II.

1.4 Outline

In this thesis, I present the design, functionality, and performance of the SAFIR-II PET insert. Chapter 2 discusses SAFIR-II's mechanical and electrical components, along with the reasoning behind specific design choices. I have evaluated the functionality of several of these components and the scanner as a whole, as well as performed measurements investigating the scanner's behaviour when detecting photons. The software programs used to control the scanner and analyze its raw data, which I adapted and tested for SAFIR-II, are detailed in chapter 3. In chapter 4, an assessment of the scanner's image reconstruction and quality is presented. I adapted various image correction methods for SAFIR-II and investigated their influence on the scanner's image quality. Also presented are the results of a first in-vivo rat measurement performed using SAFIR-II. Lastly, chapter 5 serves as a concluding statement, where I compare SAFIR-II with other PET scanners and discuss potential future studies.

2 System Design, Components, and Functionality

Before detailing the intricacies of each component of SAFIR-II, I would like to provide a brief overview of the entire system, as well as the requirements for its performance outlined before it was built. This shall serve as a guideline for explanations going forward.

2.1 The SAFIR-II High-Rate Positron Emission Tomography Insert

2.1.1 Design Requirements

There is a fundamental trade-off in PET between the three parameters of image quality, injected radiation dose, and scan time. Given a scanner's inherent detection efficiency and a fixed selected radiotracer dose, a corresponding scan time is required to reach the desired image quality. An image's quality is based on the size of its voxels and the signal-to-noise ratio (SNR) of the voxel counts. Thus, as radioactive decay is stochastic in nature, a minimum amount of counts per voxel is required to reach the desired image quality. Shortened scan times will inherently decrease image quality, which is usually undesirable as it decreases the ability to, for example, detect small lesions in oncological applications. In clinical PET, it is logical that one of the most desired properties is a scan with a minimal injected dose, which still maintains sufficient image quality. As a result, minimizing the scan time is seen as a secondary concern, with scan times of the order of 10 minutes (and sometimes more) being considered acceptable in terms of patient comfort and economic feasibility. Research into clinical PET technology is oftentimes focused on either increasing image quality at a given scan time and dose, maintaining a given image quality at a reduced dose, or a combination of the two.

However, as previously mentioned (section 1.3), not all research is achievable with scan times of several minutes. While cancerous tissue can be seen as relatively stationary, various biological processes take place on much shorter timescales. To name an example beyond just the OGI, the measurement of cerebral blood flow (CBF) involves arterial input functions that reach their peak within the first 60s and thus require measurements with timescales at least an order of magnitude smaller to be adequately resolved in time. To maintain a usable image quality while repeatedly acquiring images with very short acquisition times, one would need to increase the injected tracer dose massively. However, a scanner capable of acquiring images at such high doses is not commercially available yet.

2 System Design, Components, and Functionality

SAFIR-II, the scanner described in this thesis, was designed to make these measurements possible. As a scanner that is to be utilized in a preclinical setting for the imaging of small animals like mice and rats, it needs to be able to resolve small anatomical structures such as a mouse’s myocardium. Therefore, a minimal spatial resolution of 2 mm is required, which would be achievable using a voxel size of $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ or smaller. An image acquisition time of 5 s was desired, and to reach the count rate statistics needed for such short acquisition times, a total injected activity of 500 MBq was targeted [20].

Simultaneous acquisition of an MRI image was an additional requirement to correctly assess anatomical regions of interest in the acquired PET image. SAFIR-II was designed to be insertable and operable within a Bruker BioSpec 70/30, a preclinical 7T MRI system [21]. Therefore, it needs to be fully compatible with the MRI system, not influencing its image acquisition while ensuring that the MRI does not influence SAFIR-II either. The operation within the MRI also dictates the scanner’s inner and outer diameters, limited by the specific dimensions of the Bruker BioSpec 70/30 MRI and its RF coils to 114 mm and 199 mm, respectively.

2.1.2 System Overview

Having established key parameters and requirements for the desired scanner, the SAFIR collaboration began work on a prototype, the design, evaluation, and performance of which have been detailed extensively [22, 23, 24]. This Dual Ring Prototype (DRP) was completed in 2019 and featured a 35.8 cm long FOV, four times shorter than what was planned for the final scanner. As the DRP exhibited excellent performance, a decision was made to upgrade it by increasing its FOV to 54.2 cm, making it useful in preliminary studies until the final scanner was completed. The resulting scanner was named SAFIR detector I (SAFIR-I), and alongside its completion, a final scanner was developed under the title SAFIR-II. The design of SAFIR-II emulates that of its predecessor, though adjustments and variations exist in several key factors. A rendering of it is shown in figure 2.1, with a schematic detailing the scanner’s components and their connections being illustrated in figure 2.2. The following section presents a general overview of the design of SAFIR-II, with detailed explanations of the individual components being given later on.

SAFIR-II operates almost entirely within the bore of a Bruker BioSpec 70/30 MRI system. The only exceptions to this are an external DAQ-PC for data processing and analysis, power supplies delivering the main operating voltages, and an air cooling unit. These are connected to the scanner via glass fiber Ethernet links, twisted pair cabling, and air tubing, respectively, while being located in a separate room outside of the MRI’s Faraday cage. The scanner itself is housed in a mechanical support structure made from carbon fiber that positions the electronic components within the MRI such that the PET and MRI FOVs overlap.

SAFIR-II’s sensitive area is comprised of 11 520 LYSO crystals, which are arranged into blocks of 15×8 crystals each. The crystals of each block are grouped into one 8×8 and one 7×8 crystal matrix and kept separate within a matrix by using a sandwiched layer of enhanced specular reflector (ESR) foil and aluminum [26]. Each matrix is glued

2.1 The SAFIR-II High-Rate Positron Emission Tomography Insert




Figure 2.1: Three-dimensional rendered cross-section of the SAFIR-II insert. The rendering shows LYSO crystals coupled to SiPM arrays forming the detector head, the combined SDIP3, SBTV3 and SSPD printed circuit boards (PCBs), and the outer SPPD, SBD and SFCM ring boards. Visualized in gray is the carbon fiber mechanical support structure. (Rendering by R. Becker)

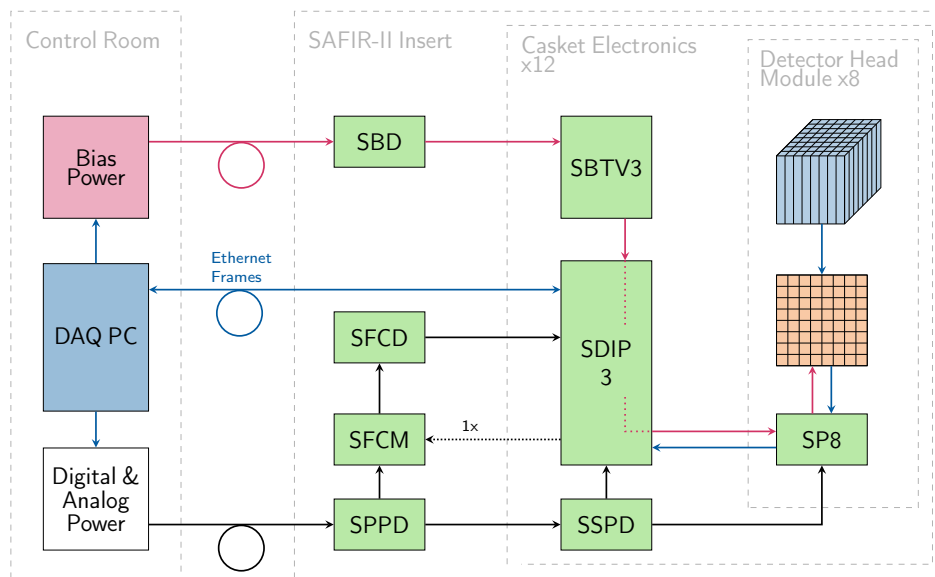


Figure 2.2: Schematic overview of all components and connections within the SAFIR-II PET insert. (Adapted from [25])

2 System Design, Components, and Functionality

to a Hamamatsu SiPM array, both of which are mounted on a single SP8 PCB featuring four PETA8 ASICs responsible for the acquisition of the SiPM analog signals. One SP8 PCB with two connected SiPM arrays and the 120 LYSO crystals glued to them form a combination called a detector head module. Twelve modules each are arranged to form a dodecagonal ring, with eight such module rings together covering an axial FOV of 144 mm.

Due to its dodecagonal shape, the SAFIR-II insert is split into twelve sectors, with each sector containing eight detector head modules. Also contained within each sector is a SAFIR Digital Interface for PETA Version 3 (SDIP3) PCB, on which the detector head modules are mounted. The SDIP3 configures and reads out the PETA-8 ASICs using a Xilinx Kintex-7 field-programmable gate array (FPGA), and features a 10GBit SFP Ethernet module to transfer data to the external DAQ-PC. It also serves as the main connection between the detector head modules and other supporting PCBs. This includes two further PCBs per sector, one being the SAFIR Bias and Temperature Version 3 (SBTV3), which is used to individually configure the bias voltage of each SiPM array, and one being the SAFIR Secondary Power Distribution (SSPD), which features power converters required for the operation of both the PETA-8 ASICs and the SDIP3.

Multiple ring boards located outside the twelve sectors are required for the distribution of electrical power and common control signals. Two SAFIR Primary Power Distribution (SPPD) boards deliver the main supply voltage of 18 V to six sectors each, while two SAFIR Bias Distribution (SBD) PCBs do the same for the bias and support voltages required by the SBTV3 and SiPM arrays. Fast control signals used to synchronize the 12 sectors are passed from a single SAFIR Fast Control Master (SFCM) to two SAFIR Fast Control Distribution (SFCD) PCBs, which in turn send these signals to all 12 SDIP3 simultaneously. The SFCM itself is connected to one of the 12 SDIP3 boards, which is used to signal the SFCM following the reception of a respective Ethernet package from the external DAQ-PC.

Several of the PCBs used for SAFIR-II are adapted from those used in SAFIR-I. SAFIR-I covered a shorter FOV by featuring only three rings of detector modules and processed its SiPM signals using the PETA6SE, an older version of the PETA-8 ASIC. The SDIP3, SP8, SSPD, and SBTV3 were all designed to accommodate these changes, and the testing of their functionality was a major part of my work. The same is true for the 10GBit SFP Ethernet modules, as SAFIR-I featured 1GBit modules instead. Unchanged from the DRP are the ring boards used for power and fast signal distribution.

2.2 Photon Detection

The previously detailed requirements for SAFIR-II's proportions and spatial resolution are the main determining factors when it comes to the size and arrangement of the scintillating crystals. As they are used to detect the photons of each annihilation process, the crystals confine each LOR by specifying the incident photons' locations. At the same time, a crystal's material dictates its attenuation length and cross-section, which is of high importance for the scanner's photon detection efficiency. The material furthermore influences both the energy resolution and the timing resolution of the scanner through factors such as its light yield and average interaction time.

2.2.1 Crystal Size and Spatial Resolution

The spatial resolution of a PET scanner is dictated by several factors apart from just the size of the detecting crystals. As positrons emitted from the tracer possess kinetic energy, they can travel a short distance before annihilating. The mean distance of this effect is known as the positron range, and it depends on the emitting isotope and surrounding material. Furthermore, the electron-positron pair is not necessarily at rest within the laboratory frame at annihilation, resulting in the annihilation photons being emitted at an angle slightly deviating from 180° . The degradation in spatial resolution of this acollinearity depends on the length of the LORs and, therefore, the scanner's radius. Additional factors include a decoding error dependent on the coupling of crystals to SiPMs, the position of the isotope within the FOV, and a sampling error owing to the non-uniform distribution of LORs in the FOV. According to Moses et al. [27], a theoretical lower limit on the spatial resolution Γ of a scanner imaging a point source (a β^+ emitter of minimal size) can be calculated as

$$\Gamma = 1.25 \sqrt{(d/2)^2 + s^2 + (0.0044R)^2 + b^2 + \frac{(12.5r)^2}{r^2 + R^2}} \text{ (mm FWHM)}. \quad (2.1)$$

Here, d denotes the crystal width, s the positron range, R the detector ring radius, b the decoding error, and r the radial distance of a source from the center. The factor of 1.25 is an empirically observed value for the sampling error, while $0.0044R$ accounts for the acollinearity. For a source placed in the center of the FOV and a decoding error $b = 0$ reached by a one-to-one coupling of crystal to SiPM, formula 2.1 simplifies to

$$\Gamma = 1.25 \sqrt{(d/2)^2 + s^2 + (0.0044R)^2} \text{ (mm FWHM)}. \quad (2.2)$$

SAFIR-II will be used in imaging studies using both ^{18}F and ^{15}O , featuring positron ranges of 0.1 mm and 0.5 mm, respectively. The inner diameter of the scanner results in a minimum radius R of 57 mm. Therefore, to reach a spatial resolution of <2 mm in the center, a maximum crystal size of ≈ 3 mm would be required.

For imaging, we would desire the <2 mm resolution not just in the center but in an extended area. Additionally, the formula given above assumes a cylindrical detector ring, while SAFIR-II's design results in a dodecagonal arrangement. The crystal distance is also extended due to the mechanical support structure employed, resulting in an extended

Scintillator Material	NaI:Ti	BGO	LSO	LYSO
Chemical Composition	NaI:Ti	Bi ₃ Ge ₄ O ₁₂	Lu ₂ SiO ₅ :Ce	(Lu–Y) ₂ SiO ₅ :Ce
Density [g cm ⁻³]	3.67	7.1	7.4	7.4-4.53
Attenuation length	29.1	10.4	11.4	11.4-25.8
light yield [ph/keV]	38	6	29	29-46
Decay constant [ns]	230	300	40	40-70
Characteristic Wavelength [nm]	410	480	420	420
Energy Resolution [%]	6.6	10.2	10.0	10.0-10.2

Table 2.1: Properties relevant to photon detection for some of the most common PET scintillators, taken from [29]. The attenuation length is given for 511 keV photons, specifically. The properties of LYSO depend on the ratio between Lutetium and Yttrium, which can vary between manufacturers.

radius of 64 mm. For these reasons, SAFIR-II’s crystals feature a width of 2 mm instead, resulting in a theoretical lower limit for the spatial resolution of 1.44 mm at the center.

2.2.2 Scintillator Material and Energy Resolution

Over the years, several potential scintillation materials have been identified as particularly useful in PET. Relevant properties of some of the most common ones are summarized in table 2.1. For SAFIR-II, the chosen material should preferably feature a short attenuation length to boost the scanner’s photon detection efficiency. This is necessary due to the length of the crystals, which is limited to 13 mm via two requirements. For one, longer crystals can lead to a higher parallax error, degrading the scanner’s spatial resolution. On the other hand, limited space is available due to the inner and outer diameter of the scanner, as well as the space needed for supporting electronics and mechanics. On top of that, the high expected rate of incoming photons requires the material’s decay time to be as short as possible to allow signals from individual hits to be kept separate. Additionally, this improves the timing resolution of the scanner. The material chosen for SAFIR-II is LYSO provided by Meishan Boya Advanced Materials [28]. LYSO is an inorganic chemical compound, with the ratio of Lutetium and Yttrium varying between manufacturers. A diagram of the cross-section for LYSO with 5% Yttrium is shown in figure 2.3. In the case of the crystals employed in SAFIR-II the percentage of Yttrium is approximately 6%.

At an energy of 511 keV, the photons created from an electron-positron annihilation are subject to three potential interaction processes: Compton scattering, Rayleigh scattering, and photoelectric absorption. These processes take place both within the studied subject as well as the scintillating material, with Compton scattering being the predominant process in both cases, though to varying degrees. Photoelectric absorption increases in significance at lower energies, while Rayleigh scattering is nearly negligible. When occurring within the body, these processes result in attenuation and scattering effects, which will be discussed in more detail in section 4.3.

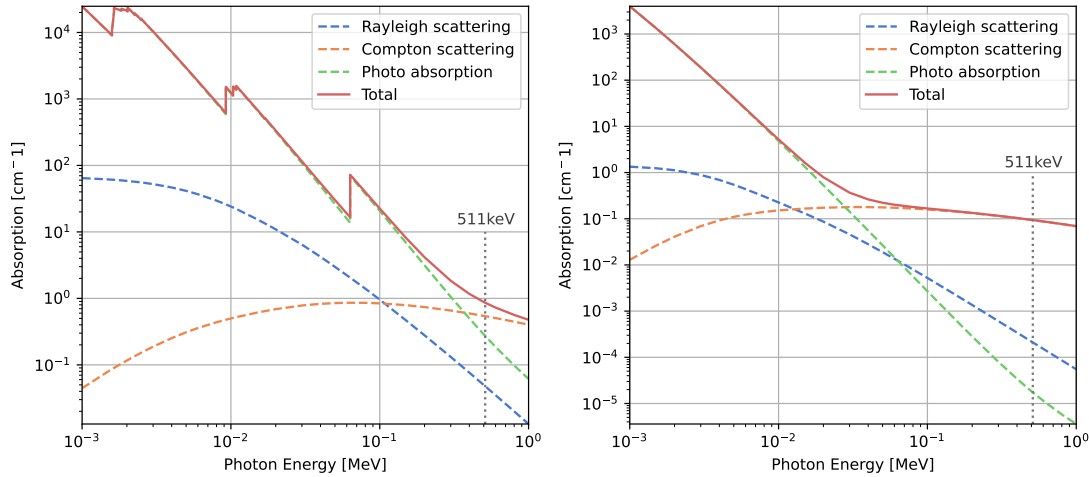


Figure 2.3: **Left:** Plot of the cross-section for LYSO with 5% Yttrium. **Right:** Cross-section for H_2O , which can be viewed as a good approximation for soft tissue. (Adapted from [22], data taken from [30]).

When occurring within a crystal, these interactions result in a measurable energy deposition. Figure 2.4 displays a typical energy spectrum observed by SAFIR-II during a measurement using ^{18}F . A PET scanner’s energy resolution typically refers to the full width at half maximum (FWHM) of the peak around 511 keV, which corresponds to the photopeak, i.e., an annihilation photon depositing its entire energy in the same crystal. This can be the case either via complete photoabsorption or a combined interaction including Compton scattering. In SAFIR-II’s case, we observed an average energy resolution of 12.1% across the entire scanner.

The energy deposited by photons undergoing Compton scattering depends on the scattering angle θ following the Klein-Nishina equation [31]

$$E'_\gamma(\theta) = \frac{E_\gamma}{1 + \frac{E_\gamma}{m_e c^2}(1 - \cos(\theta))}, \quad (2.3)$$

where E_γ, E'_γ denote the photon energy before and after scattering, m_e is the electron mass, and c is the speed of light. The maximum possible deposition of energy occurs at $\theta = 180^\circ$, resulting in $E_\gamma - E'_\gamma = 341$ keV. This matches well with figure 2.4, where we can observe the edge of the Compton spectrum at ≈ 340 keV.

Two additional effects influencing the measured energy spectrum should be mentioned. One is the presence of the isotope ^{176}Lu within the utilized LYSO crystals, resulting in the intrinsic emission of γ -radiation at energies of 202 keV and 307 keV [32]. The other is the k_α escape peak, caused by an incident photon freeing an inner electron in the LYSO crystal. For Lutetium, which has a binding energy of 63.3 keV in the inner shell, the result are peaks both at this energy and at $511 \text{ keV} - 63.3 \text{ keV} = 450 \text{ keV}$. Due to the energy resolution of 12.1%, the 450 keV peak is superimposed with the 511 keV peak.




Figure 2.4: Typical energy spectrum observed by SAFIR-II during a measurement using 6 MBq ^{18}F , averaged across all channels. Clearly visible is the photopeak at 511 keV (a) as well as the Compton edge at 340 keV (b). At low energies (c) several effects overlap, such as the spectrum dropping to 0 due to an energy cutoff intrinsic to each channel (see section 2.3) as well as the intrinsic radiation spectrum of Lutetium superimposing with the Compton spectrum.

2.2.3 Crystal Geometry and System Sensitivity

The material, size, and arrangement of the scintillation crystals are pivotal factors not only for the spatial and energy resolution of a scanner but also when it comes to the efficiency at which it is able to detect photons. Given a scintillation crystal of length d and material-dependent attenuation length λ , the probability of a photon interacting with the crystal can be calculated using the Beer-Lambert law

$$P_S = 1 - e^{-\frac{d}{\lambda}}. \quad (2.4)$$

To detect a coincidence, both photons need to interact with crystals opposite each other. Therefore, the probability of detecting a coincidence is P_S^2 . Given the 13 mm length of SAFIR-II's crystals and a LYSO attenuation length of 11.4 mm for 511 keV, this would yield a probability of 46.3%. However, the arrangement of crystals around the source is crucial, as it dictates the solid angle coverage of the scanner.

The 11 520 LYSO crystals of SAFIR-II are each $2.0 \times 2.0 \times 13 \text{ mm}^3$ in size. As mentioned in section 2.1.2, they are glued to SiPM arrays and arranged into matrices of 7×8 and 8×8 . Each SiPM array features 8×8 pixels, with the crystal matrices being arranged

such that each crystal covers one pixel. The crystals of a matrix are separated both mechanically and optically by a sandwich layer of aluminum and ESR foil, which is required to prevent light generated in one crystal from crossing over into another. The separation has a thickness of 0.2 mm and thus results in the crystals of each matrix being arranged at a pitch of 2.2 mm, a value identical to the pitch of the SiPM arrays they are attached to. Additionally, a gap of 0.5 mm width exists between each crystal matrix. The matrices are arranged as dodecagonal rings, with the surfaces of opposing crystals being separated at a distance of 128 mm. Due to this arrangement, additional triangular gaps between the crystal matrices exist in the corners of the dodecagon. Lastly, the gap between each ring is 0.6 mm wide, which results in a total axial length of 145 mm.

The arrangement of SAFIR-II's crystals results in a dodecagonal prism with several gaps in the scintillating material, which is inherently difficult to model. However, for the purposes of a rough theoretical estimate, we can assume SAFIR-II as a cylinder of length $L = 144$ mm, inner radius $R = 64$ mm, and thickness $d = 13$ mm. Depending on the incident angle, the photon will thus interact with more material, resulting in equation 2.4 changing to

$$P_S(\theta) = 1 - e^{-\frac{d}{\lambda \cos(\theta)}}. \quad (2.5)$$

Taking the solid angle into account, the probability of a photon being emitted at an angle such that it interacts with a crystal, with its decay partner being detected as well, can be calculated as

$$P_C = \frac{1}{\pi} \int_{-\theta_{\text{Max}}}^{\theta_{\text{Max}}} (P_S(\theta))^2 \cos(\theta) d\theta. \quad (2.6)$$

with $\theta_{\text{Max}} = \tan^{-1}(\frac{L/2}{R}) = 48.4^\circ$ being the maximum polar angle a photon can be emitted at to still hit a crystal.

Lastly, two multiplicative factors have to be taken into account. The individual crystals are 2 mm wide while being arranged at a pitch of 2.2 mm, resulting in a 0.2 mm gap. We can account for this by correcting P_C by a factor of 2 mm/2.2 mm. A second, more significant factor is connected to the LYSO cross-section and its Yttrium content. Naturally, the energy deposited within a scintillation crystal depends on the type of interaction process to which the photons are exposed. A photon undergoing Compton scattering might deposit just a portion of its energy before leaving the crystal, thereby becoming ineligible for the coincidence search. However, a Compton-scattered photon might also undergo a second interaction afterwards, with the LYSO cross-section for photoelectric absorption being drastically higher at lower energies. In such cases, the photon might still deposit its remaining energy in the same crystal, resulting in a valid 511 keV hit. The percentage of photons depositing their energy fully within one crystal is known as the photofraction and has been measured for LYSO at various values of up to 33% [33]. Integrating it into our equation, we arrive at

$$P_C = \frac{2}{2.2} \frac{1}{\pi} \int_{-\theta_{\text{Max}}}^{\theta_{\text{Max}}} (0.33 * (1 - e^{-\frac{d}{\lambda \cos(\theta)}}))^2 \cos(\theta) d\theta. \quad (2.7)$$

which leaves us with a rough theoretical estimate of 2.46%.

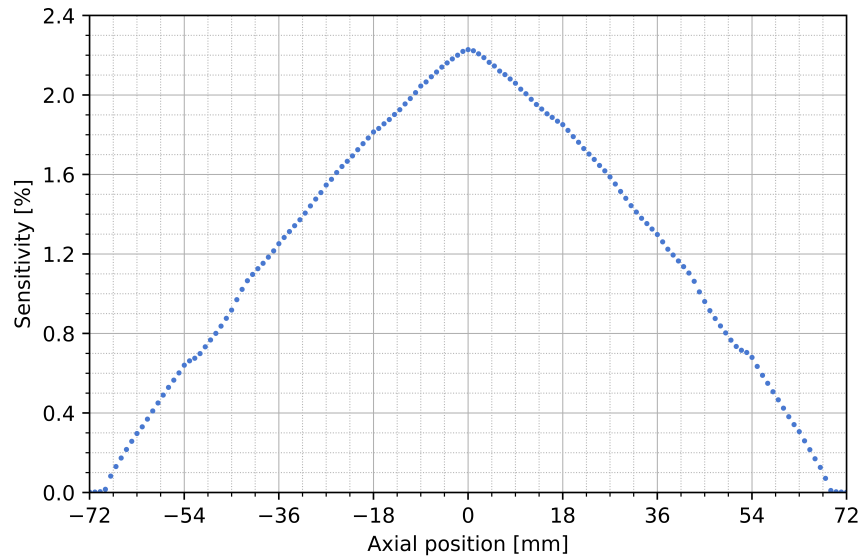


Figure 2.5: Sensitivity profile of the SAFIR-II PET insert, measured according to the NEMA NU-4 standard. A peak sensitivity of 2.23 % can be observed, with a mean system sensitivity of 1.25 %.

Measuring the System Sensitivity

Various measurements can be performed to gauge a PET scanner’s performance. To create a set of comparable measurements between scanners, the National Electrical Manufacturers Association (NEMA) described the 2008 NU-4 standard on ‘performance measurements for small animal positron emission tomographs’ [34], which includes a measurement for the system sensitivity. There, the sensitivity is defined as the percentage of a source’s decay events that are correctly detected by a scanner as coincidences. To evaluate it, data from a point source placed in the radial center of the scanner is acquired. The NU-4 standard specifies an evaluation of the sensitivity as a profile along the central axis by moving the point source in steps of at most half the crystal pitch and acquiring at least 10^4 coincidences at each position.

To evaluate the true detection efficiency of the SAFIR-II insert, I performed this evaluation using a 2.8 MBq ^{22}Na point source. The source consists of a 0.25 mm sphere embedded within a 1 cm acrylic cube. It was moved in steps of 1 mm through the axial center, and at each point, data was acquired for 60 s. The sensitivity was then calculated for each measurement as

$$S = \frac{N_{\text{Coincs}}}{0.906 \times 60 \text{ s} \times 2.8 \text{ MBq}} \times 100 \%, \quad (2.8)$$

with 0.906 being the branching ratio of ^{22}Na .

Figure 2.5 shows the observed sensitivity profile. One can observe that the peak sensitivity of 2.23 % is slightly lower than our previous theoretical estimate. The simplifica-

tions made in our modeling can, of course, contribute to this discrepancy. Nevertheless, the result seems to align rather well.

For another comparison, we can take a look at SAFIR-II's predecessor. The peak sensitivity of SAFIR-I was measured as 1.45 % at an axial FOV of 54.2 mm [35]. Accounting for the increase in solid angle caused by the longer field of view yields an expected peak sensitivity of 2.78 % for SAFIR-II. However, it should be noted that SAFIR-I's crystals were slightly larger than those of its successor ($2.12 \times 2.12 \times 13 \text{ mm}^3$) due to a difference in the thickness of the utilized separating material. Measurements made with SAFIR-I comparing both crystal sizes have shown a reduction in the measured coincidence rate of about 20 % for the smaller crystals at low activities [26]. This factor is significantly higher than just the reduction in crystal size and can be explained through an additional decrease in the photofraction due to the change in crystal size. Taking this factor into account, one would expect a peak sensitivity of about 2.22 % for SAFIR-II, which is in line with the actual observed value.

While the choice for smaller crystals reduced the system sensitivity, it was made to improve SAFIR-II's performance at higher activities. The study comparing the two matrix types showed a significant reduction in the amount of optical cross-talk for the smaller crystals, improving both the energy and timing precision of the scanner. Additionally, at high measurement activities, the impact of the crystal size on the coincidence rate decreased, amounting to only 10 %. This is likely connected to a reduced data loss from dead time effects, which will be elaborated in section 2.3.3. Lastly, a small deviation in the system sensitivity is not a major concern for SAFIR-II, as it can be easily compensated for by increasing the injected activity.

2.2.4 SiPMs and Overvoltage

While a crystal's material and shape are of high importance with respect to a detector's performance, they are not the only determining factor. Photosensors are required to convert the 420 nm photons emitted by LYSO into an electrical signal that can be digitized and analyzed. Most modern PET scanners utilize SiPM arrays for this purpose, with SAFIR-II featuring Hamamatsu S13361-2050 multi-pixel-photon counters (MPPCs), the same as its predecessor SAFIR-I. An individual SiPM pixel consists of an array of avalanche photodiodes (APDs) connected in parallel. These APDs are supplied with a reverse bias voltage ($\approx 57 \text{ V}$ for SAFIR-II), which creates a depletion region in the semiconductor material. Due to the doping in the junction region and the bias voltage, electron-hole pairs produced via the photoelectric effect are accelerated to velocities high enough to result in the creation of more pairs in an avalanche-like effect. APDs are highly sensitive to photons, being able to generate sufficiently strong signals from single photons alone.

Several parameters are used to describe the performance of a SiPM. The gain denotes the amount of charge an APD generates per detected photon and typically ranges between 10^5 to 10^7 [36]. Furthermore, a SiPM's photon detection efficiency (PDE) denotes the probability that an incident photon results in an output signal. Each SiPM has a peak-sensitive wavelength at which its PDE is highest. In the case of the Hamamatsu

MPPCs employed by SAFIR-II, this value is 450 nm, closely matched to the characteristic 420 nm emission wavelength of LYSO. Thermally generated electrons have the potential of spontaneously triggering a SiPM, resulting in ‘dark counts’ occurring at a certain rate. Lastly, the chance that a charge carrier generated in a SiPM pixel triggers an avalanche in an adjacent cell is known as the cross-talk probability.

The performance of a SiPM is one of the major factors in determining a detector’s timing precision. Commonly, an ASIC connected to the SiPM starts recording a signal the moment it surpasses a pre-selected threshold value while specifying that point as the signal’s arrival time. As such, the signal’s shape and rise time are crucial factors influencing the accuracy of this time stamp, and they are both influenced by the previously mentioned performance parameters. On the one hand, a high gain and PDE are desirable to increase the timing resolution of a detector, as they decrease the signal’s rise time. On the other hand, an increased dark count rate and cross-talk probability counteract this, as the created noise can alter the shape of a signal and result in a wrong time stamp being assigned to a hit. Influencing these parameters is the reverse bias voltage applied to a SiPM or, more specifically, the difference between the applied voltage and the diode’s breakdown voltage. This is referred to as the overvoltage (OV), and increasing it results in a higher gain, PDE, dark count rate, and cross-talk probability. It is therefore imperative to choose not just the highest overvoltage but one adapted to the desired scanner performance. Measurements done by the SAFIR collaboration using the DRP have shown that the scanner provides an excellent timing resolution at 6 V OV [22], which is therefore also the value chosen for SAFIR-II’s operation.

2.3 PETA8: Digitization of Analog Signals

To be able to detect coincidences and reconstruct an image, a PET scanner requires at least three key pieces of information: the energy, arrival time and location of each incoming photon. A SiPM’s initial output is an analog electrical signal with an amplitude and signal shape correlating to the scintillation light emitted by a struck crystal. To make it usable for later data processing, additional electronics are necessary to translate parts of this signal into digital information. For SAFIR-II, this is done using the PETA-8 ASIC, an improved version of the PETA-6-SE used in SAFIR-I.

2.3.1 The Front End: Energy and Timing

PETA-8 was designed at the Institute of Computer Engineering of the University of Heidelberg with the intended use as an ASIC for PET. It features 32 separate input channels and is split into two halves containing 16 channels each. The channels share a single centralized TDC, but are mostly separate beyond that, with each channel being used to process the energy and timing information of signals from a specific SiPM pixel. As each channel is connected to a single pixel of a single SiPM array and, therefore, a single LYSO crystal, information on the photon’s location is known via the specific channel hit. All channels are identical in functionality but can be individually configured to adjust for manufacturing variations of both the SiPM channels and the ASIC itself.




Figure 2.6: Signal processing of the PETA-8 front end for incoming SiPM signals, showing an example for a regular hit (1) and a signal not passing the energy threshold (2). (Adapted from [37]).

Within a channel, an incoming signal from the SiPM passes a fast, low input impedance amplifier before being detected by a fast comparator. This comparator features a configurable ‘timing’ threshold, which is used to filter out thermal noise and only triggers the channel for significantly high amplitudes. Following the triggering of a channel, the value of a TDC common to all channels is saved, denoting the arrival time of a channel’s detected signal. The signal’s current is accumulated using an analog integrator, with the accumulated charge therefore correlating to the energy of the detected photon. The integrator’s gain value can be individually configured for each channel, ensuring that signals generated by photons of similar energies result in a similar charge. After a short wait time, the accumulated charge is then linearly converted into a 9-bit ADC value, which is achieved by discharging the capacitor at a constant rate. Should this ADC value pass a second configurable minimum value, also denoted as the energy cut, both the ADC and TDC values are saved in the channel’s register and ready to be read out. A signal not passing the energy cut is discarded instead. Either way, the channel is reset afterwards, though in the case of an accepted signal only after a completed data readout. All these processes are supervised by a finite state machine (FSM), which is started following the channel being triggered by a signal. A diagram detailing the processing of incoming SiPM signals is detailed in figure 2.6.

The data from an accepted signal is saved as a 33-bit long ‘hit’ packet by the channel. These 33 bits include nine bits for the ADC value, an 18-bit TDC value to denote the arrival time of the signal, and six control bits for various other purposes. The 18-bit TDC is not a single counter value but is formed from three separate values combined. The 5-bit long ‘fine counter’ value is obtained from the output of a linear ring oscillator

driven by a 625 MHz clock supplied externally to the ASIC. As such, it increments at a bin width of 50 ps and has a wrap-around period of 1.6 ns. Additionally implemented is a 3-bit ‘mid’ counter incrementing by one for each period of the fine counter and a 10-bit ‘coarse’ counter doing the same for the mid counter period. Despite this separation into three counters, the resulting TDC is a singular binary encoded value with a bin width of 50 ps. The distinction is nevertheless important to keep in mind for later calibration procedures discussed in section 3.2.3. It is also one of the major differences between PETA-8 and PETA-6-SE, which featured a different TDC logic consuming significantly more power.

2.3.2 Data Readout Architecture

While a channel’s front end determines the energy and timing of an incoming photon, the specific channel detecting the hit serves as the information for the photon’s location. This is closely linked to the data readout, which, in PETA-8’s case, can be operated in two different modes. In the ‘data push’ mode, packets from channels detecting a hit are automatically sent via an 8-bit/10-bit encoding scheme as they become available, while in ‘data pull’ mode, packets are only sent following an external request using a ‘readout load’ signal. For the PETA-6 ASIC utilized in SAFIR-I, only the ‘data pull’ mode was available. Thus, this mode is used by SAFIR-II as well since the firmware architecture required for the FPGA to read data in this mode was already designed.

Each channel saves a detected hit’s data as a 33-bit long binary packet in a designated buffer register. Connected to these buffers is a linear shift register driven by an external clock at a frequency of 280 MHz. As the channels of PETA-8 are grouped into two halves, two such shift registers exist. Following the application of an external ‘readout load’ signal, the 33-bit packets from all 16 connected channels are copied into the shift register, resulting in a single combined 528-bit array. Afterwards, the shift register is read out one bit at a time, with the first bit in each packet consisting of a ‘hit’ flag denoting whether the respective channel contained valid data. Furthermore, a channel containing a hit is cleared after copying data to the shift register, making it available to receive another hit.

However, until a readout is initiated, a channel storing data in its buffer register is blocked from detecting new signals and thus cannot be triggered. Data can only be detected again after ‘readout load’ has been raised (and a short time has passed to reset the channel to its ground state). To alleviate the problem of channels being blocked from detecting hits this way, PETA-8 features an optional ‘double-buffer’ register. This register can be enabled via configuration, which makes it possible for a channel to save one additional packet for readout. In the case of two signals arriving in quick succession, this double buffer register thus prevents a secondary hit from being lost. The channel register and double buffer can only be read out one at a time and, therefore, do not increase the processing speed of the readout overall.

An additional feature exists to increase the processing speed of the readout. Upon receiving a ‘readout load’, all channels copy their current state into the shift register, regardless of a packet actually being stored. The receiving end then has to check the

incoming data stream for enabled ‘hit’ flags and discard data packets with this flag set to zero. This can be expedited using the optional ‘zero suppression’ mode, which can be enabled via configuration as well. When enabled, the additional 32 bits from data packets not containing a ‘hit’ are omitted, and only the zero-valued ‘hit’ flag is part of the data stream for the specific channel. Therefore, the total time to perform a readout is no longer a fixed $\frac{528 \text{ bit}}{280 \text{ MHz}} = 1.8 \mu\text{s}$ but instead shortened depending on the number of channels triggered. This zero suppression mode, as well as the double buffer feature, are two of the major differences between the PETA-8 ASIC used in SAFIR-II and the PETA-6-SE used in SAFIR-I. As such, SAFIR-II performs its readout operations much faster and more frequently.

2.3.3 Detector Dead Time and Data Loss

The specifics of PETA-8’s front end and readout infrastructure are important both when it comes to analyzing and calibrating the raw output data as well as implementing the FPGA firmware needed to acquire it in the first place. However, they also dictate how much time a channel requires to process a signal and how long it takes until said channel can acquire the next one. During this ‘detector dead time’, a channel rejects any additional signals, which results in data being lost. This property is a common problem in particle detection applications and a significant concern due to SAFIR-II’s high event rate.

Pile-up, Paralyzable and Non-Paralyzable Dead Time

A detector’s dead time describes the time period following a signal during which a channel is busy, and no additional signals can be recorded. It is not to be confused with ‘Pile-up’, in which two signals arriving in a channel simultaneously result in a single data point consisting of the summation of both. Typically, two modeling approaches exist for dead time behavior: *non-paralyzable*, in which a detector’s dead time following an incoming signal is fixed, and *paralyzable*, where each signal arriving during dead time is not just discarded but also extends the dead time by a certain amount [38]. Non-paralyzable dead time results in a maximum possible detection rate at saturation, while *paralyzable* dead time results in a saturation behavior where no data is recorded at all. When it comes to the PETA-8 ASIC and SAFIR-II in general, the specifics of the data readout and ASIC design result in a mixture of Pile-up and non-paralyzable dead time depending on the time difference between two incoming hits, as two separate phases take place.

The first situation occurs when two signals arrive within less than 200 ns of each other. In this case, Pile-up occurs as the second signal arrives while the charge from the first is still being integrated. Therefore, the two signals result in a single ADC value consisting of both charges combined. As a result, the resulting hit will not be useful for reconstruction, as the ADC value will not match that of a valid 511 keV photon. The 200 ns here is not an exact value but rather dependent on the energy of the first signal, as the integration time of the front end depends on the ADC value/signal charge. Instead, it reflects the average value for an initial 511 keV photon.

The second situation occurs if a signal arrives after the QDC integration has been completed. Without the ‘double-buffer’ feature, a channel storing a hit that surpassed the energy cut rejects any additional signal until it has been read out, resulting in non-paralyzable dead time. Even with the feature enabled, a short time period passes following the completed ADC integration until another signal can be detected. This complicates the concept of a simple value for non-paralyzable dead time to various different situations, as the arrival time of a signal and the timing of a ‘readout load’ are not correlated. In the best case, if the double buffer is used or a ‘readout load’ arrives immediately following a completed hit acquisition, the channel requires a total of ≈ 400 ns to reset. In the worst case, the time between two ‘readout load’ signals depends on the number of hits previously detected and can be up to 1.8 μ s. It is obvious that reducing the time until a channel is read out inherently reduces dead time, and the double-buffer and zero suppression features of PETA-8’s readout are utilized for this purpose. Nevertheless, detector dead time has the potential to significantly reduce the number of hits detected by SAFIR-II, and it is of high interest to investigate the degree to which this takes place.

Test Trigger

While a channel is usually only triggered following an incoming signal passing the timing threshold, an option exists to bypass this requirement for testing purposes. Should PETA-8 receive a signal at its ‘test trigger’ pin, all channels are triggered for data acquisition regardless of the status of the comparator. The ADC and TDC values are acquired in their current state even if the ADC does not pass the energy cut. They are then stored in the channels register along with a raised status flag indicating that the hit in question was caused by a test trigger.

Beyond making it possible to test whether a channel is responsive, this functionality is useful for two distinct purposes. On the one hand, if the channel receiving the test trigger is not currently processing a SiPM signal caused by an incident photon, then the accumulated charge at the integrator is only caused by noise. Saving this value provides helpful information about the ‘zero point’ of the ADC, being a base value that is added to all detected hits. Later analysis can subtract this value from detected hits to increase the accuracy of the energy measurement. On the other hand, if a channel receiving the test trigger is currently processing a SiPM signal, then the test trigger is discarded, with the hit being stored not exhibiting a raised status flag. This is the same behavior taking place if a regular hit arrives during dead time. Thus, a channel can effectively be queried on whether it is currently processing a signal or not. Expanding on this by repeatedly sending the test trigger signal during regular data acquisition, one can analyze the percentage of test trigger signals received to estimate the effective amount of data lost due to dead time effects.




Figure 2.7: Rendering of the line source phantom. 3D-printed cylindrical spokes are used to position the phantom within the axial center of the FOV. (Rendering by R. Becker).

Measurement

To test the validity of this method, I performed a series of measurements using SAFIR-II, during which test trigger signals were injected into all channels at a frequency of 200 Hz. The returning data was analyzed for each measurement, and the number of hits received that featured a raised test trigger flag was measured. The average percentage of such events not received/lost per second was evaluated across the entire scanner as the ‘loss rate’.

The first measurement included a thin cylindrical line source filled with up to 500 MBq ^{18}F placed along the central axis of the scanner. The phantom, depicted in figure 2.7, has a fillable volume of about 1.2 ml consisting of a cylinder with a diameter of 4.5 mm and a length of 75 mm. Measurements were performed at various source activities by letting the source decay. The average percentage of test trigger hits lost across all channels was acquired, as well as the rate of detected hits with an energy between 400 keV to 600 keV (‘singles’). During all measurements, the double-buffer and zero suppression features were enabled.

Figure 2.8 shows the observed rate of singles in relation to the measurement activity. At first glance the data appears to be linearly dependent. However, a look at the residual between the measurement and a linear fit through the three data points at the lowest measurement activities shows a decrease towards higher activities. This hints at dead time effects impacting the data acquisition of SAFIR-II, resulting in a $\approx 2.3\%$ loss of data at the highest measurement activity.




Figure 2.8: **Left:** Rate of singles detected by SAFIR-II during a line source measurement of ^{18}F . The dashed line represents a linear fit through the three data points at low measurement activities (<50 MBq), where the least amount of data loss is expected. **Right:** Residual between the measured singles rate and the linear fit.




Figure 2.9: Percentage of test trigger hits lost ('loss rate') across the entire scanner during the line source measurement. The dashed line represents a linear fit through the data.




Figure 2.10: Picture of the calibration phantom. During a measurement, the phantom is held in the center of the FOV and filled with ^{18}F diluted in water.

Figure 2.9 shows the average percentage of test trigger hits ('loss rate') received during the same measurement. Though subject to statistical fluctuation, one can observe an approximately linear increase towards higher activities. The observed loss rate at the highest measurement activity also aligns well with our previously observed loss of data.

It should be mentioned that the measurements using the line source were performed at an energy cut of 70 LSB. At this value, a significant fraction of the Compton spectrum is cut off, reducing the amount of data that the scanner transmits to the DAQ-PC and, therefore, also the time required to analyze said data. Adjusting the energy cut to a lower value would increase the amount of data and thus result in more hits being saved by and read out from each channel. It is of particular interest to observe whether an adjustment of the energy cut would impact the data loss as well. For this purpose, I performed a second measurement, this time utilizing the cylindrical 'calibration' phantom depicted in figure 2.10. The phantom features an inner diameter of 30 mm and an inner length of 50 mm, holding a total volume of about 35 ml. Measurements were performed similarly to the methodology involving the line source, with the phantom placed in the center of the FOV and filled with up to 400 MBq. At each measurement activity, the measurement was repeated for various energy cuts between 40 LSB and 80 LSB.

Figure 2.11 shows the amount of data collected by SAFIR-II for each measurement. As expected, the amount of data acquired increases for lower energy cuts, with a value of 40 LSB requiring nearly twice as much storage space as 80 LSB. Figure 2.12 shows the measured rate of singles for all measurements, as well as the observed loss rates. Surprisingly, the rate of singles does not seem to be significantly impacted by the energy cut, with the same being true for the loss rate. This is likely the consequence of the PETA-8's data readout. Specifically, the 'double-buffer' and 'zero suppression' options enable the ASIC to swiftly clear a channel once a hit is acquired. As a result, the main contributing factor to PETA-8's dead time seems to be the channel front end acquiring data.

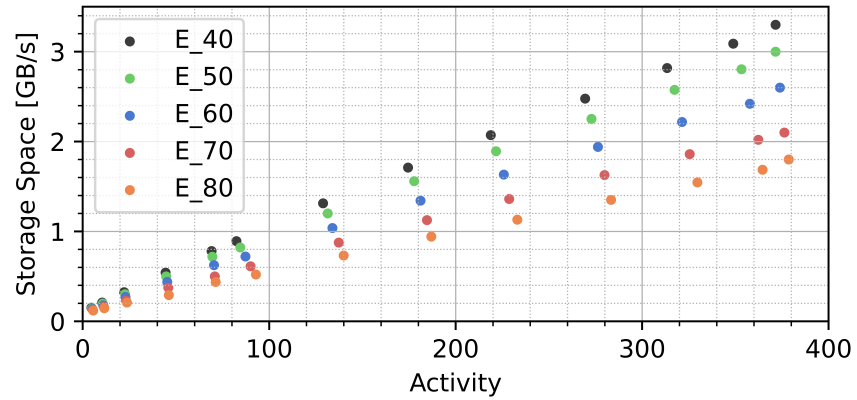


Figure 2.11: Storage space required by SAFIR-II’s data acquisition per second for various measurement activities and energy cut values. As the DAQ-PC saves all raw data it receives from the insert, the consumed storage space is equivalent to the data bandwidth across the SFP Ethernet links.




Figure 2.12: **Left:** Number of singles detected by SAFIR-II during a measurement of the calibration phantom filled with ^{18}F diluted in water, at various measurement activities and energy cuts. **Right:** Percentage of test trigger hits lost (‘loss rate’) across the entire scanner during the same measurement.




Figure 2.13: **Left:** Top and bottom view of the SP8 PCB. Visible are the four PETA-8 ASICs and three male SAMTEC ST4 connectors at the bottom, and four female SS4 connectors on the top. **Right:** Fully assembled view of the detector head module, with one 7×8 and one 8×8 LYSO-crystal matrix glued to two Hamamatsu SiPM arrays mounted on top of an SP8 PCB. Also barely visible is a copper heat sink attached to the bottom of the SP8 used to cool the PETA-8 ASIC.

2.3.4 SP8: The Detector Head Module

The combination of scintillator crystals, SiPMs, and readout ASICs forms the functional core of a PET scanner, as these are the main components necessary to detect and digitize the characteristic 511 keV photons. Going forward, this combination will be referred to as the ‘detector head’, and in the case of SAFIR-II specifically, as the ‘detector head modules’. The detector head of SAFIR-II consists of a cylindrical arrangement of 96 such modules, with each module grouping together the same number of crystals, SiPM arrays, and ASICs.

The basis of each detector head module is the SP8 PCB, which is shown in figure 2.13. The SP8 is a small (18 mm by 39 mm) rectangular PCB featuring seven SAMTEC ST4/SS4 connector plugs as well as four PETA-8 ASICs. The seven connectors are split into three male plugs at the bottom of the SP8 and four female plugs at the top. Two Hamamatsu SiPM arrays, each featuring two male connectors and carrying LYSO crystal matrices bonded on top of them via optical glue, can be plugged into the top connectors of the SP8 to connect them to the PETA-8 ASICs. This leaves a gap of 0.5 mm between the matrices and, therefore, widens the pitch between the crystals across this gap. Additionally, the crystal matrices glued on top of the SiPM arrays are not identical. Instead, one array carries an 8×8 matrix, and one holds a 7×8 matrix, leaving one row of SiPM pixels empty and resulting in a combined 15×8 crystals. As previously mentioned, the 2 mm by 2 mm crystals of each matrix are separated and covered by a 0.2 mm thick sandwich layer of ESR foil and aluminum. An image of a fully assembled detector head module is visible in figure 2.13 as well.

2.4 Supporting Electronics

The crystal matrices, SiPM arrays, and PETA-8 ASICs form the functional core of SAFIR-II, acquiring the data needed for later analysis and image reconstruction. However, said data needs to be transferred from the ASICs and passed to a PC capable of processing it. Furthermore, each ASIC needs to be properly configured and calibrated, as well as supplied with fast control signals, to ensure a synchronous operation of the timing counters. Lastly, both the SiPM arrays and the SP8 PCBs need to be supplied with power, with the bias voltage for each SiPM being adjusted individually to ensure a proper operating voltage. Therefore, several supporting PCBs are present within the insert with the explicit purpose of executing these tasks.

2.4.1 Digital Control: SDIP3

If the detector head is the eye of SAFIR-II, collecting information about events it observes, then the SDIP3 PCBs are its brain, processing said information and controlling how the scanner operates. Twelve SDIP3 PCBs are present within SAFIR-II, each serving as a mounting point for eight detector head modules. The SDIP3 also serves as a connection point for all other supporting PCBs, being in direct connection with the SBTV3, SSPD, and SFCD, as well as in indirect connection to the power distribution ring boards SBD and SPPD. An image of the full SDIP3 assembly is shown in figure 2.14.

The detector head modules plug into SAMTEC connectors at the front of the SDIP3, enabling the swift replacement of a module in case of an electrical failure. Through this connection, both the configuration and data readout of the PETA-8 ASICs are performed using a Xilinx Kintex-7 FPGA. The data readout is handled via a 280 MHz serial data connection, with firmware in the FPGA accumulating valid hits from all eight SP8 modules. This firmware takes advantage of the ‘zero suppression’ and ‘double-buffer’ features of PETA-8 mentioned in section 2.3 to perform read operations as frequently as possible. It also bundles all received hits into Ethernet frames to be sent to the DAQ-PC. To ensure proper allocation of hits to their respective locations during analysis, each event in an Ethernet frame is supplied with a 9-bit ‘ID code’ indicating the ASIC and respective channel from which it was read out. Ethernet frames are sent to the DAQ-PC via a 10Gbit SFP+ module (Finisar FTLX8573 [39]), which is located next to, connected to, and controlled by the FPGA. Configuration of the ASICs occurs through the same connection but in the inverse direction, with the DAQ-PC sending the necessary data via the Ethernet link and the FPGA passing it on to the ASICs. This configuration is performed using a JTAG interface between the FPGA and the ASICs, which also makes it possible to read out configuration registers from the ASICs to check if a configuration was sent successfully.

All SDIP3 require four fast control signals to ensure proper and synchronous operation of the scanner. These signals are distributed to each SDIP3 via LVDS signaling from the SFCD PCBs, which are connected via standard RJ-45 ribbon cables. The signals in question consist of a power enable signal to start the scanner, a common 31.25 MHz clock



Figure 2.14: Full assembly of the SDIP3, including eight detector head modules on the left, the attached SBTV3 PCB, SSPD PCB and DC-DC power converters on the right, as well as the Kintex-7 FPGA towards the center. Not visible is the 10G SFP+ Ethernet connector and glass fiber, hidden below power converters and the SBTV3 PCB.

signal, a synchronously delivered reset signal, and the previously mentioned test trigger. ‘Power enable’ is raised on system startup and lowered for system shutdown, controlling the power distribution to the FPGA and other components. The 31.25 MHz clock is used as a base signal to generate various further clock signals for the SDIP3. The clock signals consist of a 312.5 MHz signal required for the Ethernet transmission, several internal clock signals for the FPGA operation, and a 625 MHz clock required by the PETA-8 ASIC to drive its timing counters. Ensuring the synchronous operation of these timing counters is the reset signal, which is used to set all counters to zero simultaneously. When received via the RJ-45 cable, it is passed through a clock fanout buffer IC ("ADCLK846") to all detector head modules. Lastly, the test trigger signal is distributed in the same fashion. While the clock signal is constantly supplied by the SFCD, the other three are triggered following a user command from the DAQ-PC. For this purpose, one of the twelve SDIP3 is additionally connected to the SFCM PCB via transistor-transistor logic (TTL) signaling and designated as a ‘master’ board. An Ethernet command from the DAQ-PC results in this master board signaling the SFCM, which in turn simultaneously distributes the desired signal to all 12 SDIP3 via the SFCD.

The FPGA firmware features additional functionality beyond just the data and configuration transfer. Temperature sensors placed along the SDIP3 are periodically read out, and their data is relayed to the DAQ-PC to help monitor the scanner. Furthermore, an additional timing counter is implemented, which is added to each Ethernet frame containing data from the PETA-8 ASIC. This timing counter, dubbed the ‘epoch counter’, is reset along with the TDC of PETA-8 and increments whenever the TDC overflows, effectively extending the TDC wraparound period. Lastly, the FPGA is capable of adjusting the speed of the ASIC data readout, being able to change the readout frequency in steps of 10 MHz. However, this feature is not used during regular operation, and the readout frequency is kept constant at 280 MHz.

As it is one of the most complicated PCBs in the scanner, the SDIP3 underwent a significant amount of changes from the SDIP2 PCB used in SAFIR-I. These changes include five additional mounting points for detector head modules along with the corresponding required additional signal layers and changed layout. Furthermore, adjustments to the PCB layout and FPGA firmware had to be made to accommodate the new PETA-8 ASICs. The firmware and layout also had to adapt to the 10Gbit SFP+ module, which replaced a 1Gbit link (Avago AFBR-354 57R5APZ).

2.4.2 Power Conversion: SSPD

With the exception of the SiPM arrays themselves, all electronic components mounted on the SDIP3 operate at voltages between 1 V and 4.1 V. At these respective voltages, the scanner requires a current of more than 100 A. Supplying these voltages externally is infeasible, as it would require unreasonably thick cables. Combined with the 10 m length needed to connect the power supplies located outside of the MRI's reach, this would also cause significant losses and produce large amounts of heat across these cables. Therefore, SAFIR-II features MRI-compatible DC-DC converter modules within the insert itself [40].

Attached to each SDIP3 is the SSPD PCB. It receives a singular input voltage of at least 12 V from the SPPD PCB and features four converters with an output voltage of 2.4 V each. These converters are used to supply the analog and digital voltages required by the PETA-8 ASICs. Two converters are utilized for digital power and two for analog power, respectively, resulting in each converter serving four SP8 modules. These four converters power up sequentially, as PETA-8 requires digital power to be supplied first to ensure proper operation. Failure in a converter used for digital power causes the respective analog power converter to be shut down as well to prevent damage to the PETA-8 ASICs. The SSPD also passes the 12 V input voltage to the SDIP3 itself, which features one 4.1 V and one 2.4 V converter to provide the voltages required by the FPGA and SFP module. Lastly, on the SDIP3 itself, several low-dropout (LDO) regulator integrated circuits (ICs) are installed to further convert the DC-DC converter voltages into whatever is required by the assembled components.

The split of the four converters between analog and digital power supply is another change between SAFIR-II and its predecessor. For SAFIR-I, one power converter was used per detector head module to supply both analog and digital power. This would have resulted in SAFIR-II utilizing eight converters per SSPD, requiring a much longer PCB. Additionally, as the PETA-8 ASIC used in SAFIR-II only consumes about half as much power as the PETA-6-SE, fewer power converters were required overall.

2.4.3 Bias Voltage Control: SBTV3

The operating voltages of the individual SiPM arrays are not identical but vary from component to component in a range from 50.5 V to 52 V due to manufacturing imperfections. Furthermore, these voltages are temperature dependent at a coefficient of 54 mV/°C. Thus, they might vary depending on the measurement activity, as higher count rates result in the SiPM's heating up. This poses a problem since measurements have shown that both the energy and timing resolution of these matrices, as well as their PDE, are highly correlated to the applied voltage, with an OV of 6 V above the breakdown voltage being optimal for SAFIR-II's applications [22]. Thus, SAFIR-II requires precise control of the bias voltage applied to each SiPM array and the ability to swiftly adapt it as necessary in case of temperature changes to ensure optimal performance of the scanner at all times.

This task is not performed by the SDIP3 itself, which only serves to pass these voltages to the detector head modules via the SAMTEC connectors. Instead, mounted above the




Figure 2.15: The SBTV3 bias control PCB used in SAFIR-II.

SDIP3 and SSPD atop 3D-printed PLA structures lies the SBTV3 PCB. This PCB, depicted in figure 2.15, features eight output voltage channels driven by a 14-bit DAC (AD5648 [41]), each connected to one detector head module through the SDIP3. Only eight channels are required since, for a given detector head module, both SiPM arrays are matched such that their average operating voltages are within 0.01 V of each other. Control of the DAC is facilitated by a microcontroller IC (ATMega328 [42]), which in turn receives commands from the SDIP3 FPGA via a connecting cable. The SBTV3 is supplied with power separately from the SDIP3 via the SBD PCB. It receives a 50 V ‘shift’ voltage as well as two additional ± 12 V floating voltages, which are utilized by the DAC to generate the required bias voltage for each SiPM at a resolution of about 1.2 mV. Lastly, the SBTV3 also features the readout of up to eight external LMT01 temperature sensors, with two further sensors placed on the SBTV3 itself. However, in the current design of SAFIR-II, no external temperature sensors are used, and only the sensors on the SBTV3 itself are read out.

The SBTV3 is another PCB that saw adjustments between SAFIR-I and SAFIR-II, mainly as it required additional channels due to the extra detector head modules.

2.4.4 Fast Control Signals: SFCM & SFCD

The twelve sectors within SAFIR-II start and operate almost fully autonomously. While they are supplied with power from the same source and sent commands by the same DAQ-PC, the individual SDIP3 PCBs do not communicate with each other. However, they need to operate in a synchronous fashion for a few key functionalities, which is ensured using four control signals shared between them. The SFCM and SFCD PCBs depicted in figure 2.16 are designed to deliver these signals to all twelve sectors quickly and simultaneously. They work in tandem and are located at the back flange of the scanner. The SFCM is connected to a single ‘master’ SDIP3 via small coaxial cables and raises a control signal whenever required by the SDIP3. This is the case whenever the master SDIP3 receives an Ethernet frame requesting such a signal from the DAQ-PC. Control signals generated by the SFCM are passed to the two SFCD PCBs, which distribute the signals to six SDIP3 boards each using clock fanout buffer ICs ("ADCLK846") at a jitter of 100 fs and regular Ethernet ribbon cables.

These signals, as mentioned in section 2.4.1, consist of the ‘power enable’ to start




Figure 2.16: The SFCM (left) and SFCD (right) PCBs used in SAFIR-II.

the scanner, a common 31.25 MHz clock, as well as the reset and test trigger signals. The 31.25 MHz clock signal is always sent by the SFCM regardless of control via the DAQ-PC, being generated by a crystal oscillator and distributed at a jitter of less than 100 fs (LMK00301 [43]). Both the test trigger and the reset are generated following TTL signals from the master SDIP3. They are distributed from the SFCM at a jitter of less than 100 fs as well (IDT8P34S1208I [44]). ‘Power enable’ is not timing-critical and is the only signal not distributed from the SFCD using the fanout buffer ICs. Instead, it is automatically generated by the SFCM at startup and passed to each sector sequentially via the SFCD with a 0.85 s delay between sectors. As two SFCDs are used, two sectors are always started simultaneously, resulting in the power-up sequence completing in 5.1 s. Power-down follows an inverse procedure, with the ‘Power enable’ signal being lowered following an Ethernet command by the DAQ-PC and being passed to sectors sequentially.

2.4.5 Power Distribution: SPPD & SBD

The voltages required for SAFIR-II’s operation are generated outside of the insert itself by six commercial power supply units located in a room outside of the MRI’s Faraday cage. These power supplies are connected to the insert via 10 m long cabling and grouped into three pairs, each pair providing a different set of voltages for the scanner. Within the scanner itself, ring boards located at the back end distribute these voltages to the individual detector sections, similar to the previously mentioned SFCD & SFCM boards. Two boards of each type are present, each connecting to six sectors. This way, an individual supply in a pair provides power to one half of the scanner, being connected to one of the two distribution PCBs.

The main functional power for SAFIR-II’s digital electronics is provided by two switching mode power supplies (TDK Lambda Z+ [45]) operating at 18 V, and distributed by two SPPD ring PCBs. The SPPD, depicted in figure 2.17, is a passive PCB that only features capacitors for voltage stabilization and protective diodes to prevent voltage reversal. Beyond that, it merely forwards the 18 V main power to the six SSPDs. One of the two SPPDs additionally supplies power to the SFCM, with the SFCM itself featuring a 3.3 V DC-DC converter and connections for the two SFCD PCBs, which operate at the same voltage.




Figure 2.17: The SPPD (left) and SBD (right) PCBs used in SAFIR-II.

The three bias voltages required by the SBTV3 and the SiPM arrays are distributed in the same fashion, though they share a singular board in the SBD, also shown in figure 2.17. The 50 V shift voltage is provided by two linear single-output power supplies (Keysight E3645A, 60 V, 1.3 A max), while the ± 12 V floating voltages are provided by two dual output supplies of the same version (20 V, 1.5 A max). The SBD itself does not feature additional components and supplies these voltages to the connected SBTV3 with a common ground.

The SBD and SPPD, as well as the previously described SFCM and SFCD PCBs, have remained unchanged from SAFIR-I. The only notable change consists in the power supply units and their cabling, as SAFIR-I consumed less power and thus only required one of each type, while SAFIR-II uses two of each.



Figure 2.18: Carbon fiber mechanical support structure of the SAFIR-II insert. The 12 ‘caskets’ in the central part are additionally covered with carbon fiber lids (not visible here), with the entire structure being surrounded by carbon fiber covers as well. Already mounted and visible on the right are the SPPD, SBD and SFCD ring boards.

2.5 Mechanical Support

A mechanical support structure is required to position the various measurement electronics utilized by SAFIR-II and ensure the desired geometrical arrangement of the crystals. This structure, displayed in figure 2.18, is for the most part made of molded carbon fiber. It is constructed around a 69 cm long cylinder with an inner diameter of 114 mm and a thickness of 2 mm. Mounted around this cylinder are 12 ‘caskets’ containing the fully assembled SDIP3 boards, including crystal matrices and DC-DC converters. Also mounted around this cylinder, towards the end of the SDIP3 boards, is a 3D-printed ring structure used to position the power and fast signal distribution boards. Capping of the ends of the scanner at both sides are flanges molded out of the same material as the cylinder itself. These flanges are required to position the scanner within the MRI, containing extendable rubber feet that can be used to lock the insert in place. They also hold the scanner’s covers and contain smaller holes for cables to pass through, as well as large holes to allow airflow through the scanner for cooling.

2.5.1 Air Cooling and Insulation

During data acquisition, SAFIR-II consumes about 580 W of electrical power and consequently produces the same in heat. Most of that heat is generated by the 384 PETA-8 ASICs utilized in the detector head, which themselves amount to about 310 W through their analog circuitry alone. The obviously resulting risk of overheating, combined with the increased thermal noise resulting from the close proximity of the ASICs to the SiPM arrays, shows that proper cooling of the SAFIR-II insert is of high importance.

Due to this, SAFIR-II is integrated into a closed-loop air cooling circuit driven by a side channel blower (SCL K05-MS-2.2kW [46]). The blower draws hot air out of the scanner, which is then directed into a 1.6 kW cooling unit (MLT Type 311-W-RB400 [47]).

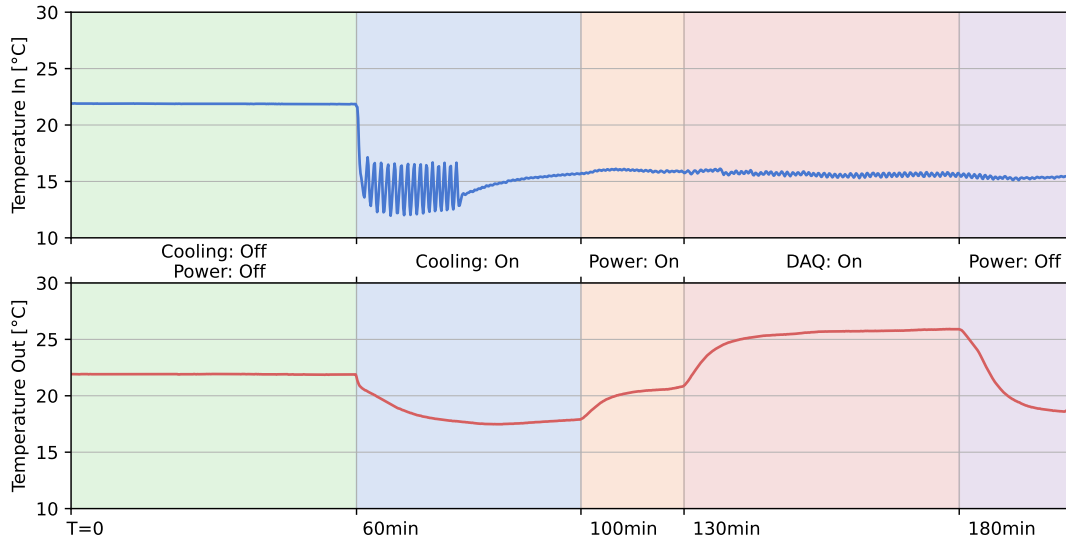


Figure 2.19: Air temperatures of SAFIR-II’s closed loop cooling circuit, going into (top) and out of (bottom) the scanner, during various states of system operation. Following its start ($T=60$ min), the cooling unit requires some time to stabilize to the target temperature of 15°C . The scanner is supplied with power at $T=100$ min, and analog power to the ASICs is supplied at $T=130$ min.

Following this, the cooled air is passed through tubes back into the scanner, completing the cycle. The airflow through the scanner is oriented front-to-back, meaning that cool air meets the detector head including the SP8 PCBs first, since they generate the most heat. Additionally, this helps to keep the SiPM and ASIC temperatures as stable as possible. To better dissipate heat from the PETA-8 ASICs, custom copper heat sinks are screwed to the bottom of the SP8 PCBs, being in thermal contact with the ASICs via thermal paste (Thermal Grizzly Kryonaut [48]). After cooling the detector head, the air traverses the rest of the SDIP3 and, lastly, passes over the ring boards in the back flange before leaving the scanner again. USB controllable sensors (Yocto-Meteo-v2, [49]) are placed within the cooling circuit to monitor the incoming and outgoing air temperatures of the insert, as well as measure the air humidity. These sensors are continuously read out during operation via the DAQ-PC, allowing the scanner to be shut down should an unexpected rise in the air temperature or other concerning behavior be observed. Figure 2.19 shows the temperature of the air entering and exiting the scanner during several states of operation.

Current measurements using SAFIR-II, including all measurements mentioned in this thesis, have been performed at an input air temperature of 15°C , which has proven sufficient for the proper operation of the scanner. Future measurements might be conducted at lower temperatures, but proper insulation and sealing of the circuit are required to prevent the buildup of condensation water, as these lower temperatures might cross the dew point in the laboratory. The inner cylinder of the insert has already been insulated

utilizing polystyrol, as it is vital to prevent the air around the animal (and thus the animal itself) from being thermally influenced by the scanner.

2.5.2 Shielding and MRI Compatibility

An essential requirement for SAFIR-II's construction is its compatibility with the MRI. Due to its nature and design as a PET-MRI insert, it is imperative that SAFIR-II does not perturb the MRI measurement and vice versa. Modern MRI research has brought forth a large number of available pulse sequences and techniques capable of acquiring images with various contrasts and foci. However, going into detail on such a vast field of research would exceed the scope of this thesis, so the interested reader is referred to Haacke et al. [50] for further information. The main point of interest here is that all MRI imaging is based on a complicated interplay of radiofrequency (RF) pulses and temporary gradients applied to a strong, homogeneous base magnetic field known as the B_0 field.

RF Emissions

Most of the mechanical structure of SAFIR-II has been constructed using carbon fiber as a material. To be more specific, molds of the individual components were created and filled with sheets of carbon fiber mesh, as well as a mixture of epoxy resin and graphite dust, before undergoing a thermal curing process to harden them. The utilization of carbon fiber and dust was a deliberate choice to make use of their RF shielding properties [51]. As it employs digital logic for its operation, it is inevitable that SAFIR-II generates RF emissions originating from the clock signals of the components used, data transmission lines, and similar sources. Shielding such emissions is one way of ensuring that they do not influence the MRI's RF coil. A second option relies on controlling the emitted frequencies.

During an MRI pulse sequence, the RF coil sends and receives signals with frequencies close to the Larmor frequency ω of an atomic element present within the patient. This frequency depends on the specific element and on the strength of the MRI's B_0 field. For hydrogen within the $B_0 = 7\text{ T}$ field of the Bruker BioSpec 70/30, the Larmor frequency ω can be calculated as

$$\omega = \gamma \times B_0, \quad \gamma(H_2) = 42.58 \frac{\text{MHz}}{\text{T}}. \quad (2.9)$$

This shows that emissions of frequencies around 298 MHz have to be avoided as much as possible. For this purpose, no clock signal utilized within SAFIR-II's electronics operates at frequencies between 290 MHz and 310 MHz.

However, potential emissions can also occur from the data readout lines and power conversion. Such emissions do not exhibit a singular characteristic peak, instead covering a much wider frequency range, which may intersect with the RF coil's sensitive frequencies. Since not all of these sources can be easily eliminated, it is necessary to add RF shielding to the scanner to ensure that as few of these signals as possible reach the coil. In addition to the carbon fiber used for this purpose, thin sheets of copper have been placed around the central cylinder interlaced with insulating Kapton tape, as shown



Figure 2.20: Copper-Kapton shield surrounding the inner cylinder of SAFIR-II. Each copper strip is cut to be no larger than 2 cm in width. While all strips overlap such that no gaps in the cylindrical structure are present, the strips are electrically insulated from each other by the layers of Kapton.

Operating Condition	Normalized SNR	Deviation
Baseline (insert powered off)	$(3521 \pm 5) \text{ mm}^{-3}$	0.0 %
Digital power on	$(3257 \pm 30) \text{ mm}^{-3}$	7.5 %
Digital & Analog power + data readout	$(3196 \pm 12) \text{ mm}^{-3}$	9.2 %

Table 2.2: Normalized SNR values for various states of SAFIR-II’s operation. The reported values are the mean of 5 separate measurements, with the mean average error reported as the uncertainty.

in figure 2.20. The DC-DC converters used for power distribution within the insert have been designed for MRI compatibility as well [40], featuring an additional copper housing for RF shielding.

While shielding is helpful in reducing the amount of RF emissions reaching the MRI receiver coil, the coil’s high sensitivity will inevitably still observe a small degree. To evaluate the interference of SAFIR-II’s operation on the MRI image acquisition, I utilized a quality assurance sequence provided by the MRI manufacturer to observe the MRI-SNR. This sequence involves imaging a cylindrical phantom homogeneously filled with a calibration solution provided by the manufacturer (H_2O , with $1 \text{ g l}^{-1} \text{ CuSO}_4$ and $3.6 \text{ g l}^{-1} \text{ NaCl}$), and comparing the signal within the phantom to the voxel values outside it. Measurements were taken and used to compare the SNR both during and without SAFIR-II’s operation. The results are summarized in table 2.2. It can be observed that SAFIR-II’s operation leads to a degradation in the MRI-SNR of about 10 % during full data acquisition. I consider this value to be acceptable, as it is still well above the SNR minimum of 2300 mm^{-3} quoted by the manufacturer [52]. The observed images from this sequence also show no artifacts, which would have been an indication of a strong emission at a frequency close to 298 MHz.




Figure 2.21: Maps of the B_0 field (a) without SAFIR-II installed, and (b) with SAFIR-II installed. The dotted line indicates the shimmed area. (a) shows a maximum deviation of ± 0.54 ppm, (b) deviates by a maximum of ± 0.62 ppm

B_0 field Homogeneity

To ensure proper acquisition of an image, the B_0 field of an MRI has to be as homogeneous as possible in the region of interest. However, the presence of the SAFIR-II insert's electronics and the materials utilized in its assembly are bound to influence this magnetic field. Fortunately, modern MRI systems are capable of applying a 'shimming' to the B_0 field, meaning they can correct inhomogeneities using additional coils designed for this purpose. However, these corrections have their limits, and thus, the materials utilized within SAFIR-II are mostly nonmagnetic, with the electronics containing next to no coils.

It is nevertheless imperative to see whether the B_0 field retains its homogeneity when inserting SAFIR-II. Therefore, I performed a second evaluation using a quality assurance sequence capable of mapping the B_0 field within the MRI. It involved the same phantom as mentioned before and was measured after shimming with SAFIR-II inserted and removed from the MRI bore. Figure 2.21 shows maps of the MRI's B_0 field before and after inserting SAFIR-II. It can be observed that the deviations of the magnetic field are less than 1 ppm in both cases. Inserting SAFIR-II shows a slight increase in the maximum deviation by less than 0.1 ppm, but this can just as easily result from physical movement of the phantom during the insertion and is not deemed as significant.




Figure 2.22: **Left:** Schematic sequence diagram of MBEST (Modulus-blipped echo-planar single-pulse technique), an early EPI technique. **Right:** Traversal of k-space by MBEST. Frequency encoding is represented along the horizontal axis, phase encoding along the vertical. (Adapted from [53]).

Eddy Currents

Even when properly shielded and kept nonmagnetic, electric components and circuit boards are inherently conductive. Hence, they are also subject to inductive eddy currents, which can be created by the switching magnetic gradients applied during MRI imaging sequences. These eddy currents can inversely influence the magnetic fields and thus interfere with the imaging sequence.

To investigate the significance of the eddy currents created within SAFIR-II within the scope of MRI compatibility, I used a third quality assurance metric provided by the MRI manufacturer. The sequence ‘EPI-QA1’ acquires an image of the previously described phantom while incorporating fast echo-planar imaging (EPI) techniques. EPI is based on sending an RF pulse, which is quickly followed by continuous application and inversion of the gradient field. Due to the strong reliance on switching gradients, this method is very susceptible to eddy current influence and is thus also useful in gauging such. A schematic diagram for a typical EPI sequence protocol is displayed in figure 2.22.

The EPI-QA1 sequence was run once with an empty bore as a baseline and once after inserting SAFIR-II. Figure 2.23 shows the resulting images of the two acquisitions. Following the insertion of SAFIR-II, artifacts known as Nyquist N/2 Ghosts [54] are clearly visible in the EPI image. This indicates that the eddy currents induced by SAFIR-II are influential enough to be problematic for image acquisition. Following the completion of SAFIR-II, these artifacts were a surprising finding, as the previously built SAFIR-I insert did not exhibit them. The cause can be identified as the SDIP3 PCB’s ground planes, which in previous designs were fragmented into smaller sections, but for SAFIR-II were kept as continuous planes to create a uniform ground and reduce ‘ground bounce’ [55].

As a simplification, eddy currents induced within a PCB’s ground plane can be regarded as individual resistor-inductor (RL) circuits. Following the application of a magnetic field gradient, a set eddy current starts to decay as external magnetic flux is no longer present. This decay is not immediate due to Lenz’s law but rather follows the equation

$$V = V_0 e^{-t \frac{R}{L}}. \quad (2.10)$$




Figure 2.23: Images taken of a cylindrical water phantom using the EPI-QA1 sequence (a) without inserting SAFIR-II into the MRI bore, and (b) with SAFIR-II inserted. Clearly visible in image (b) are the Nyquist ghosting artifacts caused by eddy currents, which are lacking in (a).

The ground plane, in this case, can be viewed as both the resistance R due to its inherent, material-dependent electric resistance and the inductance L in the form of a circular conductive loop. Simplifying further, the inductance of a long coil can be calculated via

$$L = \mu_0 \cdot \mu_r \cdot \frac{N^2 \cdot A}{l} \quad (2.11)$$

with N being the number of windings, A the area enveloped by the coil's loops, l its length, μ_0 the magnetic permittivity of vacuum, and μ_r the permittivity of the filling material. While the eddy currents within a ground plane deviate significantly from a long coil, it is still reasonable to assume that the inductance L increases when the current is circulating over a larger area. This larger inductance then, in turn, increases the RL circuit's decay constant.

The strength and timing of the magnetic field gradients applied during an EPI acquisition are highly dependent on the desired image's dimensions, with a larger image or smaller voxel size requiring stronger and more frequently applied gradients. A plausible explanation for the newly observed Nyquist ghosting consists of the eddy currents persisting longer than before and not decaying fast enough for the selected image dimensions. However, there is also a quick solution to the observed imaging artifacts. By decreasing the number of voxels used for a given image acquisition, the eddy currents are given additional time to decay, reducing the influence of such artifacts. To test this, the EPI sequence was adjusted to create a smaller image, and the acquisition was repeated. An image of this new EPI acquisition is visible in figure 2.24. One can observe that the adjusted sequence no longer shows ghosting artifacts.

While this adjustment does decrease the FOV somewhat, it shows that only small changes are required to adapt sequences to SAFIR-II's presence. Eddy currents are an unavoidable fact of MRI image acquisition, and as such, better and more sophisticated methods exist to compensate and correct for them. One such option is active eddy




Figure 2.24: Image taken of a cylindrical water phantom using the adjusted EPI-QA sequence with smaller image width (64 mm \Rightarrow 55 mm) and SAFIR-II inserted. A lack of Nyquist ghosting artifacts shows the usefulness of this adjustment.

current compensation using coils specially designed for this purpose. However, these adjustments require intervention by an MRI service engineer and were unavailable to me at the time of writing. Future measurements with SAFIR-II might be performed following such an intervention, thereby improving the MRI compatibility of the scanner. Should additional SAFIR scanners be constructed, they will undoubtedly include an adjustment to the SDIP3 ground planes, as previous measurements using the SAFIR-I insert show this to resolve the issue [24].

Influence of the MR on SAFIR-II

While it is important to ensure that SAFIR-II does not influence the MRI acquisition, it is equally as important that the MRI does not influence the operation of SAFIR-II. The complicated interplay of switching magnetic gradients, RF pulses, and the strong B_0 field can influence SAFIR-II's electronics. In the worst case, this could cause issues during data acquisition. To gauge the MRI's influence on SAFIR-II, I performed several measurements using a point source placed in the scanner's FOV. During these measurements, various MRI sequences were executed, and the rate of observed coincidences was recorded. The results, shown in table 2.3, show that SAFIR-II's operation seems to be unaffected by the MRI.

To summarize, SAFIR-II is compatible with the Bruker Biospec 70/30 within which it operates. While some effects on the MRI operation were observable, most of these influences are negligible in effect. The B_0 field homogeneity can be restored via shimming, and the MRI-SNR still resides well above the values required by the manufacturer, thanks to the shielding of SAFIR-II's RF emissions. Influence on the switching gradients during EPI due to eddy currents caused by SAFIR-II is present but can be remedied through adjustments to the acquisition sequences. These results certify the compatibility of SAFIR-II and the MRI system, paving the way for simultaneous PET-MRI acquisition.

MRI Sequence	Coincidence Rate [s ⁻¹]
None (Baseline)	5.519×10^4
B_0 -Mapping	5.53×10^4
Localizer	5.526×10^4
EPI-QA	5.526×10^4
T2-Turborare	5.525×10^4

Table 2.3: Coincidence rate observed by SAFIR-II during various MRI sequences. A ^{22}Na source was placed within the FOV next to the MRI coil, and measured for 60 s per measurement.

Statement of personal contribution: All measurements and results shown in this chapter have been performed and evaluated by the author. Further contributions include the electrical and functional testing of the mentioned SP8, SDIP3 and SBTV3 PCBs, as well as any firmware and software adjustments necessary for the testing and operation of said PCBs. Additional testing included the SFP modules, crystal matrices, power supplies and their control, as well as the temperature monitoring. SAFIR-II's design and components were created by members of the SAFIR-collaboration. The evaluation of SAFIR-II's MRI-compatibility was performed with the assistance of Anita Siebert, as was the adjustment of the EPI-QA1 sequence. Software to read out and control the Yocto-Meteo-v2 temperature sensors was created in cooperation with Anastasia Doinaki. The idea of measuring the data-loss using test-trigger injections was initially proposed by Prof. Peter Fischer, and implemented by the author. Lastly, the author supervised and either executed or assisted in large parts of the scanner's assembly process, including mechanical components, cabling, heat sinks, fibers, insulation and shielding.

3 System Control, Data Acquisition, and Raw Data Analysis

In the previous chapter, we have looked at all the components that comprise the SAFIR-II PET insert, as well as their functionality, connections, and the reasoning behind specific design choices. Missing from these discussions has been the functionality of a major and crucial component without which the scanner would be unable to start, much less acquire data. Any measurement operation by a person using SAFIR-II has to be started, controlled, and analyzed via the DAQ-PC. This workstation, located in a separate room outside the MRI's Faraday cage, receives data from and sends control signals to the insert via the 12 glass fiber Ethernet links connected to the 12 SDIP3 PCBs. It is also used to control the power supply units and analyze the collected raw data for coincidences using custom software.

To handle the significant processing demand arising from these tasks, the workstation features a 32-core central processing unit (CPU) (3.5GHz, AMD Ryzen Threadripper PRO 3975WX [56]) along with 256GB of DDR4 Memory (Kingston 9965694-026.B00G [57]). Three Quad Port 10Gbit SFP+ Ethernet network cards (Intel X710 [58]) connect the DAQ-PC to the SDIP3 PCBs, with three 7.68TB NVME solid-state drives (SSDs) (Samsung PM9A3 [59]) used to store data at a maximum write speed of 4 GB/s.

But even though the DAQ-PC is located in a basement room away from the MRI and its laboratory, an individual performing measurements using SAFIR-II does not have to be. Instead, the workstation can be controlled remotely via an SSH connection using a separate laptop (Lenovo ThinkPad E14 [60]). This way, the entire scanner can be started and stopped remotely as well.

3.1 Control Software

The software utilized to control SAFIR-II via the DAQ-PC closely matches what was used in the predecessor system SAFIR-I. For SAFIR-II, I have modified this software in several areas to account for differences in the scanner's design. While the implementation for SAFIR-I has been detailed extensively in [22], I wish to provide an overview of the various functionalities for ease of understanding, updated for SAFIR-II.

The control software consists of three server programs running in parallel on the DAQ-PC. Each server is responsible for controlling a specific part of SAFIR-II's operation:

1. Power supply control: Handling the output voltages and monitoring the current consumptions of the six power supply units.

3 System Control, Data Acquisition, and Raw Data Analysis

2. Ethernet control: Processing data traffic to and from the 10Gbit optical Ethernet links.
3. Bias Control: Used to control the applied bias voltages via the SBTV3 PCBs.

Communication between the server programs is possible via the Apache Thrift™ open-source protocol. Furthermore, three corresponding programs featuring graphical user interfaces (GUIs) allow a user to control the servers.

3.1.1 Power Supply Control

The six power supplies utilized for SAFIR-II's operation reside in the same room as the DAQ-PC and are connected to it through regular USB cables. The power supply control Server ('pscServer') program is utilized to send commands to the power supplies to activate their outputs, set their applied output voltages and read back the respective currents. The corresponding GUI allows the user to activate the output of units and observe the power supply status even when operating remotely.

The target output voltages for each power supply are stored in and read from a configuration file, preventing the user from applying a erroneous voltage by accident. Additionally, a polarity check routine is implemented, which has to be performed before full system activation. The routine temporarily sets all power supply outputs to a low voltage (1 V), observing the output current to check for electrical shorts, issuing an error and disabling full power-up of the scanner should such an irregularity be detected.

Powering of the scanner must be performed in a specific order, beginning with the 18 V main functional power, followed by the ± 12 V floating voltages for the SBTV3, and finally the 50 V bias shift voltage. As explained section 2.4.5, two power supplies are utilized for each voltage. For the main functional power, the software controls a strict powering sequence. As one of the two SPPD PCBs additionally supplies power to the SFCM, it is imperative that the power supply unit connected to it is started first, since the SFCM needs to be powered before or simultaneously with the SDIP3 for proper operation. Such a sequence is not required for the other two voltages, and they are enabled simultaneously.

An additional step involving the power supply control is not implemented in the software but rather externally through additional hardware depicted in figure 3.1. The power



Figure 3.1: Image of the power control panel.

plugs of all power supply units, as well as those of the cooling unit and the blower, are plugged into a singular, wall-mounted control panel. It features circuit breakers for each connection and can be controlled remotely to start or stop the cooling units and power supplies, both through a web interface as well as physical buttons located in the laboratory room outside the MRI's Faraday cage. In case of a failure of the cooling system, power to the power supplies is cut to avoid overheating of the scanner.

3.1.2 Ethernet Control

The Ethernet control server ('ecServer') software is responsible for managing the transfer of data to and from SAFIR-II via the 10Gbit optical Ethernet links. It receives raw data packages arriving from the insert, performs cyclic redundancy checks (CRCs), and saves the data to the hard drives during acquisition. Simultaneously, it is used to send control commands and configuration data, including packets to reset or power down the scanner, send test trigger signals, and initiate the PETA-8 ASIC readout. Similar to the power supply control, a corresponding GUI allows the user to interact with the server, instructing it to send the desired control signals or configuration.

Control Commands

Using the GUI, an operator can initiate the sending of several commands to one or all SDIP3 boards. They are processed by a MicroBlaze™ microprocessor implemented within the FPGA firmware, which then executes the desired functionality.

The following control commands are available:

- Enable/disable LDOs supplying analog power to a detector head module. During regular operation all LDOs are enabled, but an option exists to control each LDO individually. This way, individual detector head modules can be disabled for testing purposes.
- Request the sending of a fast control signal from the SFCM/SFCD (reset, test trigger, power down). An additional option exists to initiate a continuous sending of the test trigger signal at a set frequency of 200 Hz.
- Perform a single readout from all ASICs or initiate continuous readout operations. A single readout acquires data from all PETA-8 channels once. In continuous mode, a new readout operation is initiated immediately once the previous one is complete.
- Adjust the frequency of the ASIC readout in a range from 200 MHz to 300 MHz. Currently unused, but necessary should MRI sequences targeting something other than hydrogen be desired.

Beyond that, communication with the SBTv3 is handled through the Ethernet links as well. However, a separate additional program interfacing with the 'ecServer' is involved in this process, which will be discussed in detail in section 3.1.3.

ID	Register	Parameters	Length [bits]
06	Global Config	28	78
08	DAC settings	27	285
09	Temperature ADC	4	12
10	Neighbour logic	8	96
12-43	Per-channel	9	35

Table 3.1: Selected JTAG instructions used to configure PETA-8.

ASIC Configuration

Similarly to the control commands, configuration data for the PETA-8 ASIC is sent following user request and processed by the MicroBlaze™, which passes it to the ASICs themselves. The configuration data is stored within a file on the DAQ-PC and has to be loaded into the GUI manually before it can be sent. Within the GUI, each individual parameter can be adjusted as needed before sending it to the scanner. The sending of configuration data is done after powering up the scanner but before initiating the data readout and is not required more than once during regular operation.

The PETA-8 ASIC is configured via a Joint Action Test Group (JTAG) protocol (IEEE Std. 1149.1). All four ASICs on a given SP8 PCB are integrated into a single JTAG chain, allowing them to be configured simultaneously through the same interface. Additionally, the JTAG protocol enables the readback of registers, which can be used to examine whether the configuration was successful. The configuration data includes a total of 84 parameters, most of which are identical across all ASICs. These parameters are grouped into five separate JTAG instructions, detailed in table 3.1. Four instructions affect the entire ASIC, while the fifth is used to configure each channel individually. This per-channel configuration is further split into 32 instructions, one for each channel. Two parameters within the per-channel configuration are individually adjusted for each channel in order to adapt to manufacturing variations. However, while the adjustments for these parameters are performed by the Ethernet control software, they will be detailed later in section 3.2, as they are part of the scanner’s calibration.

PETA-8 features additional JTAG instructions with purposes not connected to the ASIC configuration. These include the readback of several device status parameters, an instruction to trigger individual channels similar to the test trigger, and a readback of the device ID. While functionality for these instructions exists in the Ethernet control software, it is not utilized during regular operation.

Data Acquisition

The data received from the SDIP3 boards can be divided into two types: slow control data, which includes temperature readout values from the PCB, and measurement data, which is mainly comprised of the hit data acquired by the ASICs. Both types of data are independently transmitted by the FPGA as User Datagram Protocol (UDP) Ethernet

Epoch Counter Packet

ID (0xFF)	Epoch Counter Value
39 – 32	31 – 0

Hit Packet

Chip ID	Channel ID	Test Hit?	TDC valid?	ADC value	fine counter	mid counter	coarse counter
39 – 34	33 – 29	28	27	26 – 18	17 – 13	12 – 10	9 – 0

Figure 3.2: Format of data blocks within the payload of a raw data packet received from SAFIR-II. For each readout operation, one epoch counter packet is saved along with a large number of hit packets depending on the number of hits detected.

packages and saved to disk in separate files for each of the 12 SDIP3 boards during data acquisition. The saved data is only comprised of the accumulated UDP package payload, not containing headers or CRCs.

Slow control data is continuously acquired following the power-up of the scanner and can be monitored by the user via a secondary GUI program. This data includes information from two temperature readout ICs (LTC2990 [61]) placed on the SDIP3 itself, as well as any response data from the SBTV3. While both are received and saved by the Ethernet control software, only the former is monitored during regular operation.

Most of the Ethernet packages received from the insert consist of measurement data, and its acquisition only occurs upon user request. Via the GUI, the user can opt to either perform a single readout operation, acquiring data from all PETA-8 channels once, or a continuous readout, meaning data is acquired constantly and as quickly as possible until a stop is issued. The payload of Ethernet packages containing measurement data is comprised of a large number of 40-bit data strings. Most of these strings are hit packets, which contain data acquired by the PETA-8 channels for detected hits. The first 11 bits of a hit packet are used to indicate which channel of which PETA-8 received the hit using the 'chip ID' and 'channel ID'. The remaining 29 bits contain information about the hit, as shown in figure 3.2. If the first byte of a 40-bit string is set to a specific hexadecimal value (0xFF), this indicates that the remaining 32 bits instead contain an updated value of the 'epoch counter'. This counter, as mentioned in 2.4.1, is implemented in the FPGA firmware and used to extend the length of the coarse counter. It increments synchronously with the overflow of the coarse counter and is added to the measurement data packages between readout operations.

All twelve files containing measurement data are saved on the same SSD. Should the measurement exceed the available storage space on a given drive, the software automatically switches to the next drive to keep the acquisition running. An option exists to split the files across all three drives, with each drive saving the data from four links. This can be used in the case of high-activity acquisitions where an individual drive's write speed

of 4 GB/s is not sufficient to handle the data load of all twelve links simultaneously. For measurement activities of up to 500 MBq and an energy cut of 70 LSB this is not needed. However, it might be necessary for even higher activities or lower energy cuts when compared with the results of the previously detailed calibration phantom measurement in figure 2.11 of section 2.3.3.

In addition to handling data storage and display, the Ethernet control software is used to perform basic analysis tasks of the raw data during acquisition. This includes monitoring the incoming data for potentially lost packages, and it also features the ability to evaluate the raw ADC values of events acquired by the PETA-8 ASIC channels.

3.1.3 Bias Control

Once correctly configured, all ASIC channels acquire data from 511 keV photons with approximately constant ADC values. However, this mandates the stability of the SiPM gain, which is significantly temperature-dependent. To be precise, the SiPM breakdown voltage changes at the previously mentioned coefficient of $54 \text{ mV}/^\circ\text{C}$, which effectively changes the OV if the applied bias voltage stays the same. As an increased APD current, caused by higher measurement activities, will inevitably heat the SiPM, the resulting change in OV will impact the SiPM gain and, therefore, the mapping of photon energies to ADC values. Likewise, the other SiPM performance parameters introduced in 2.2.4, such as the PDE and dark count rate, will change, which would ultimately result in the degradation of the scanner's timing resolution. For these reasons, changes in the SiPM bias voltage are continuously accounted for by adjusting the applied bias voltages via the SBTV3.

The SBTV3 is controlled indirectly via Ethernet packages sent to the SDIP3. Similar to the control commands and configuration data, such packages are first received by the MicroBlaze™. The relevant information, detailing what voltage should be applied to which SP8 module, is then passed to the SBTV3 via a connecting cable. However, in this case, the Ethernet control server is only used to pass these packets to the SDIP3, while a secondary program dictates their contents.

The bias control server ('bcServer') interfaces with the Ethernet control software via the Apache Thrift™ open-source protocol. Along with its corresponding GUI, it is used to communicate with the SBTV3, applying or canceling the bias voltage. The breakdown voltages of each detector head module are saved in a file to disk, with the software determining the voltages to be applied by adding the user-defined OV. In order to adjust for temperature-dependent changes of the SiPM gain, the Ethernet control software continuously monitors the average position of the photopeak within the ADC value spectrum of each detector head module. This is compared to the average peak positions at low measurement activities, which are saved for all detector head modules in a file as well. Should a deviation be observed in the raw data, the bias control software will gradually adjust the applied voltages until the observed raw data aligns with the saved values again. The speed of this adjustment is dependent on the rate of detected hits, with higher rates and thus measurement activities seeing more frequent updates.

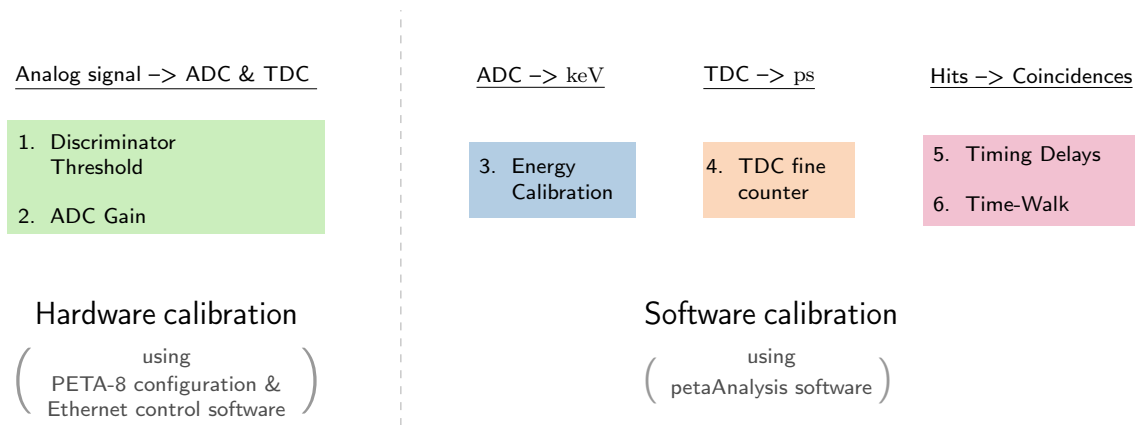


Figure 3.3: Overview of the order of calibration procedures performed for SAFIR-II.

3.2 Data Analysis and Calibration

The SAFIR-II insert is designed to acquire and digitize raw binary hit data created from scintillation processes before transmitting them to the DAQ-PC. During a measurement, the locations, ADC, and TDC values of these hits are recorded on an SSD by the Ethernet control software. Following this, the data needs to be converted into keV energies and ps time stamps before finally being filtered for coincidence pairs. These tasks are handled by the ‘petaAnalysis’ software, which is executed following a completed data acquisition. Like the previously discussed control software, petaAnalysis was initially developed for SAFIR-I [22] and I have modified it to be able to process SAFIR-II data.

Various calibration procedures are required to convert a given set of raw data accurately. These calibrations are carried out in a specific order, as detailed in figure 3.3. While the first two are performed using the Ethernet control software, they are elaborated in this section for ease of understanding. All remaining calibrations are being conducted by the petaAnalysis software. The following section gives an overview over the various calibration procedures and steps of the analysis.

3.2.1 Energy Filtering

Not all hits acquired by the PETA-8 ASICs are cleanly detected 511 keV photons stemming from an annihilation process. As mentioned before (section 2.2.2), intrinsic radiation of the utilized LYSO crystals can trigger a channel. Photons can also interact with matter before reaching a crystal, thereby losing energy and deviating from their original path. While the ‘Energy Cut’ configuration parameter of PETA-8 can technically be used to remove hits below a specific energy, isotopes such as ^{22}Na can possess additional decay modes with energies higher than 511 keV. Additionally, effects like pile-up can also result in detected hits with higher energies, which must be considered invalid for coincidence pair matching. For these reasons, the analysis software is required to first convert the arriving ADC values to keV energies. A cut on the data is performed afterwards, allowing




Figure 3.4: Example energy spectrum of a ^{22}Na source, averaged across all channels of one SP8 module. The red lines indicate the energy window of 400 keV to 600 keV. The data has undergone the energy calibration of the `petaAnalysis` software. Said calibration seems to be accurate for lower energies, such as the 511 keV photopeak, but less so for the 1275 keV line of ^{22}Na .

only hits within a specified energy window to be considered for coincidence sorting.

For SAFIR-II, this energy window ranges from 400 keV to 600 keV, as depicted in figure 3.4. Considering the previously discussed energy resolution of 12.1% FWHM, about 3σ of the peak are included within this range. The K- α line at ≈ 450 keV is included in the window as well, while the lower energy cutoff is still above the Compton edge.

To convert the raw ADC values accurately, a calibration is required. This calibration is performed in two steps, first at the hardware level via configuration of the PETA-8 ASIC, and afterwards at the software level by the ‘`petaAnalysis`’ software.

Channel Gain of the PETA8 Front End

It is desirable to have signals of similar energy result in similar ADC values across all channels. This approach ensures that the channels will exhibit a consistent performance when detecting hits. Specifically, as the channel’s front end determines the ADC value for a hit by linearly discharging a capacitor, said value dictates the required processing time. On the other hand, smaller ADC values for a given hit impact the scanner’s energy resolution, as overall hits with different energies become less distinguishable.

The ADC gain of the PETA-8 front end can be individually configured for each channel. This can be used to compensate for variations in the SiPM and ASIC channel gain caused by manufacturing tolerances, thereby ensuring that hits with similar energies result in

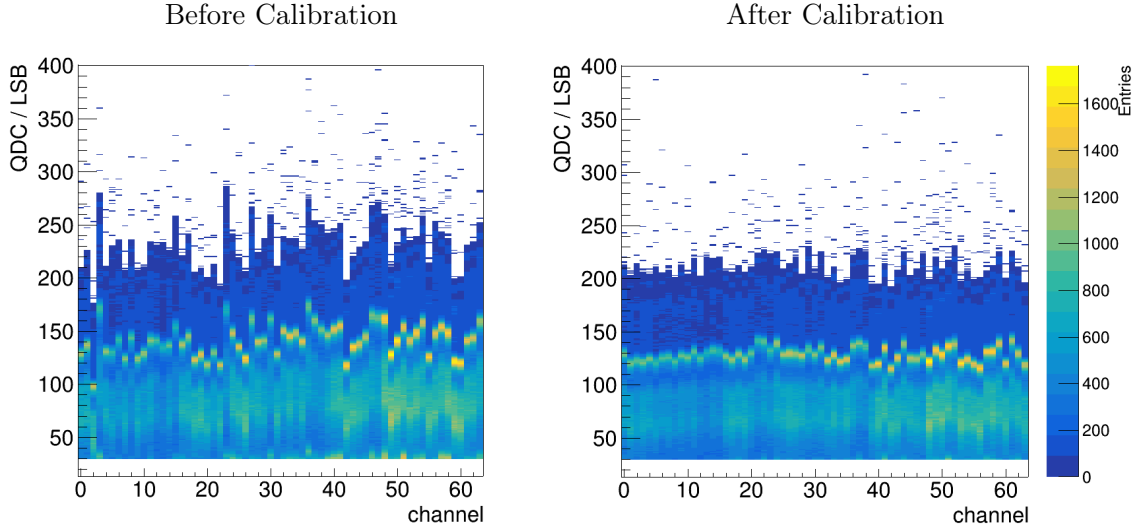


Figure 3.5: Two-dimensional histogram of raw data QDC(ADC) values measured using a ^{22}Na point source before and after calibration of the ADC gain of the PETA-8 front end. Observable before calibration is the photopeak of each channel varying between values of 95 LSB to 175 LSB while showing a much smaller spread afterwards.

similar ADC values. These gain values are not calibrated by using the `petaAnalysis` software but instead via the previously discussed Ethernet control software.

The calibration routine is performed using a ^{22}Na point source placed within the scanner. It functions by incrementally raising or lowering the ADC gain configuration value of each channel until the spectra of all channels roughly align. At a given bias voltage, which is held constant and not automatically adjusted, data from all channels is acquired for 20 s. Afterwards, the data is analyzed to find the ADC value of the highest energy peak in every channel's individual spectrum. As the measured source is made of ^{22}Na , this peak corresponds to the 1275 keV emission line. Its value is compared to a reference identical to all channels (200 LSB); the gain is increased if it is lower and reduced if it is higher. Each adjustment of the gain configuration value occurs in steps of only 1 LSB, and the procedure is repeated ten times. Following this calibration, the 511 keV peak is located at around 120 LSB for each channel's spectrum. Figure 3.5 shows the spectra for several channels before and after this calibration.

The manufacturing variations in each channel's gain are obviously hardware-dependent. As such, the calibration of the ADC gain values needs to be done only once for any given hardware and does not have to be repeated unless a component, such as a detector head module, is exchanged.

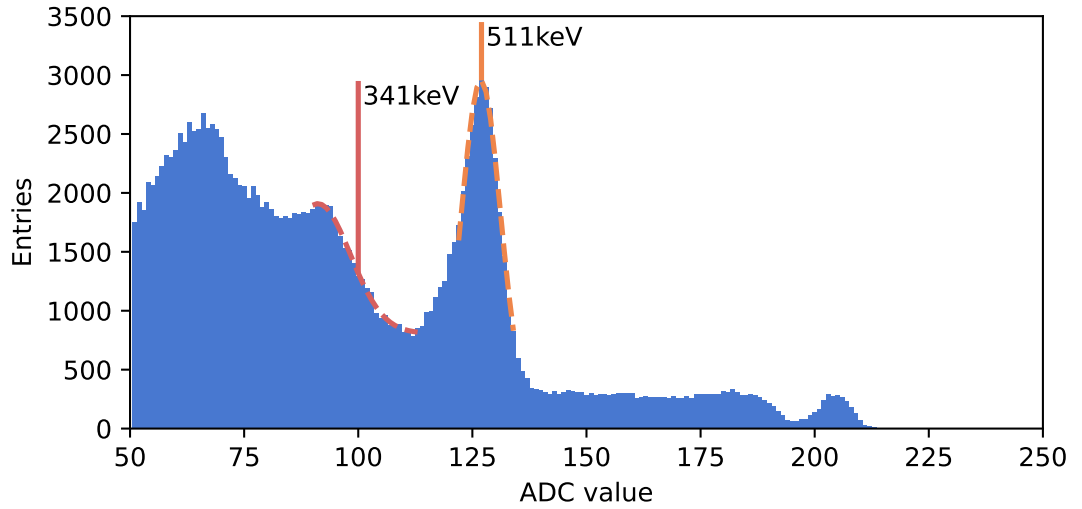


Figure 3.6: Typical energy spectrum observed by a PETA-8 channel during the calibration measurement of a ^{22}Na point source. For the energy calibration the ADC value of the photopeak is determined via a Gaussian fit. The same is done for the Compton edge using half of a Gaussian.

Analysis Energy Calibration

The `petaAnalysis` software utilizes the following formula to convert raw ADC values Q_γ to keV energies E_γ :

$$E_\gamma = a_i \cdot (Q_\gamma - b_i)^{c_i}. \quad (3.1)$$

b_i accounts for the ADC offset of each channel caused by thermal noise and other effects. The parameter a_i reflects the linear dependency of Q_γ and E_γ , while c_i models the SiPM's exponential saturation behavior.

While the raw ADC spectra and 511 keV peak positions for each channel are roughly similar following the calibration of the front end gain, they are by no means identical. Each channel is influenced by variations not only in the SiPM gain but also in the ASIC analog front end manufacturing and the optical coupling of crystals and SiPM pixels. Therefore, the `petaAnalysis` software calibrates the parameters a_i , b_i , and c_i individually for each channel.

b_i is determined via a calibration measurement using PETA-8's test trigger feature. The measurement involves disabling the bias voltage, sending 10.000 test trigger signals, and determining the average ADC value of the received hits for each channel. By switching off the bias voltage, a channel's ADC value is equal to the offset b_i . Therefore, sending a test trigger signal and reading out the data can be used to measure this value.

The parameters a_i and c_i are determined afterwards in a separate measurement using the ^{22}Na point source, by obtaining reference energies from the energy spectrum of each channel. For this, the source is placed within the center of the FOV and measured for a

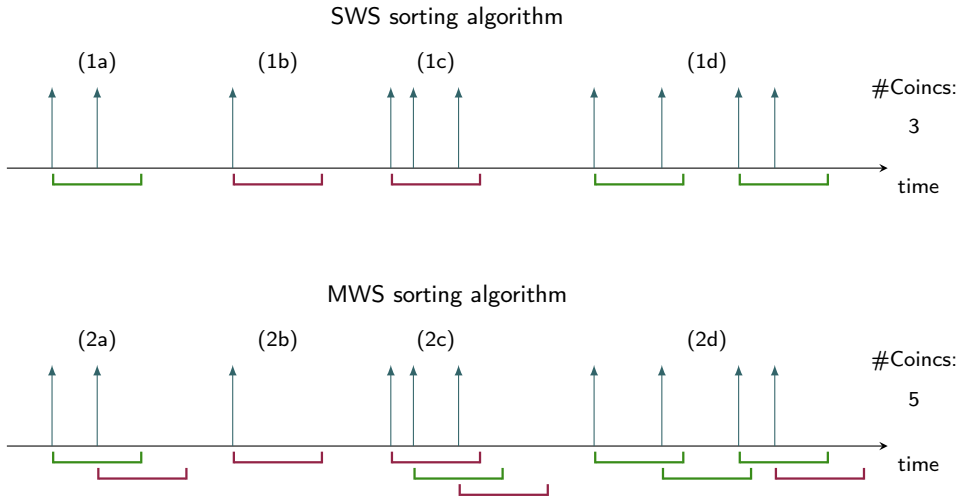


Figure 3.7: Schematic overview of the SWS and MWS algorithms. Each arrow represents a valid single, while the horizontal brackets denote the algorithm’s timing window. A green bracket signals a window containing a coincidence according to the algorithm, while a red bracket does not. One can observe the SWS algorithm rejecting a triple hit (1c), while the MWS algorithm accepts two of the three singles as part of a coincidence (2c), though without guarantee that this selection is correct.

total of 15 min to gather sufficient statistics. The 511 keV photopeak is determined using a Gaussian fit to the energy spectrum, as shown in figure 3.6. Similarly, the Compton edge at 341 keV is fitted using the right half of a Gaussian. These reference energies and their ADC values can then be used to determine the two values a_i and c_i .

3.2.2 Coincidence Sorting

Once the available raw data has been calibrated and filtered for valid 511 keV hits (‘singles’), it can be sorted for coincidences. Intuitively, it makes sense to accept two singles with similar arrival times as two parts of a coincidence. The single-window sorting (SWS) and multi-window sorting (MWS) algorithms are two common approaches to this idea [62], and their functionality is rather similar to each other. Given a set of data, the singles are first sorted in time. Afterwards, the data is examined sequentially as the algorithms compare the time differences between subsequent singles.

A sketch visualizing the two algorithms is shown in figure 3.7. For the SWS algorithm, a given timing ‘window’ is initially applied to the first detected single. If the time difference between this and the next single is larger than the timing window, the first is rejected, and the process is repeated. If the time difference is shorter than the timing window, the two singles are considered part of a coincidence, and the window is moved to the next single after them. Should more than one single fall into the window, all hits are

discarded, as coincidences between more than two singles are invalid, and it is impossible to determine which two singles might form a pair. MWS opens a timing window for every detected single, even if said single is already within the window of a previous hit. As such, windows can overlap, potentially creating two coincidences for three detected singles.

Both methods have potential upsides and downsides. For an event with three singles, SWS might miss the ‘true’ coincidence, while MWS would detect both it and a ‘false’ coincidence. The `petaAnalysis` software used for SAFIR-II employs SWS, as previous studies using the DRP showed the choice of algorithm to have had little impact on the scanner’s performance [22].

It should be noted that geometric effects have to be considered for coincidence detection as well. While a coincidence between two crystals on the same detector head module can potentially be detected, the resulting LOR will not pass through the scanner’s FOV, much less the subject under study. As such, the analysis software will only accept a coincidence if the two crystals detecting it are separated at a tangential angle of 90° or more.

3.2.3 Timing Calibration

An obviously important parameter influencing the effectiveness of the SWS algorithm is the length of the applied timing window. Photons created by an annihilation process are not necessarily detected at the exact same time, as one photon might travel longer than the other before reaching a crystal, depending on the decay’s position. As such, a window that is too small could result in valid coincidences not being detected, resulting in a loss of sensitivity for the scanner. Inversely, the larger the window is, the higher the probability of two unrelated photons being falsely detected as a coincidence or a valid coincidence being rejected due to a third single falling into the window. In order to reasonably minimize the length of the timing window, it is imperative that the TDC values of each single reflect the arrival time as accurately as possible. However, hardware-specific effects can deteriorate this accuracy, and therefore, the `petaAnalysis` software features several calibration procedures to correct each hit’s TDC value.

TDC Conversion

In basic principle, the PETA-8 ASIC stores a hit’s timing information as a single, continuous 18 bit TDC value with a bin width of 50 ps. However, as mentioned in section 2.3.1, it is actually comprised of three separate counters combined, which differ in implementation.

The 5 bit long ‘fine’ counter consists of the converted output of a voltage controlled ring oscillator. Said ring oscillator features delay elements, which, due to manufacturing variations, result in the individual bins of the fine counter not being perfectly equal in width. In order to accurately convert the TDC values to ps timestamps, the width of each fine counter bin needs to be determined for each channel. This is done using the same point source calibration measurement also used for the energy calibration. As radioactive




Figure 3.8: Coincidence timing distribution of the SAFIR-II insert determined using a ^{22}Na point source after calibrating for the fine counter bin width.

decay is a random process, one would expect a uniform distribution of all fine counter values if they were of equal bin width. Therefore, by observing the actual distribution and accounting for the fixed fine-counter wraparound period of 1.6 ns, the bin widths can be determined.

In addition to the fine counter, a 3-bit ‘mid’ counter is utilized, incrementing once every 1.6 ns along with each wraparound of the fine counter. Likewise, the 10-bit ‘coarse’ counter increments with each wraparound of the mid counter. Both of these counters are implemented twice for a given ASIC. Since hits arrive at random, the possibility exists for a counter to be transitioning between values as a channel is being triggered, which would result in an invalid TDC value being saved for the hit. Therefore, the two instances of each counter are phase-shifted by half the bin width, assuring that, for any given moment, at least one counter displays a stable value. Which of the two instances is selected for a given hit depends on the value of the preceding counter, i.e., the selected coarse counter depends on the mid counter value, and the selected mid counter depends on the fine counter value. The counter values at which a selection is swapped are selected within the ASIC’s configuration and need no further calibration.

Following the calibration measurement of the fine counter bin width, the analysis software is capable of converting binary TDC values into ps timestamps. Sorting the data for coincidences yields a broad, non-Gaussian distribution, as displayed in figure 3.8. This is caused by timing delays, which can result in phase shifts between different ASICs and channels.

Correcting for Timing Delays

The SFCM and SFCD PCBs are designed to distribute the reset and clock signals to all boards with minimal jitter. This is achieved using ICs such as the ‘IDT8P34S1208I’ [44],

which feature jitter below 100 fs. However, these signals are subjected to propagation delays before they reach a PETA-8 ASIC, depending on the length of the signal path. As the lengths of these paths vary between different ASICs, the respective TDCs operate at a fixed phase shift to one another. Adding to this are variations in the length of the analog signal lines between individual SiPM and ASIC channels, which delay the arrival time of an input pulse. Various other minor effects affect this as well, such as optical coupling between the crystals and SiPMs, slight differences in the breakdown voltages of each pixel, and other manufacturing variations between boards. All these effects ultimately result in timing delays between channel pairs, which need to be calibrated to ensure accurate timing information.

The delay calibration routine of the `petaAnalysis` software uses the same data as the previously discussed energy and timing calibrations, measured using a ^{22}Na point source. By placing this source within the axial and transaxial center of the FOV, the created annihilation photons are bound to interact with crystals on opposite ends of the scanner at the same time. Sorting the data between these detector pairs for coincidences and histogramming the differences in arrival time for each coincidence will yield a Gaussian distribution, with the Gaussian mean being equal to the difference $\Delta_{i,j}$ of the two detector channels' timing delays. Given the individual, unknown timing delays d_i for each of the 11 520 channels, a linear system of equations can be formed using

$$\Delta_{i,j} = d_i - d_j. \quad (3.2)$$

The various $\Delta_{i,j}$ are determinable from the coincidence timing distribution of detector pairs with a sufficiently large number of coincidences. For SAFIR-II's analysis, only detector pairs with 50 coincidences or more are considered, although this value is adjustable if desired.

If the correct timing delay values are known, the quadratic sum

$$x = \sum_i \sum_j (\Delta_{i,j} + d_j - d_i)^2 \quad (3.3)$$

will result in $x = 0$. The '`petaAnalysis`' software utilizes a minimizer algorithm to determine the values for d_i , and in any subsequent measurement, all timestamps recorded in data are adjusted to address these delays. This calibration is repeated thrice on the same dataset, grouping channels in the first two repetitions to accelerate the algorithm's convergence.

In the first calibration step, the data is analyzed with a timing window of 5 ns. All channels of a given ASIC are grouped together before histogramming, resulting in the calibration routine determining the average offsets between TDC pairs. These TDC delay values are taken into account in the second calibration step, which uses the same timing window of 5 ns. This step calibrates systematic offsets within the SP8 modules, as the layout of signal paths between the ASIC channels and SiPM pixels is identical for all SP8. Lastly, the routine is performed for the individual channels with a smaller timing window of 1.5 ns, incorporating the previously determined delays.

Following the determination of all delays present within the scanner, the calibration values are saved within a file. All coincidence sorting performed afterwards incorporates




Figure 3.9: Coincidence timing distribution of the SAFIR-II insert determined using a ^{22}Na point source after calibration of the timing delays. The red line indicates a Gaussian fit applied to the peak, which was used to determine the timing resolution as the FWHM of the peak. A timing resolution of 221.2 ps was observed.

these values, and uses a timing window of 500 ps. Figure 3.9 displays a histogram of the timing distribution of the SAFIR-II insert following these calibrations.

While, in principle, these calibration values can be reused for each measurement, an additional non-constant delay source exists, which needs to be considered. The TDC counter value for each ASIC is set to zero following a reset signal. But as PETA-8 samples the reset asynchronously, the TDC counter values being set to zero might not occur on the same clock cycle for every ASIC. This results in a potential phase shift of one clock period between different ASICs, equal to 1.6 ns. Thus, the first step of the timing calibration needs to be redone following a reset, which is performed after each power cycle of the scanner.

Sparse linear solvers

For SAFIR-I, the algorithm used to minimize the quadratic sum in equation 3.3 was the TMINUIT-MIGRAD minimizer implemented within the ROOT analysis framework [63]. MIGRAD is based on a variable metric method [64] and, among other things, operates by repeatedly approximating the Hessian matrix for the function to minimize [65]. This minimizer package is frequently employed in particle physics data analysis and worked sufficiently fast for its purpose in SAFIR-I and the DRP, completing the calibration within about 10 minutes.

A careful reader proficient in computational arithmetics might however have already caught on to a problem with this approach with respect to SAFIR-II. The fourfold increase in the number of crystals for SAFIR-II results in a sixteen-fold increase in the

number of channel pairs and parameters, which would significantly increase the required computation time. Furthermore, this approach to the minimization is highly inefficient, as the linear system of equations formulated in equation 3.2 is rather sparsely populated. This sparsity is caused by the central position of the point source since only crystals in positions roughly opposite each other actually contribute. Indeed, initial tests with SAFIR-II showed a minimum computation time of four hours just for the last step of the calibration procedure, which, depending on the conversion speed of the minimizer, even extended up to a full 24 hours in some cases. As the calibration of timing delays needs to be performed preceding any measurement following a power cycle (or reset) of the scanner, it is apparent that an alternative solution was required.

As mentioned, the linear system of equations arising from 3.2 is sparsely populated while featuring a large number of parameters d_i , meaning a matrix of the observed timing differences $\Delta_{i,j}$ features a significantly larger number of zero elements than non-zero elements. The MIGRAD minimizer processes this system as if it were densely populated, resulting in a large amount of computational resources being spent on zero elements. A minimizer optimized for solving sparse systems would promise significant improvement to this problem. Therefore, I adjusted the calibration routine for SAFIR-II by using the ALGLIB open-source library [66]. More specifically, an implementation of the sparse LSQR algorithm [67] is utilized, which solves

$$\min \|Ax - b\|_2 \tag{3.4}$$

for x , where x is the vector of individual channel delays d_i , b contains the timing differences $\Delta_{i,j}$, and A is an $m \times n$ matrix with elements $a_{i,j} \in [-1, +1]$. Tests of the new algorithm proved its suitability, performing the calibration routine in less than 10 s.

3.2.4 Threshold of the PETA-8 Front End

In addition to the gain of the PETA-8 front end, a second configuration parameter has to be individually calibrated for all channels. The value of the TDC comparator threshold is crucial to the timing performance of the scanner. Ideally, this value would be set as low as possible, ensuring that an analog signal from the SiPM is detected as early and accurately as possible. However, due to the hardware implementation, threshold settings below a specific value will result in the channel not triggering at all. Additionally, for a certain value range, thermal noise within the channel itself will trigger the comparator regardless of whether a signal has arrived, rendering the channel useless. Due to manufacturing variations, the exact value ranges for these effects vary from channel to channel and thus need to be individually determined. They are, however, approximately constant for a given hardware. As such, these values only need to be determined once following the scanner's assembly and can be used for all subsequent measurements as long as no detector head module is exchanged.

The base values for noise and comparator triggering are determined via a calibration routine embedded in the DAQ-Software. Disabling the bias voltages usually supplied to the SiPMs ensures that a channel only acquires data when triggered by thermal noise. Simultaneously, by setting the energy cut to zero, a triggered channel can easily be




Figure 3.10: Timing resolution of the SAFIR-II insert at various measurement activities up to 500 MBq, measured for six different values of the relative TDC comparator threshold.

detected, as it is sure to contain a data packet to read out. Applying a threshold value, waiting for a short time, and then reading out data from all triggered channels can thus be used to determine if a given value falls within the noise range. To determine the level at which a channel no longer triggers due to thermal noise, a controlled sweep is performed across threshold values within a predetermined range (900 LSB to 1300 LSB). Afterwards, the software records the range of values at which noise triggers occur for each channel within a configuration file on the DAQ-PC.

However, the SiPMs themselves are also subject to thermal noise as well as optical and electrical crosstalk. Therefore, the chosen value for the discriminator threshold is not directly above the base noise level. Instead, the median value of the noise range is combined with a ‘relative’ threshold to determine the value actually applied in the configuration. This relative threshold is identical for all channels and has a significant impact on the scanner’s performance, with the optimal value depending on the measurement activity.

To examine this, I performed a measurement using a thin cylindrical line source filled with up to 500 MBq ^{18}F solute in water. The performed measurement is the same as mentioned in section 2.3.3, though data was recorded six times for each measurement activity, with a different relative threshold applied each time. Figure 3.10 shows the observed values for the scanner’s timing resolution during this measurement. It can be clearly observed that the lowest selected threshold value of 100 LSB yields optimal performance for lower activities, with the timing resolution degrading for higher thresholds. However, this situation quickly changes for higher activities, with a threshold of 140 LSB to 160 LSB yielding the best performance, while both lower and higher thresholds end up as suboptimal.

Another aspect to consider in this is the scanner’s energy resolution, as the comparator threshold is responsible for triggering the integrator and can thus influence the energy




Figure 3.11: Energy resolution of the SAFIR-II insert at various measurement activities up to 500 MBq, measured for six different values of the relative TDC comparator threshold.

resolution. However, figure 3.11 shows that SAFIR-II’s energy resolution is relatively stable across its targeted activity range, degrading by less than 10 %. While the threshold value seems to have a slight impact on the energy resolution, especially at higher activities, the deviations are less than 0.1 %. This exemplifies a significant difference between SAFIR-II and its predecessor, which observed a significant degradation in the energy resolution at higher activities [24]. This change can be attributed to the improved separation layers between crystals in SAFIR-II’s detector head, as less optical crosstalk is able to contaminate signals at high activities. A second factor could be the improved cooling utilized in SAFIR-II, which might significantly reduce thermal noise in both the ASICs and the SiPM. Given the observed results, I concluded that a threshold value in the middle of the selected range would be the most suitable for the activities SAFIR-II is targeting. Subsequent measurements mentioned in this thesis have been performed at a relative comparator threshold value of 150 LSB.

Time-walk

The comparator threshold also gives rise to an additional systematic effect concerning the timing performance. While its constant value results in a channel triggering at the same amplitude for all signals, it has to be noted that the time a signal requires to reach said threshold amplitude, or the Time to Threshold (TT), is dependent on the signal’s total amplitude. The TT for a high amplitude signal is considerably shorter than for a low amplitude signal, leading to a systematic shift in the time-stamp referred to as the time-walk. Figure 3.12 shows an illustration of this principle.

For this purpose, the analysis software features an additional time-walk calibration routine. It operates by observing the different energies between two hits of a coincidence, as well as the average timing difference for coincidences with similar energy differences.




Figure 3.12: Schematic visualization of the principle behind time-walk. Due to the different amplitudes of the two signals, the comparator threshold gets triggered later for the smaller signal.

To adjust for the time-walk, a quadratic function of the form $\Delta t = p_0(\Delta E)^2 + p_1$ is fitted to these energy differences, with the value of p_0 being saved to calibrate future measurements.

The observed gain in timing resolution obtained from using the time-walk calibration is relatively small; for SAFIR-II, an improvement of about 4 ps was observed. This can be attributed to the fact that the energy window of 400 keV to 600 keV is chosen to be relatively small, as the effect of time-walk on the timing resolution is more substantial for larger energy differences. However, its correction becomes more significant as soon as hits of lower energies are also considered.

3.2.5 Inter Crystal Scatter Recovery

The coincidence sorting process of the `petaAnalysis` software described so far has focused on hits resulting from photons depositing their entire energy within a single crystal, exemplified by the utilized energy window of 400 keV to 600 keV. The reason for this is apparent, since 511 keV photons undergoing a Compton scattering process might deposit only a fraction of their energy. It is not inherently possible to distinguish such hits from other background radiation. However, there exists the possibility of a Compton-scattered photon interacting a second time, depositing its remaining energy in a crystal close to the initial scattering point. This makes it possible to register the two received hits as the result of a single 511 keV, thereby recovering the original single. In the technique known as inter-crystal scatter recovery (ICSR), the analysis software considers hits with lower energies that were detected within a short time span of each other in a set maximum distance. If their energies sum up to ≈ 511 keV, it can be concluded that they originated from a single annihilation photon. This provides additional singles for the coincidence sorting algorithm, ultimately boosting the system sensitivity.

Concerning the software implementation of the ICSR algorithm, several parameters need to be considered to optimize the scanner's performance. As Inter-crystal scatter results in a single ≈ 511 keV depositing its energy in different crystals, it is unclear which crystal should be considered the initial interaction point. Furthermore, the various hits are bound to be assigned different timestamps. Therefore, the combined hits arrival time

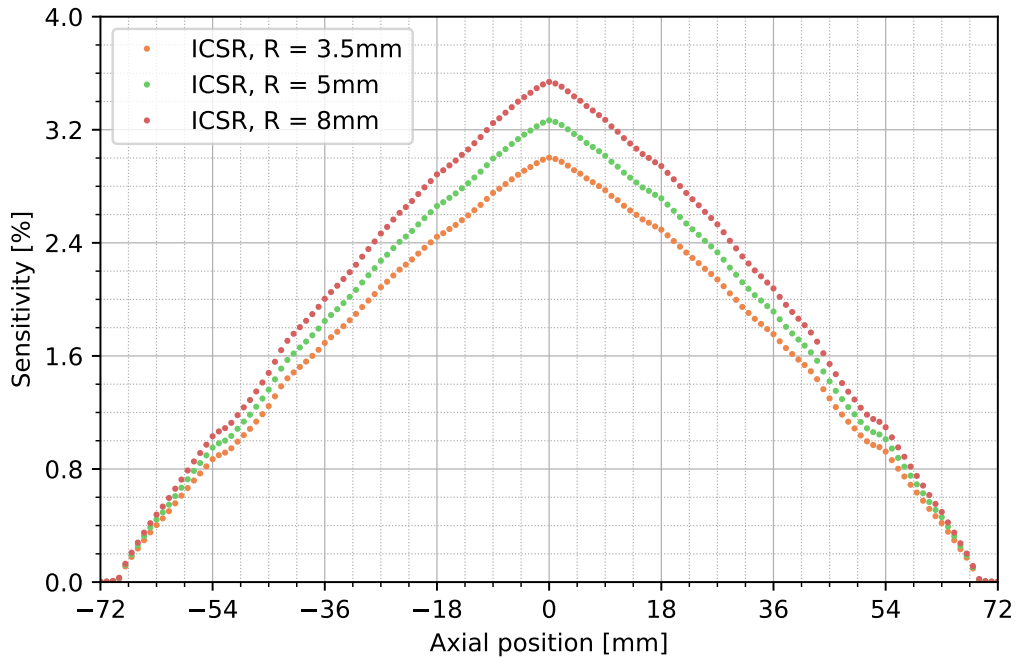


Figure 3.13: Sensitivity profile of the SAFIR-II detector, after utilizing ICSR at various recovery ranges to recover additional singles.

has to be determined. Lastly, the maximum distance between hits considered by the algorithm, known as the recovery range, is an important parameter that can affect the system sensitivity and degrade its spatial resolution. The same is true for the maximum number of hits combined for a recovered coincidence, as a single photon might scatter multiple times, or both photons might undergo scattering.

The `petaAnalysis` software implementation considers the hit with the highest energy as the initial interaction point. To generate a timestamp t_r for the combined hit, the timestamps are weighted by each hit's respective energy following the formula $t_r = \frac{t_1 E_1 + t_2 E_2}{E_1 + E_2}$, with the timestamp of the secondary interaction additionally being corrected for the travel distance between the two crystals. These parameters were determined empirically via measurements of the SAFIR-Collaboration using the DRP scanner and have been detailed more extensively in [22]. The maximum distance between scatter hits and the maximum number of hits to be considered as part of a scattered coincidence have also been determined previously. Measurements with SAFIR-I utilizing ICSR featured a recovery range of 5 mm and considered a maximum of three hits per coincidence [24]. However, a re-evaluation of these parameters for SAFIR-II is of significant interest, as the adjusted crystal geometry might have influenced their effect on the scanner's performance. Since the individual crystals are smaller, a given photon travels through less scintillating material per unit length and might thus travel farther overall. Therefore, I reanalyzed the data of the sensitivity measurement previously discussed in section 2.2.3 using the ICSR algorithm, considering three events per coincidence at various recovery




Figure 3.14: Schematic illustration of the method utilized to determine the hardware-based spatial resolution used to evaluate the ICSR impact. The LORs *i)* and *ii)* pass through the middle of the FOV, while *iii)* does not, giving rise to a distribution along the central axis. The FWHM of this distribution is quoted as the ‘axial resolution’.

ranges. Figure 3.13 displays the observed sensitivity profiles, showing a significant gain in the overall system sensitivity when compared to the previously measured peak of 2.23 % in section 2.2.3.

While the gain in sensitivity by using ICSR is bound to be beneficial, its effect on the spatial resolution of the scanner has to be gauged as well. However, the spatial resolution of a PET scanner is significantly dependent on the reconstruction algorithm used to obtain images, which will be elaborated in more detail in chapter 4. In order to evaluate the spatial resolution degradation of the ICSR on a hardware level without being influenced by reconstruction-specific parameters, the average axial position of each detected coincidence can be used instead. For a point source placed within the center of the FOV, the middle between the two interaction points of a coincidence’s photons should, on average, be in the FOV center as well. Figure 3.14 shows a schematic illustrating the principle along the axial direction, evaluating the spatial resolution via the position of hits along the detector rings. However, due to effects influencing the scanner’s spatial resolution, like larger recovery ranges, one instead observes a distribution around the center instead. I determined the effect of the ICSR algorithm on the scanner’s spatial resolution by observing the width of this distribution along the axial direction.

Table 3.2 compares the (peak) sensitivity values obtained for a point source measured within the center of the scanner’s FOV with the evaluated spatial resolution. Comparisons were made with the previously selected recovery ranges, while also considering either three or four hits per scattered coincidence. A surprising find is the fact that the number of hits considered per coincidence seems to affect the spatial resolution much less than in the previous scanner. This might be attributed to the improved optical separation between crystals for SAFIR-II, which has been shown to improve the scanner’s

Recovery Range	None	3.5 mm		5 mm		8 mm	
No. of Hits	2	3	4	3	4	3	4
Sensitivity [%]	2.23	3.00	3.16	3.27	3.51	3.54	3.89
Axial Resolution [mm]	3.35	3.38	3.39	3.47	3.50	3.61	3.68

Table 3.2: Peak system sensitivity and axial resolution FWHM for various settings of the ICSR algorithm. The first column details a measurement with ICSR disabled.

energy resolution by reducing optical crosstalk and might thus make the algorithm less susceptible to noise events. Nevertheless, an increased recovery range degrades the spatial resolution, and thus, an experimenter needs to consider the tradeoff between sensitivity and spatial resolution when utilizing ICSR.

The sensitivity gain obtained through using ICSR would seem to suggest that the scanner should generally operate with some form of this option enabled. However, another factor in the form of the measurements required storage space needs to be taken into consideration before making such a decision. Without ICSR, the PETA-8 ASIC’s Energy Cut configuration can be selected at a value of 70 LSB, reducing the amount of data transferred from the insert to the PC. For ICSR to work, this value needs to be lowered significantly, preferably to 40 LSB or even 30 LSB, to increase the recovered number of singles. But while the three high-speed SSDs of the DAQ-PC might be capable of handling the incoming data bandwidth, their storage space is limited. Studies planning to perform several high-activity acquisitions over longer time periods have the potential to approach or exceed this storage space, and thus, users should take this possibility into account when planning their measurements.

Statement of personal contribution: *All measurements and results shown in this chapter have been performed and evaluated by the author. The DAQ control and `petaAnalysis` softwares were originally developed by C.Ritzer. All adjustments and modifications to the software required for the compatibility to SAFIR-II, such as the adjusted configuration scheme, changed data format, implementation of sparse linear solvers, and inclusion of additional power supplies and ASICs, were either done or assisted by the author. Additional contributions include the setup of the DAQ-PC as well as the selection of its parts.*

4 From Coincidences to Images

The specific design of every component influences the performance of SAFIR-II. The size and arrangement of the scintillation crystals dictate how effective the scanner is at detecting events, the electronics play an essential part in gathering digitized information, and the analysis software is responsible for extracting as many coincidences as possible without introducing noise. Another crucial aspect in this regard is the reconstruction of actual images from the obtained coincidences. The field of PET image reconstruction has brought forth many algorithms, such as the early analytic methods, modern expectation-maximization methods, and, more recently, neural network-based approaches. These various techniques and approaches are able to significantly influence the resulting image even when using the same underlying coincidence data, as some do not consider the statistical nature of PET data, while others exhibit varying levels of noise depending on specific hyperparameters [68].

The image reconstruction for SAFIR-II is performed via the open-source Software for Tomographic Image Reconstruction (STIR) [69]. STIR is a widely used tool for PET image reconstruction, featuring various usable algorithms and multiple utilities for data analysis. For this thesis, STIR version 3 was utilized, specifically, to allow for the use of various SAFIR specific features previously implemented by other members of the collaboration. However, efforts are underway to utilize up-to-date versions of STIR in the future.

It should furthermore be noted that the image reconstruction for SAFIR-II is not performed on the previously mentioned DAQ-PC workstation. While its powerful processor and large amount of RAM would seem to make the DAQ-PC an obvious choice, it is limited in the number of images it can process simultaneously. A typical in-vivo measurement for SAFIR-II would see 45 continuous minutes of data being acquired, which are then split into shorter blocks following coincidence sorting. At the targeted timescale of 5 s per image, this would yield up to 900 individual images, a number that would obviously benefit from parallelization. Therefore, the data is transferred to a cluster computing system, where the reconstructions are run in parallel. The cluster features 64 identical nodes, each using two 10-core Intel[®] Xeon[®] CPUs operating at 2.20 GHz and 32 GB of RAM.

The following chapter focuses on the various steps in SAFIR-II's image reconstruction process. Beginning with a brief insight into the basics behind PET reconstruction algorithms, it details metrics used to gauge the scanner's image quality. Several effects impacting the image quality will be examined, along with the methods that are used to correct images from their influence. Lastly, a first in-vivo rat measurement with SAFIR-II is presented.

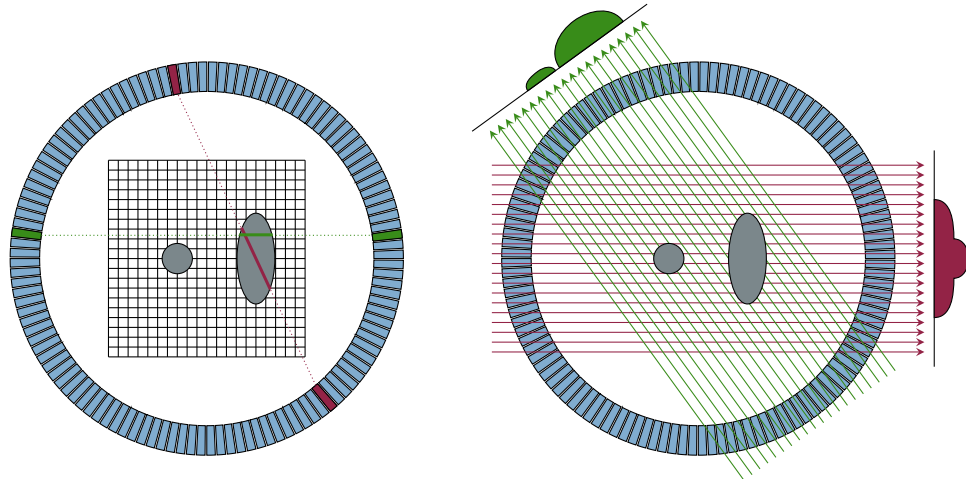


Figure 4.1: Schematic diagram detailing the principle behind PET projection data. **Left:** LORs (green, red) crossing through the underlying activity distribution (gray) in the image space effectively integrate the amount of activity found along them. **Right:** Regarding the resulting data from parallel LORs results in a projection of the distribution along that direction.

4.1 PET Image Reconstruction

Given a three-dimensional distribution of PET tracer material, the accumulated coincidence events acquired during a scan can be interpreted as individual elements of multiple two-dimensional projections for this distribution. For a given crystal pair, the number of coincidences detected between them correlates to the amount of tracer material found along the corresponding line of response (LOR) [70], as shown in figure 4.1. Fundamentally, the problem of reconstructing PET images consists of inverting this projection back into a three-dimensional distribution. However, for any given PET scanner architecture, a subset of possible projection angles can not be acquired, and the statistical nature of radioactive decay will inherently lead to statistical fluctuations for the projection angles that can be. Thus, the true tracer distribution can never be fully recovered, and reconstruction algorithms instead yield approximate estimations. Nevertheless, given sufficient counting statistics, several algorithms have been shown to yield sufficiently accurate images.

4.1.1 Projection Data

Following the successful calibration and post-processing of raw measurement data, the final output of the analysis software consists of a time-sorted list of coincidence events. Within this so-called ‘listmode’ file, each coincidence’s arrival time is specified along with the crystal pair that detected it. To reconstruct an image over a given time period, the corresponding events first need to be sorted into what is known as projection data. In a projection data file, the accumulated number of coincidence events for each LOR is

specified, forming the elements of the previously mentioned two-dimensional projections of the tracer distribution. However, as several angles exist along which such projections can be formed, one can interpret projection data as a four-dimensional histogram, with each LOR ‘bin’ being indexed by four parameters:

- *Segment number*: For a scanner with multiple detector rings, the segment number specifies the ring difference for a given LOR. As such, it represents the LOR’s polar angle.
- *View angle*: Specifies the azimuthal angle of a given LOR, discretized to the number of crystals in a detector ring.
- *Axial and tangential position*: Given a three-dimensional angle specified by the segment and view, all parallel LORs form a two-dimensional projection of the tracer distribution. The position of each LOR is then specified along the axial and tangential direction according to the crystal hit.

For future discussion, it is important to note that the four-dimensional nature of projection data makes it difficult to visualize as a whole. Therefore, usually, only a subset of the data is displayed in the form of a two-dimensional histogram, with two of the four parameters being fixed to a desired value. Two commonly used options for such parameter choices are the viewgram and the sinogram, visualized in figure 4.2. A viewgram is given for a fixed segment number and view angle, while a sinogram is given for a fixed segment number and tangential position. As such, a viewgram represents a two-dimensional projection of the tracer distribution along one specific three-dimensional angle, while a sinogram represents a rotating slice taken through the object.

4.1.2 Filtered Backprojection

Given that the projection data of a PET-scan represents a projection of the original tracer distribution, a simple approach for reconstructing an image would be to inversely project said data back onto the image space. This process is equivalent to the inverse Radon-transform and is commonly referred to as back-projecting an image. An adjusted form of the back-projection algorithm, known as filtered back projection (FBP), was commonly used early on for PET reconstruction.

In mathematical formulation for a two-dimensional example, the Radon-transform [71] $p(\xi, \gamma)$ of a distribution $f(x, y)$ is given as

$$p(\xi, \gamma) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \cdot \delta(x \cos(\gamma) + y \sin(\gamma) - \xi) dx dy \quad (4.1)$$

while the process of back-projecting onto the image space can be formulated as

$$g(x, y) = \int_0^{\pi} p(\xi, \gamma) d\gamma = \int_0^{\pi} p(x \cos(\gamma) + y \sin(\gamma), \gamma) d\gamma. \quad (4.2)$$

Here, γ denotes the projection angle (view angle ϕ for projection data), while ξ denotes the detector element (tangential position r). Figure 4.3 illustrates this process.

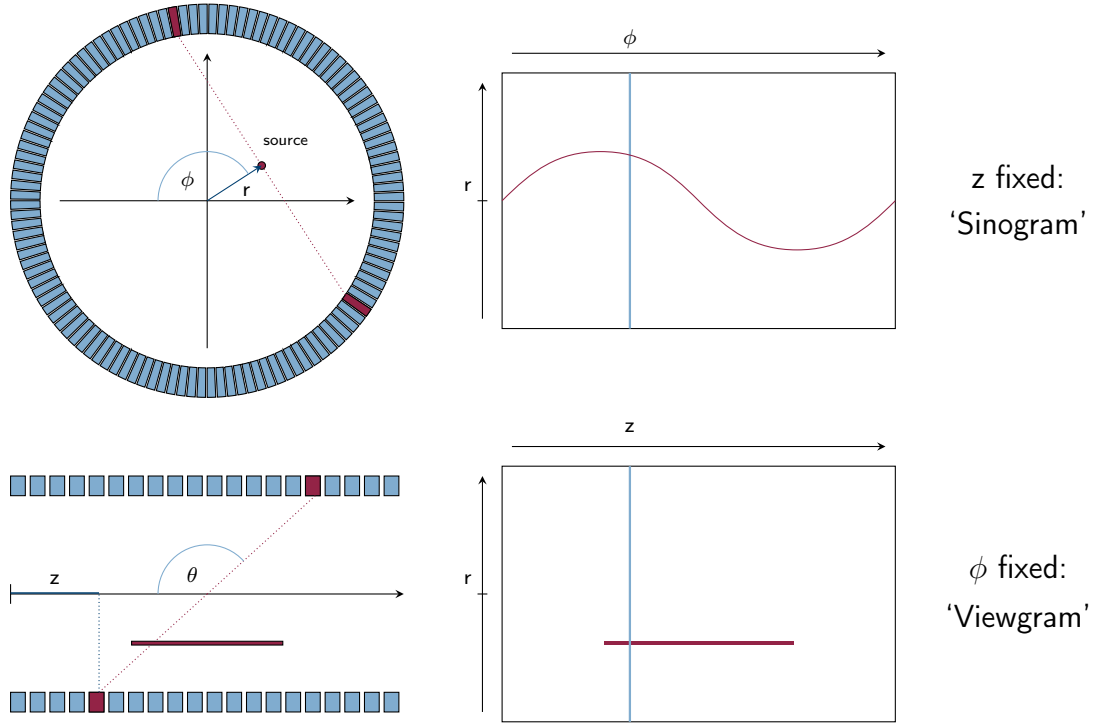


Figure 4.2: Illustration of the functional principle behind sinogram and viewgram formation. Any LOR can be described via the segment number/polar angle θ , view/azimuthal angle ϕ , axial position z , and tangential position r . Keeping the segment number fixed along with either the axial position or the view angle results in a two-dimensional histogram of the coincidence counts per LOR, as either a sinogram or a viewgram respectively.

In principle, this algorithm should result in the back-projected image $g(x, y)$ closely resembling the original distribution $f(x, y)$. However, in practice one will observe $g(x, y)$ to be a significantly blurred representation of $f(x, y)$. Cause for this is not only the limited number of projection angles γ that can be realistically acquired, but also the distribution of acquired data in Fourier space [73].

Following the Central Slice Theorem, the Fourier transform $P(q, \gamma) = \mathcal{F}_1(p(\xi, \gamma))$ of an image projection is equivalent to a slice through the origin of the image's 2D Fourier transform $F(u, v) = \mathcal{F}_2(f(x, y))$. Therefore, one can alternatively obtain the original image $f(x, y)$ through the inverse Fourier transform

$$f(x, y) = \int_0^\pi \int_{-\infty}^\infty P(q, \gamma) e^{2\pi i q \xi} dq d\gamma. \quad (4.3)$$

where $\xi = x \cos(\gamma) + y \sin(\gamma)$ [73]. Within this method, the issue of blurred images can be resolved through the application of a high-pass filter $|q|$, resulting in




Figure 4.3: Schematic diagram detailing the principle behind FBP reconstruction. Overlapping several view angles highlights the area in which tracer material is present. The larger the number of viewing angles, the closer the reconstructed image to the original tracer distribution. (Adapted from [72]).

$$f(x, y) = \int_0^\pi \int_{-\infty}^{\infty} P(q, \gamma) e^{2\pi i q \xi} |q| dq d\gamma. \quad (4.4)$$

The FBP algorithm based on this principle has been shown to be simple and computationally efficient and is widely used in CT image reconstruction. For PET images, it can unfortunately lead to streak artifacts due to the limited sampling inherent to the technique, as shown in figure 4.4. At this point, it has mostly been replaced by iterative methods but is still utilized for certain performance evaluations dictated by the NEMA-NU4 standard.

4.1.3 Iterative Reconstruction

Iterative methods, in contrast to FBP-based reconstruction, account for the statistical nature of PET data. Initially introduced with the maximum likelihood expectation maximization (MLEM) algorithm [74], these methods are based on the concept of estimating a likely tracer distribution and generating projection data based on said estimate. By comparing this projection to the measured data, back-projecting the difference between the two onto the estimated tracer distribution, and repeating the process, an algorithm can iteratively maximize the probability that the measured projection data is produced by the estimated tracer distribution. A diagram illustrating this principle is shown in figure 4.5.




Figure 4.4: **Left:** Image of the calibration phantom reconstructed using the FBP algorithm. The desired smooth circular image is degraded by streak artifacts both inside and outside. **Right:** The same image, but reconstructed using the iterative OSMAPOSL algorithm.

In mathematical formulation, iterative reconstruction aims to determine the distribution of tracer material λ that maximizes the probability $P(n|\lambda)$ of observing the measured projection data n . Given a distribution of tracer material λ_h discretized to voxels $h = 1, \dots, H$, the probability of observing k counts in said voxel follows the Poissonian distribution

$$P(k) = e^{-\lambda_h} \frac{(\lambda_h)^k}{k!}. \quad (4.5)$$

An emission generated in voxel h is detected by the LOR between crystals i and j at a probability $p_{ij,h}$, with the expected number of emissions detected from that voxel then being

$$E[n_{ij,h}] = \lambda_h p_{ij,h} \quad (4.6)$$

such that the probability of detecting $n_{ij,h}$ counts in LOR ij from voxel h becomes

$$P(n_{ij,h}|\lambda_h) = e^{-(\lambda_h p_{ij,h})} \frac{(\lambda_h p_{ij,h})^{n_{ij,h}}}{n_{ij,h}!} \quad (4.7)$$

The likelihood $L(\lambda)$ to observe n is then

$$L(\lambda) = P(n|\lambda) = \prod_{ij,h} e^{-\lambda_h p_{ij,h}} \frac{(\lambda_h p_{ij,h})^{n_{ij,h}}}{n_{ij,h}!} \quad (4.8)$$

with the partial derivative of the log-likelihood $l(\lambda) = \log(L(\lambda))$ being

$$\frac{\partial l(\lambda)}{\partial \lambda_h} = \sum_h \left(\frac{n_{ij,h} p_{ij,h}}{\sum_{h'} \lambda_{h'} p_{ij,h'}} - p_{ij,h} \right) \quad (4.9)$$

Based on this, Shepp and Vardi [74] proposed an iterative scheme to maximize the likelihood following

$$\lambda_h^{n+1} = \lambda_h^n \sum_{ij} \frac{n_{ij,h} p_{ij,h}}{\sum_{h'} \lambda_{h'}^n p_{ij,h'}}. \quad (4.10)$$

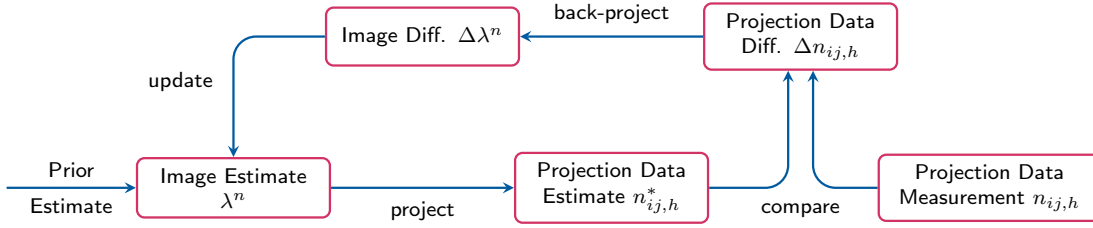


Figure 4.5: Flow-chart diagram detailing the principle behind iterative image reconstruction. The output image consists of the image estimate reached after the desired number of iterations. (Adapted from [75]).




Figure 4.6: **Left:** Schematic diagram illustrating the Siddon ray tracing algorithm (Adapted from [76]). **Right:** Considering multiple interaction points spread over the crystal volume can result in more accurate estimations.

The individual steps of the algorithm are then performed as depicted in figure 4.5:

1. Create a projection of the estimate λ^n as $n_{ij,h}^* = \sum_{h'} \lambda_{h'}^n p_{ij,h'}$
2. Compare to the measured data via $\Delta n_{ij,h} = \frac{n_{ij,h}}{n_{ij,h}^*}$
3. Project the result back and update the estimate: $\lambda_h^{n+1} = \lambda_h^n \sum_{ij} p_{ij,h} \Delta n_{ij,h}$

Iterative algorithms utilize an (often uniform) prior estimate λ^0 and are usually employed until the estimate λ^n converges. Any further iterations generally do not improve the image and can instead introduce additional noise [77].

An integral part of this algorithm is the probability $p_{ij,h}$, also known as the system matrix. It specifies the individual probabilities that an emission contained within a voxel contributes to a given LOR. Several methods exist to calculate the individual matrix elements, with one of the most common ones being Siddon's ray tracing algorithm [78]. For a given LOR and voxel, the algorithm considers the length of the intersecting segment, as shown in figure 4.6. To note here is that, to achieve more accurate estimations for $p_{ij,h}$, it is beneficial to consider the volume of the crystal as well. Usually, several potential interaction points within the crystals are considered, though a higher number of points also results in an increased computational complexity.

Apart from just considering the statistical nature of PET data, a significant additional advantage of iterative algorithms is the fact that estimations for image-altering physical effects can be incorporated into the iterative loop. The system matrix can be adjusted for e.g attenuation-based effects [79], which will be discussed in detail in section 4.3. As a result, the estimated image/tracer distribution will effectively be corrected for such effects, providing an estimate even closer to the true distribution.

4.1.4 Ordered Subset Processing

Iterative methods have been shown to yield images of significantly improved quality for PET scanners and have found widespread use since their inception. However, the improved image quality comes at a price, as the MLEM algorithm has been shown to be computationally expensive compared to analytical methods like FBP. Since its inception, various methods have come forth to accelerate and improve the original algorithm, most notably the inclusion of ordered subset processing [80]. The ordered subset expectation maximization (OSEM) algorithm splits the measured data n^* into k (preferably equally sized) subsets, with the number k being chosen by the user. Each iteration of the MLEM scheme is then performed using only one of these subsets, though a different one each time. One can thus differentiate between sub-iterations, equal to one iteration using one subset, and full iterations, which are completed after k sub-iterations when each subset has been processed once. OSEM reconstruction can accelerate the convergence of the MLEM algorithm approximately at a factor of k . However, studies have shown that an increased number of subsets also has the potential to increase noise in the final image [81]. Therefore, one has to make a tradeoff between reconstruction speed and image quality when using OSEM-based reconstruction. Additional methods exist to address the increased noise, such as the usage of a Gaussian filter applied to the image between iterations. Utilized for SAFIR-II is the ordered subset maximum a posteriori one-step-late (OSMAPOSL) algorithm, an adjusted form of the OSEM algorithm implemented in STIR [82, 83].

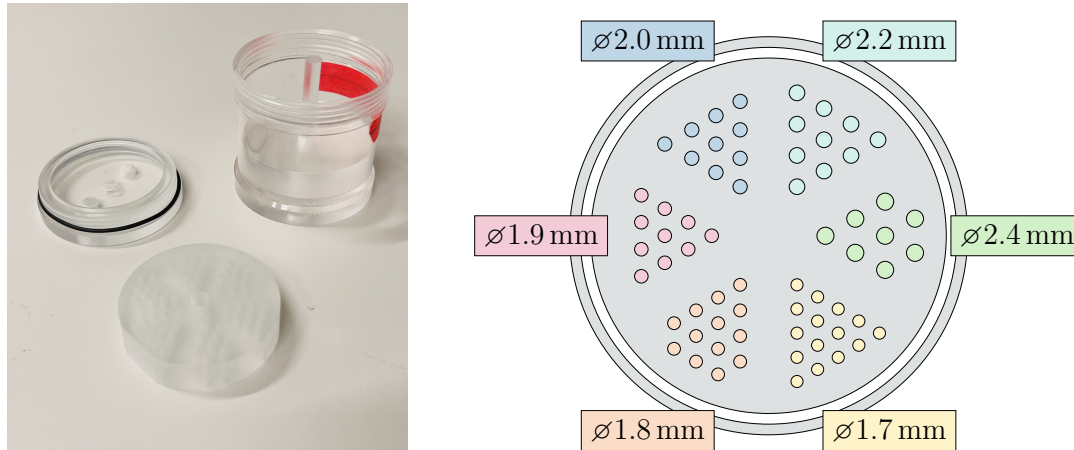


Figure 4.7: **Left:** Image of the utilized Derenzo hot-rod phantom. A circular disk is located in the middle of the phantom, containing the fillable hot rods. **Right:** Schematic of the hot-rod insert. The rod diameters range between 1.7 mm and 2.4 mm, covering the expected range of SAFIR-II's spatial resolution.

4.2 Image Quality

The main focus of an experimenter using a PET scanner is the resulting image quality. This single term incorporates numerous desired parameters, such as the prevalence or absence of noise, the ability to distinguish small structures within the image, or the relative accuracy with which various activity concentrations are displayed.

4.2.1 Derenzo Phantom Measurement

The way an image is reconstructed is one of the main factors determining the image quality. This extends beyond just the chosen algorithm and includes parameters such as the voxel size, the number of utilized subsets, and the number of iterations. Iterative reconstruction methods are generally employed until they converge, but the utilized number of subsets can impact the number of iterations required to reach that point.

To evaluate how many subsets should be selected for SAFIR-II's image reconstruction, and how many iterations would be required, I performed a measurement using a Derenzo hot-rod phantom. This phantom, displayed in figure 4.7, features areas with fillable rods of varying diameters. Rods of equal diameter are grouped in distinct sections and separated at a pitch of twice the rod diameter. This way, one can evaluate the minimum size of structures that can be differentiated by the scanner.

The phantom was filled with ≈ 12 MBq of ^{18}F solute in water, and data was acquired and analyzed as regular. The resulting projection data was then reconstructed using the OSMAPOSL algorithm, using a voxel size of 0.55 mm and six different selections for the number of subsets. Each reconstruction was allowed to operate for ten 'full' iterations, though for evaluation purposes image estimates were also saved for preceding iterations.




Figure 4.8: Images of the Derenzo hot-rod phantom, reconstructed via the OSMAPOSL algorithm while using the specified number of subsets. Each reconstruction was performed over ten full iterations.

Figure 4.8 shows the final images for each subset number for a visual comparison. One can clearly observe the improved image quality for higher subset numbers, with five, six and nine subsets yielding comparable results. This can be explained by lower subset numbers not having fully converged yet, a fact which could also be observed upon close investigation of the images for preceding iterations.

Further supporting this observation is an investigation into line profiles through the images at various iterations, some of which are displayed in figure 4.9. Such line profiles are commonly used to quantitatively evaluate the Derenzo-phantom images beyond just regular visual inspection. For this, the average height of the peaks and valleys in the profile, corresponding to the voxel value intensities for the rods and the spaces between them, are compared in what is known as the valley-to-peak ratio (VTP). This ratio, shown in figure 4.10, clearly still decreases for three or fewer subsets at higher iterations while already having converged after about six iterations for nine subsets. For this reason, I decided to reconstruct any future images for this thesis using nine subsets and ten iterations. This should leave enough remaining iterations in case a different tracer distribution requires longer to converge while not deteriorating the image yet.




Figure 4.9: Line-profiles through the Derenzo hot-rod phantom. These profiles show the absolute voxel values along a line passing through a selected group of rods, in this case for the 2.2 mm diameter. The profiles for images reconstructed with five subsets and nine subsets show the latter having converged, while the former still seems to be improving.

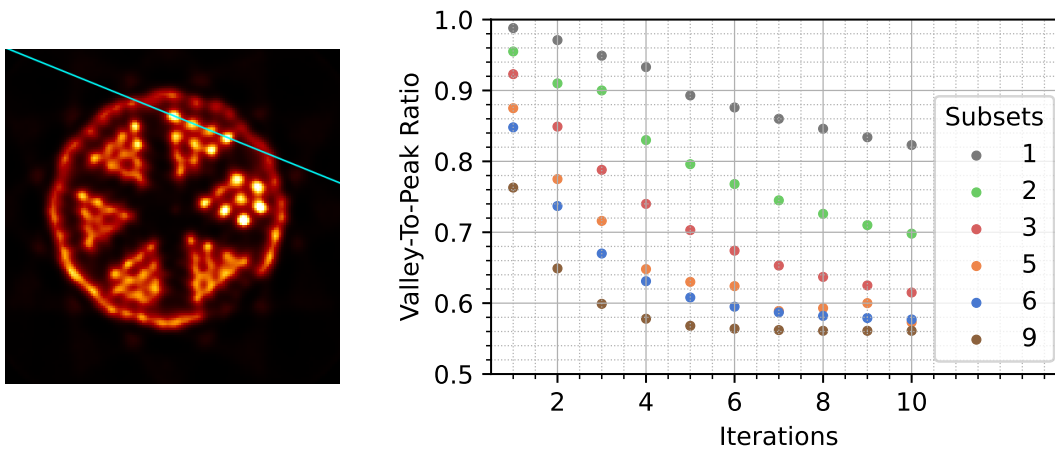


Figure 4.10: **Left:** Line through the 2.2 mm rods, along which the VTP is evaluated. **Right:** VTP for the Derenzo phantom's 2.2mm rods, evaluated for six different subset selections and up to 10 iterations. Easily observable is the quick convergence when using nine subsets.

4 From Coincidences to Images

It should be noted that the time needed to reconstruct each image was roughly equivalent for all subset numbers, requiring about six days and five hours on average. While this immense processing time can be attributed to the low computational power of the used cluster system as well as inefficiencies in the STIR software, it nevertheless shows the necessity for a higher number of subsets.

On reconstruction speed and software

The results of the Derenzo phantom measurement showed two surprising findings of significance for users of the SAFIR-II insert. One concerns SAFIR-II's ability to resolve smaller structures, as rods with a diameter of 2 mm or less were barely distinguishable, if at all. The second issue was the long time required to reconstruct an image, with all selected subset numbers requiring several days to converge. While the long reconstruction time is affected by the low computational power of the cluster system, its main cause are inefficiencies inherent to the STIR software.

STIR is widely used by research groups to reconstruct images and has received a plethora of features for various detector types, reconstruction methods, etc. Due to this it was selected as a reconstruction tool for the SAFIR project, and used in the predecessor system SAFIR-I. The issue of the long reconstruction time became apparent during the testing of SAFIR-II, and it relates to the specifics of the OSEM implementation in STIR. To execute the forward and backward projections necessary for the iterative reconstruction process, the system matrix mentioned in section 4.1.3 is required. The size of this matrix is dependent on the number of LORs and voxels for a given image, which in turn scale quadratically with the number of crystals and linearly with the size of the FOV. As SAFIR-II features more than double the crystals and a much larger FOV compared to SAFIR-I, this matrix becomes excessively large. STIR computes required elements of this system matrix during the reconstruction process, a computationally expensive process aggravated by the fact that it has to be repeated for each iteration. However, a solution to this problem exists, which was utilized in an alternative software named Fast Tomographic Reconstruction (FTR). It was developed by C. Ritzler, a member of the SAFIR collaboration [22].

FTR circumvents the issue of repeatedly having to calculate elements of the system matrix at run time by requiring it to be calculated and stored on a hard drive prior to reconstruction. As the system matrix is static for a given detector geometry and image space, it can be reused for all subsequent reconstructions as well. The downside to this is the large amount of disk space required to store the matrix. Said storage space is also optimally located on a hard drive with a fast read speed, such as an SSD, to maximize the gain in computation speed. As a figure of reference, the system matrix for SAFIR-II and an image space with a voxel size of 0.55 mm per side requires a total of 944 GB of disk space.

To compare the performances of STIR and FTR for SAFIR-II, I reconstructed the Derenzo-phantom data set with both softwares. The resulting images are shown in figure 4.11, having been reconstructed on the previously mentioned DAQ-PC instead of the cluster system due to the storage space constraints for the system matrix. Immediately




Figure 4.11: Images of the Derenzo phantom, reconstructed using the STIR software (**Left**) and the FTR software (**Right**). Both reconstructions utilized the same dataset, and were reconstructed using nine subsets and ten full iterations each. The FTR reconstruction completed within two hours and 50 minutes, while STIR required more than 87 hours.

noticeable was the significantly reduced reconstruction time when using FTR, being more than 30 times faster than STIR. Beyond that, however, a noticeable improvement is also visible in the resolution of hot rods within the phantom. While STIR struggles to resolve rods with a diameter of less than 2 mm, even the 1.7 mm rods are clearly distinguishable using FTR. This could also be observed via a quantitative evaluation using the VTP. To determine if a scanner is able to resolve a given rod diameter, a common approach is to require the VTP to be below the Rayleigh criterion value of 0.735 [84]. In the image reconstructed using STIR, the 2 mm rods exhibited a VTP of 0.639, while for 1.9 mm it already exceeded the criterion at a value 0.748. On the other hand, FTR managed to achieve a VTP of 0.593 for the 1.7 mm rods. The cause for this is likely connected to the respective implementations of the reconstruction algorithms, specifically concerning the ray tracing for each given LOR. While both STIR and FTR use some form of Siddon's ray tracing algorithm, FTR considers the three-dimensional volume of the crystals for its ray tracing. STIR, on the other hand, only considers the tangential direction.

While these facts pose a strong argument for the utilization of the FTR software in the future, this software does not support all of the image corrections mentioned later in this chapter at the time of writing. Likewise, efforts are being made by the developers of STIR to continuously improve and accelerate their software, with recent tests having shown that the newest version can be comparable with the FTR in terms of speed. For the time being, STIR will therefore be employed for reconstructions of SAFIR-II data.




Figure 4.12: **Left:** Image of the utilized image quality phantom. **Right:** Schematic cross section of the phantom, detailing the hot rods, uniform region and cold rods, as specified by the NEMA-NU4 standard.

4.2.2 Image Quality Phantom Measurement

The 2008 NEMA-NU4 standard introduced a specialized phantom along with a corresponding measurement procedure. The purpose of this procedure was to establish a set of parameters common to all preclinical PET scanners that could be used as a base for image-quality-related comparisons. This image quality phantom, displayed in figure 4.12, is comprised of a Polymethylmethacrylate (PMMA) cylinder, featuring cavities which are to be filled with a liquid isotope, in this case ^{18}F solute in water. It is split into three sections, each of which can be used to determine an image quality parameter of interest.

The phantom has an outer diameter of 33.5 mm and is 63 mm long. Within the central region is a 15 mm long cylindrical cavity with a diameter of 30 mm. As the phantom is filled, this region will contain a homogeneous distribution of activity. Any image acquired of this phantom should therefore also show a uniform distribution of activity within this region. The parameter used to evaluate this behaviour is simply called the ‘uniformity’.

Towards one end of the phantom, this cylindrical cavity extends for an additional 15 mm. However, the uniformity is disturbed in this region, as two cylindrical ‘cold region’ chambers are contained within. These chambers are both 15 mm long and 8 mm in diameter and are not filled with activity. Instead, one contains air and the other water. Ideally, any image reconstructed of the phantom should show no activity within these chambers at all, but photon scattering and other effects can impact this. The intensity of voxel values within these chambers in relation to the uniform region specifies the spill-over ratio (SOR), another important image quality parameter.




Figure 4.13: Cross-sections of the reconstructed image quality phantom at 500 MBq. Marked in blue are the regions of interest used to evaluate the various image quality parameters.

Lastly, the other end of the phantom is mostly solid. An exception to this are five 20 mm long rods of varying diameters, which connect to the uniform region and are filled with activity as well. The rod diameters are 1 mm, 2 mm, 3 mm, 4 mm, and 5 mm, respectively. An ideal reconstruction would show a voxel value identical to the uniform region for these rods. However, degrading image quality impacts this, and therefore, the recovery coefficients (RCs) are used to compare the voxel intensity between the rods and the uniform region.

The determination of these parameters from an image of the IQ-Phantom follows a standardized pattern:

1. **Uniformity:** A cylindrical region of interest (ROI) (22.5 mm diameter, 10 mm length) spanning the phantom's central uniform region is examined (see figure 4.13). The Uniformity is reported as the percentage standard deviation of the voxel values contained within.
2. **SOR:** Two ROIs (4 mm diameter, 7.5 mm length) spanning the center of the air- and water-filled chambers are examined. Reported as the SOR is the ratio between the mean voxel value in each chamber and the mean value in the uniform region.
3. **RC:** For each of the five rods, the image slices along the central 10 mm are averaged into a single slice. The ratio between the maximum voxel value of this slice and the mean value in the uniform region is reported as the RC for that rod.

The NEMA standard specifies that these parameters should be evaluated for a measurement performed over a time period of 20 min at an activity of 100 μCi (3.7 MBq). However, as SAFIR-II targets higher measurement activities, it makes sense to evaluate these parameters for higher activities as well. Therefore, I filled the phantom with over 500 MBq of ^{18}F solute in water, and acquired data several times as it decayed. Measurements were taken at 500 MBq, 400 MBq, 300 MBq, 200 MBq, and 100 MBq, with measurement times of 9 s, 11 s, 15 s, 22 s and 44 s respectively. The data analysis was performed as regular, featuring a timing window of 500 ps, an energy window from 400 keV




Figure 4.14: Image quality parameters for uncorrected images, evaluated at various activities.

to 600 keV, and with ICSR disabled. Afterwards, images were reconstructed via the OSMAPSL algorithm within STIR, using nine subsets and ten iterations. The results are displayed in figure 4.14, showing the three image quality parameters relative to the measurement activity. It can be observed that almost all parameters degrade towards higher activities, with the only exception being the RC for the 1 mm rod. However, this rod is also not visible in figure 4.13. This might be due to the fact that it is smaller than SAFIR-II's spatial resolution, and therefore not resolvable. It is also possible that this is the result of a bubble being present within the rod, as it is difficult to fill. Therefore, this part of the phantom might not actually be filled with activity, resulting in this value instead reflecting the image background. The remaining RC values all show a relative degradation of less than 12% between 100 MBq and 500 MBq, as does the uniformity parameter. A stronger degradation is observed for the SOR values (17% for air, 26% for water). The SOR values for air and water also show a significant discrepancy at all activities.

These observations shall serve as a baseline to compare the effects of various image correction methods discussed in section 4.3.

4.3 Correcting for Artefacts

The ultimate aim of the SAFIR-II insert is the ability to successively acquire several detailed and quantitatively accurate images comprised of only a few seconds worth of data. Hindering the efforts towards this goal are a plethora of physical effects degrading the scanner's effectiveness, which have to be remedied through corrections and adjustments like those detailed in the previous chapters. The process of reconstructing an image using the acquired data is subject to many such effects as well. Implemented into SAFIR-II's image reconstruction are correction methods accounting for:

- attenuation effects, as photons might interact with material before reaching a detector element,
- detector normalization, as LORs might exhibit varying efficiencies when detecting coincidences,
- scattered coincidences, as photons undergoing Compton scattering might result in mispositioned LORs, and
- random coincidences, as the analysis software might falsely extract uncorrelated photons as coincidences.

While corrections for each of these effects existed for the SAFIR-I insert, I had to adjust each implementation in order to make it useable for SAFIR-II. I furthermore investigated the effect each correction method had on the image quality phantom data.

4.3.1 Attenuation Effects

One of the most significant aspects affecting a PET image is the presence of the subject or phantom under study in which the tracer is accumulated. The absorption and scattering of 511 keV photons does, of course, not occur exclusively within the scintillation crystals but also within the surrounding tissue, bones, and blood. As photons attenuated in this way either deviate from their original path or do not reach a detector element at all, the number of counts for a given LOR is reduced depending on the type and distribution of attenuating material found along it. For a given LOR between detector elements i and j , the amount of coincidences detected is reduced by a factor

$$A_{ij} = e^{-\int_i^j \mu(x) dx}. \quad (4.11)$$

given a distribution of attenuating material with attenuation coefficients $\mu(x)$ [23], also known as the 'attenuation map'. If $\mu(x)$ is known, measurement data can be corrected for attenuation effects by multiplying the individual bins of a projection data file with the corresponding attenuation correction factors A_{ij}^{-1} .

Several methods can be employed to obtain an attenuation map for a given measurement. A common approach is to perform a CT scan of the study's subject and extract the required attenuation coefficients from there [85], which is an obvious and often employed solution for PET-CT scanners. This is not as easily translatable for a PET-MRI

4 From Coincidences to Images

scanner, as the data extracted from an MRI image does not directly correlate to the distribution of attenuating material. However, by segmenting an MRI image [86] into regions containing various types of tissue, an attenuation map can be estimated. While not utilized in this thesis, this method will be employed to correct future animal studies with the SAFIR-II insert.

For the measurements discussed in this chapter, I generated the attenuation map manually based on the known geometry and material compositions of the utilized phantoms. It is, however, important to consider that additional attenuating material beyond the measured phantom is present due to the detector's construction, particularly the inner carbon fiber cylinder of SAFIR-II's mechanical support structure. While this was not considered for the SAFIR-I detector, the increased axial length of SAFIR-II's FOV causes this to become more significant since LORs with a larger ring difference will pass through more attenuating material of this cylinder. Therefore, it has been included in the utilized attenuation maps.

4.3.2 Detector Normalization

In an idealized PET scanner, the probability of detecting a given coincidence would be independent of the pair of detector elements that was used to detect it. However, in reality, each detector element has an intrinsic, hardware-dependent efficiency to detect a photon dictated by physical effects such as the coupling between crystal and SiPM. Additionally, a photon's incident angle will impact the probability of detecting a hit, creating a systematic error affecting the LOR detection efficiency. For this reason, a normalization correction method is utilized in which the count values for each LOR are multiplied with previously determined normalization factors to emulate uniform detection efficiency.

The determination of these normalization factors can be achieved through various methods. For the DRP and SAFIR-I, an adjusted form of the 'direct' method was utilized. By exposing each LOR to an equal amount of activity, the individual detection efficiencies can be determined from the amount and variation of coincidences detected per LOR [87]. I.e, the normalization factor N_{ij} for an LOR between detector elements i and j can be determined as the inverse of the number of detected coincidences n_{ij}

$$N_{ij} = \frac{1}{n_{ij}}. \quad (4.12)$$

However, equal exposure to activity for all LORs is rather difficult, usually being achieved through several measurements using a thin plate of uniformly distributed ^{68}Ge . An easier alternative is the usage of a large cylindrical phantom [88] filled with uniformly distributed activity, such as ^{18}F solute in water. If placed within the center of the scanner, the amount of activity each LOR is exposed to is equivalent to the intersection length d_{ij} of the LOR with the phantom. Therefore, the normalization factor can then be calculated as

$$N_{ij} = \frac{d_{ij}}{n_{ij}}. \quad (4.13)$$

While the relative simplicity of this method is a clear benefit, it comes at the cost of a few other detrimental factors. For one, the size of the phantom and the corresponding amount of water included in it is a significant source of attenuation, which needs to be corrected for to receive accurate normalization values. Likewise, additional effects discussed later on, such as the random and scattered coincidences in section 4.3.3, need to be corrected as well. Another issue, which is significantly more impactful for SAFIR-II, is the statistical accuracy of the determined normalization factors. Given the statistical nature of radioactive decay, the normalization factors uncertainty correlates to the number of detected coincidences following

$$\frac{\delta N_{ij}}{N_{ij}} = \frac{\sqrt{n_{ij}}}{n_{ij}}. \quad (4.14)$$

To reach a relative uncertainty of 3% or better for all normalization factors, an average of 1000 coincidences would need to be collected per LOR. For this reason, the normalization measurement for the DRP was performed at a starting activity of 93 MBq, acquiring the required data while this activity decayed.

Determining the normalization factors of SAFIR-II this way is possible but inherently unpractical. While the peak sensitivity of SAFIR-II is slightly more than double that of the DRP, the number of LORs has increased by a factor of 16. The increase in measurement activity needed to compensate for this would be too high for SAFIR-II and would also be subject to strong data loss due to detector dead time. Furthermore, a single measurement like the one used for the DRP requires a total of 6.5 TB of disk space. This means that even if one performs multiple measurements at <100 MBq, it would result in massive amounts of data that need a considerable amount of time to be analyzed.

To resolve this issue, I used an alternative to the direct method for SAFIR-II instead. The ‘component-based’ method [89] relies on modeling the normalization factors as

$$\tilde{N}_{ij} = \frac{1}{\epsilon_i \epsilon_j g_i g_j} \quad (4.15)$$

where ϵ_i and ϵ_j denote the individual photon detection efficiencies of detector elements i and j , while g_i and g_j denote geometric coefficients accounting for the incoming photon angle. A major advantage of the component-based method lies in the fact that the individual coefficients are subject to symmetries within the crystal arrangement. Thus, they can be determined at a significantly increased accuracy compared to the factors of the direct method.

To visualize this effect, one can look at figure 4.15. While each of the 180 individual crystals in a ring is located in a different position, rotational symmetry dictates that a minimum of 12 crystals will be subject to the same amount of incident photons. Therefore, the detection efficiencies ϵ_i can be determined from the amount of coincidences from all LORs and not just a single one. Likewise, the geometric coefficients are bound to be identical for various LORs as long as they are all oriented at the same angle towards their respective crystals.




Figure 4.15: Schematic visualizing the principle behind the component-based normalization correction. **Left:** Though the detector elements marked in green and red might vary in detection efficiency, they are exposed to the same amount of photons. A difference in detected counts will thus reflect the detector efficiencies. **Right:** The LORs marked in red are rotationally invariant from each other, and thus share the same geometric coefficient. A different coefficient is present for those marked in green, as well as those in black.

Therefore, I determined the component-based normalization factors for SAFIR-II using the following procedure:

1. Acquire data by measuring a normalization phantom filled with 100 MBq ^{18}F solute in water. The phantom, depicted in figure 4.16, has a fillable volume with a radius of 37 mm and a length of 160 mm. Therefore, if placed within the scanner's center, it covers the entire length of the FOV.
2. Correct the data for attenuation effects and the effects of random and scattered coincidences (see section 4.3.3).
3. Determine the detection efficiencies ϵ_i by averaging the amount of coincidences acquired for detector elements in rotationally symmetric positions and comparing it to each element.
4. Correct the data for the individual detection efficiencies, then determine the geometric coefficients g_i by comparing the average number of coincidences for rotationally symmetric LORs to the average of all LORs.
5. Calculate the normalization factors \tilde{N}_{ij} for each LOR using equation 4.15.



Figure 4.16: Image of the Normalization phantom.




Figure 4.17: Comparison of a normalization phantom sinogram, the correction factors calculated via the direct method and via the component-based method.




Figure 4.18: Histogram of the difference in normalization values between the direct and the component-based method. Displayed in orange is a Gaussian fit, though a χ^2_ν of $5.9 \cdot 10^6$ suggests that the values follow a more complex distribution.

To compare the two methods, I also used the data from step 2 to calculate normalization factors with the direct method, utilizing equation 4.13 for each LOR instead. Figure 4.17 displays sinograms of the normalization factors determined via the component-based and direct methods from the same data set. As one can see, the component-based method yields a significantly less noisy result while not exhibiting any artifacts deviating from the direct method. The histogram depicted in figure 4.18 reinforces this observation, showing the difference in determined values following a distribution similar to a Gaussian without side-peaks.

Additionally, the normalization factors of both methods were utilized in the reconstruction of the same aforementioned image-quality data set to compare their effect. However, no significant difference in the image quality parameters could be observed when comparing the results. For these reasons, the implemented component-based method was judged to be adequate for utilization in SAFIR-II, and has been incorporated into the reconstruction pipeline going forward.




Figure 4.19: Comparison of a normalization phantom sinogram, the correction factors calculated via the direct method and via the component based method for a dataset with an incorrect timing calibration. The incorrect calibration results in a section with reduced coincidences in the sinogram (marked in red), which the direct method compensates for through increased correction factors, while the component-based method only results in a higher average for the region.

While these results seem to suggest that the presented implementation of the component-based method is inherently superior to the direct method, a significant factor has been left out. The component-based method improves on the statistical uncertainty of the direct method by working with averaged values of several LORs. However, this can result in the elimination of systematic effects as long as these effects stem from specific LORs instead of individual detector elements.

Figure 4.19 shows a sinogram of the normalization phantom obtained following an erroneous timing calibration of the data, along with the corresponding correction factors from the two separate methods. As one can observe in the original sinogram, certain LORs exhibit a reduced number of coincidences. Despite its statistical uncertainty, the direct method accounts for this imbalance by increasing the correction factors for the specific LORs. The component-based method, on the other hand, is incapable of this, averaging over the various LORs and resulting in a smearing effect for the region surrounding the artifact instead. While this issue stemmed from an incorrect calibration of the data and has since been remedied, it should be noted that, in principle, the component-based method trades systematic accuracy for statistical accuracy. A careful comparison of the results of both methods should thus be performed before a choice is made concerning which method to utilize.

4.3.3 Random & Scattered Coincidences

Whereas the attenuation and normalization corrections represent multiplicative factors that attempt to account for lost coincidences, the corrections for random and scattered coincidences are subtractive factors that serve to remove the influence of unwanted events falsifying our data. Scattered coincidences are the result of one or both photons of a positron-electron annihilation undergoing Compton-scattering before reaching a detector element. Given a small enough scattering angle, these photons will still possess a high




Figure 4.20: Schematic visualizing the three major types of coincidences. For a true coincidence (a), the two emitted photons reach their respective crystals uninhibited, resulting in an LOR crossing through the tracer. A random coincidence (b) occurs if only one photon each from two separate decay events is detected, resulting in an incorrect LOR. A scatter (c) is the result of one or both photons from a single emission deviating from its original path, shifting the LOR away from its true position.

enough energy to fall within the analysis software’s energy window. Therefore, they will be detected as a valid coincidence but contribute to a different LOR from their original one. Random coincidences, on the other hand, are a direct result of the statistical nature of radioactive decay, the detector’s limited detection efficiency, and its finite timing resolution. Since multiple decay events can also occur in quick succession, cases exist in which two photons of separate decays are detected within the same coincidence window, resulting in a false contribution to an LOR. The principles behind both types of ‘false’ coincidence are illustrated in figure 4.20.

Correcting for Random Coincidences

Given a pair of detector elements detecting single photons at rates $S_{i/j}$, and an algorithm sorting the resulting data for coincidences with a timing window τ , the rate R_{ij} at which uncorrelated events are falsely detected as coincidences can be estimated [90] via

$$R_{ij} = 2 \tau S_i S_j. \quad (4.16)$$

This ‘singles rate’ method is one of two approaches usually employed for the estimation of random coincidences in PET instrumentation, the other being the ‘delayed window’ method [91]. Both methods have benefits and drawbacks, with the singles rate method having the tendency to overestimate the true random coincidence rate, especially at higher activities [23]. Therefore, random estimation for SAFIR-II is done using an adjusted form of the singles rate method, called the ‘singles prompt’ method [92]. This method overcomes the problem of overestimation by taking information concerning the rate of prompt coincidences into account, with prompts in this context referring to detected coincidences, be they random or otherwise. The equation 4.16 is adjusted to

$$\tilde{R}_{ij} = 2 \bar{\tau} \bar{S}_i \bar{S}_j, \quad (4.17)$$

where $\bar{\tau}$ denotes the *effective time coincidence window*

$$\bar{\tau} = \tau \frac{e^{-\tau(\lambda+S)}}{(1-2\tau\lambda)^2} \quad (4.18)$$

and $\bar{S}_{i/j}$ are the *effective singles count rates*

$$\bar{S}_i = S_i - P_i e^{\tau(\lambda+S)}. \quad (4.19)$$

$P_i = \sum_j P_{ij}$ is the rate of prompts (coincidences) acquired by detector element i , and $S = \sum S_i$ is the total rate of all singles detected by the scanner. λ solves the equation

$$2\tau\lambda^2 - \lambda + S - P e^{\tau(\lambda+S)} = 0, \quad (4.20)$$

where $P = \sum P_i$ is twice the total rate of prompts detected by the scanner. To correct an image, a sinogram containing the estimated number of random coincidences for each LOR is generated using equation 4.17, which is then later integrated into the reconstruction process.

Correcting for Scattered Coincidences

Due to the finite energy resolution inherent to each PET scanner, photons scattered at a low angle can not be differentiated from those that did not undergo scattering. While it cannot be known which coincidences included photons scattered from their original path a posteriori, an estimation can be made based on the observed data and information about the distribution of attenuating material within the FOV.

For SAFIR-II and its predecessor systems, the number of scattered coincidences is estimated using the single-scatter-simulation method [93] implemented within STIR [94]. This method is based on generating an estimated distribution of activity given the measurement's projection data and simulating scattered coincidences using the Klein-Nishina equation [31] and the attenuation map from section 4.3.1. This process is performed iteratively, looping over the following steps:

1. Generate an estimate of the activity distribution by reconstructing the data using the OSEM algorithm. Include available corrections for attenuation, normalization and randoms.
2. Determine the voxels in which scattering can potentially occur from the attenuation map. For all LORs, calculate the probabilities that a coincidence scattered in a given voxel is detected in said LOR.
3. For a given LOR, store the sum of all voxel probabilities as the scatter estimate. The projection data resulting from all LORs is saved as the estimate of this iteration and integrated into the reconstruction of step (1) for the next iteration.

This repeated reconstruction and simulation is computationally very expensive. Therefore, it is performed on a downsampled image and projection space, i.e., the voxel size of the image and attenuation map is multiplied by a factor of two in each direction. The final scatter estimate is then upsampled and scaled using trispline interpolation to make it compatible with the original data [23].

4.3.4 Results

To investigate the effects of each correction method, I applied them to the image quality phantom data mentioned in section 4.2.2. Additionally, I reconstructed and corrected images of the cylindrical calibration phantom previously mentioned in section 2.3.3. It should be noted, however, that the corrections were applied successively instead of considering each individually. This was done due to the fact that the presented correction methods should not be considered functionally separate from one another. Estimating the scattered coincidences requires knowledge of the attenuation map, and the estimate is furthermore improved if random coincidences and normalization factors are taken into account. Similarly, the measurement used to determine the normalization factors is subject to attenuation effects, random coincidences, and scattered coincidences as well, such that correcting it for these values is beneficial. It thus only makes sense to perform a scatter correction if an attenuation correction is also performed. Therefore, the reconstructed images featured the following configuration of corrections:

1. Attenuation correction only, calculated from a manually generated attenuation map.
2. Attenuation and normalization correction, with the normalization factors being generated using the component-based method. The normalization phantom data was corrected for attenuation effects only.
3. Attenuation, normalization, and random correction, the latter being evaluated using the singles prompt method. Here, the normalization data was additionally corrected for random events.
4. All corrections, with the scatter correction being generated using the single-scatter-simulation method. The simulation incorporated attenuation, random and normalization corrections. The normalization data was corrected for attenuation, random and scatter events.

For the image quality phantom, I evaluated each parameter for each of these configurations. Furthermore, the parameters were evaluated across an activity range between 100 MBq and 500 MBq, as some corrections vary in effectiveness depending on the activity. The calibration phantom data is used to visualize the effect of each correction on the reconstructed image.




Figure 4.21: Visualization of the various correction methods' effect on a reconstructed image of the calibration phantom. Each successive image has the preceding corrections applied as well, i.e. 'scatter' has been corrected for attenuation, normalization, random and scattered coincidences.

Visual Comparison

The calibration phantom was measured at an activity of 376 MBq for a total of 60 s. Images were reconstructed using STIR's OSMAPOS algorithm featuring nine subsets and ten full iterations, incorporating each of the correction configurations mentioned above.

Figure 4.21 shows a coronal cut through the respective reconstructed images. As can be observed, the attenuation and normalization corrections seem to have the strongest impact, removing a visual artifact within the center of the phantom. The presence of this artifact, as well as its removal, can be readily explained. As the phantom is placed within the center of the FOV, LORs with a large ring difference cross through its center as well. These LORs pass through more attenuating material, thus receiving fewer total counts, resulting in a darker image. Likewise, the detection of such LORs is less likely since they reach their respective detector crystal at a sharp angle. The first effect is remedied by the attenuation correction, while the second is reflected in the geometric coefficient of the normalization correction.

Uniformity

Judging from the previously discussed visual comparison of reconstructed images, one would expect the attenuation and normalization corrections to have a significant impact on the uniformity parameter of the image quality phantom. This is confirmed when observing figure 4.22, in which the uniformity values are shown at various measurement activities for each successive correction applied. The scatter correction, on the other hand, seems to degrade the uniformity again, which is likely brought on by a slight increase in noise. The same can be said for the random correction, though to a much lesser degree. These results are very much in line with observations previously made for the DRP [23].




Figure 4.22: Uniformity of the image quality phantom at various activities, compared when applying the various correction methods. Each successive image has the preceding corrections applied as well, i.e. ‘scatter’ has been corrected for attenuation, normalization, random and scattered coincidences.

Spill-over Ratio

The observed SORs for the uncorrected image quality phantom displayed a significant discrepancy between the air and water chambers. This is to be expected, as the difference in material also represents a difference in attenuation coefficient, which also contributes to increased photon scattering. It thus makes sense to expect the SOR values to change most through the application of the attenuation and scatter corrections. This hypothesis is confirmed in figure 4.23, which shows the values for each chamber and applied correction. The final SOR values between the water and air chamber after all corrections still show a slight discrepancy, though it is significantly reduced from the uncorrected image.

Recovery Coefficients

Observing the RC values displayed in figure 4.24 shows each correction varying in effectiveness depending on the rod diameter. For the larger diameters, a significant improvement is visible when utilizing the scatter correction, while the other methods seem less impactful. For the 2 mm and 3 mm rods, however, the attenuation correction seems to gain in importance. The 1 mm rod displays significantly different behaviour. However, as mentioned, this rod was not visible in the image, either due to SAFIR-II’s spatial resolution or a potential bubble in the rod. As such, the observed effects might also be reflective of the image background in which no activity is present.

The applied corrections also seem to decrease the activity dependence of the RC values, with all diameters except for 1 mm exhibiting a relative degradation of less than 5% after all corrections.




Figure 4.23: SOR of the image quality phantom at various activities, compared when applying the various correction methods. Each successive image has the preceding corrections applied as well, i.e. ‘scatter’ has been corrected for attenuation, normalization, random and scattered coincidences.




Figure 4.24: RCs of the image quality phantom at various activities, compared when applying the various correction methods. Each successive image has the preceding corrections applied as well, i.e. ‘scatter’ has been corrected for attenuation, normalization, random scattered coincidences.




Figure 4.25: Values for the CF for various measurement activities.

Quantitative Calibration

For the purposes of SAFIR-II and this thesis, an additional parameter of interest exists that warrants examination. As the scanner is intended to be used for quantitative imaging, an image's voxel values need to be translated into activity concentrations. This is done using the cylindrical calibration phantom, which is measured while being homogeneously filled with activity. The ratio between the activity concentration and the observed voxel values from this measurement, also known as the calibration factor (CF), can then be used for other measurements. Ideally, this ratio would be constant irrespective of the total measurement activity.

However, this is hindered by several detrimental factors affecting the scanner at higher activities. For one, the detector dead time mentioned in section 2.3.3 decreases the number of detected coincidences depending on the measurement activity. Similarly, random coincidences can also result in the loss of coincidences, as a 'true' coincidence is rejected due to the presence of a third detected hit. As such, the calibration factor is bound to be dependent on the total measurement activity.

To investigate the calibration factor I measured the calibration phantom for a range of activities between 500 MBq and 6 MBq. A cylindrical ROI, 44 mm long and 22.5 mm in diameter, was placed within the center of the reconstructed phantom image after applying all corrections. The CF is given as the activity concentration within the phantom multiplied by the measurement time and divided by the average voxel value within the ROI. As can be observed in figure 4.25, it is dependent on the total measurement activity. Images acquired using SAFIR-II should thus always be calibrated using a factor acquired at the respective activity.




Figure 4.26: Images of the studied Sprague-Dawley rat, acquired 40 minutes after tracer injection. The images were reconstructed over time frames with different lengths. The scale is given in arbitrary units identical for all images, and is depicted along with the first image. The bright red circular spot on the left of the rat's head corresponds to one of the Eppendorf tubes used to assist in image positioning.

4.4 SAFIR-II In-vivo: Zeus the Rat

The aim of the SAFIR collaboration is to create a device capable of imaging fast metabolic processes within small rodents. While the measurements presented in this thesis have shown that SAFIR-II fulfills the requirements outlined at the project's start, there is no better way of showing this than a measurement featuring a live animal.

I performed an experiment featuring a Sprague-Dawley rat (nicknamed 'Zeus') placed within the scanner. It was positioned such that the brain was located within the center of the FOV, and an MRI image was acquired using a T2-Turborare sequence. To assist in the image positioning, four Eppendorf tubes were placed next to the head of the rat; each filled with 1.3 MBq FDG. An injection of 2 ml FDG at an activity of 283 MBq was administered into the femoral vein, after which data was acquired using SAFIR-II for a continuous 45 min. During analysis, the data was split into smaller time-frames to investigate the image quality. Images were reconstructed without any correction methods and using the FTR software, as no attenuation map was available at the time.

Figure 4.26 shows three images acquired 40 minutes after tracer injection, reconstructed over a 10 s, 30 s, and 5 min time frame each. The image reconstructed from 5 min of data clearly shows small anatomical structures within the rat. The horseshoe-shaped cerebral cortex is located in the image's center, along with the optic nerves behind the eyes. Also




Figure 4.27: Coronal view of the rat brain in vivo. **Top:** MR images acquired using a T2-Turborare sequence. **Middle:** PET images acquired over the timespan from 40-45 minutes. **Bottom:** PET images fused with MR images. (Image provided by S. Saghamanesh.)

observable is the myocardium at the bottom, even though it is located towards the edge of the FOV.

Observing the images taken over shorter time frames, it is clear that the noise level increases due to the decreasing statistics. However, the aforementioned anatomical structures are still distinguishable even for the 10s image. As this image was acquired for an activity of ≈ 218 MBq (283 MBq starting activity, decaying over 40 minutes), we can expect similar images for a 5s time frame at ≈ 450 MBq.

Figure 4.27 shows the acquired MR images for the measurement, along with the PET image from figure 4.26 reconstructed over 5 min for the same area. Overlaying the PET image on top of the T2-Turborare acquisition demonstrates the close match in the anatomical structures.

Figure 4.28 shows images acquired 0s, 30s, 10 min, and 40 min after tracer injection, each using a 30s time frame. Clearly observable is the tracer distribution changing over time. FDG is progressively accumulated within the brain before concentrating within the myocardium as well. Once these images can be quantitatively calibrated, Time-Activity Curves tracking the accumulated tracer material within a region of interest (such as the brain) should be easily obtainable.

Overall, the observed images are of satisfactory quality, showing SAFIR-II's potential for the purposes of quantitative kinetic studies.




Figure 4.28: Images of the studied Sprague-Dawley rat, acquired over 30s each. The images were reconstructed at different starting times. The voxel values of all images follow the same scale, showing the evolving tracer distribution.

***Statement of personal contribution:** All measurements and results shown in this chapter have been performed and evaluated by the author. The in-vivo rat experiment was performed with the assistance of Dr. Matthias Wyss and Dr. Afroditi Eleftheriou, who handled the animal and tracer injection. STIR is an open-source library, and the FTR software was developed by C. Ritzer. The attenuation and scattered coincidence correction methods were initially part of the STIR library and made compatible with the Dual-Ring-Prototype detector by Dr. Parisa Khateri. Dr. Khateri also implemented the direct normalization correction, while the random correction was developed by Josep F. Oliver. I adapted all correction methods for use with SAFIR-II and additionally implemented the component-based normalization correction. Dr. Pascal Bebie assisted in adapting the random correction method.*

5 Conclusion

SAFIR-II represents the culmination of more than ten years of research and development to construct a PET scanner for a very specialized task. The design steps taken to create the final scanner and the considerations behind those steps, both of which have been the focus of this thesis, have resulted in a uniquely powerful device.

The hardware components and system design of SAFIR-II, presented in chapter 2, facilitate measurements at tracer activities far exceeding those of all other preclinical PET scanners. While the observed peak sensitivity of 2.23% at low activities is less than initially expected, the adjusted scintillation crystal size that caused this has helped improve the scanner's performance at higher activities. It has been shown that detected events are processed fast enough that less than 5% of events are lost due to dead-time effects, even at the highest targeted activity. An effective cooling solution in the form of a closed loop air cooling circuit operating at 15 °C is utilized to dissipate the 580 W of heat that the scanner generates. The scanner's influence on the Bruker MRI within which it operates has been shown to be acceptable. A 9.2% degradation of the MRI SNR still lets the system perform well above the manufacturer's specifications, and Eddy-current distortions during EPI acquisitions can be compensated for by adjusting the utilized sequence.

The data analysis software discussed in chapter 3 is crucial in obtaining coincidence events. At low activities, a high timing resolution of 221.2 ps FWHM has been observed, facilitated by the various employed timing calibration procedures. The proper selection for a timing threshold has also been shown to conserve a good resolution for higher activities, which is crucial in reducing the number of random coincidences. ICSR techniques increase the peak sensitivity to up to 3.89% while deteriorating the spatial resolution by less than 10%. This is aided by the high observed energy resolution, which ranges between 12.1% and 13.1% FWHM across the targeted activity range.

Investigations into the image reconstruction and quality from chapter 4 have demonstrated that SAFIR-II can resolve structures down to a 1.7 mm diameter, as measured using a Derenzo hot-rod phantom. Correction methods for attenuation effects, detector normalization, random coincidences, and scattered coincidences help to improve the image quality, as shown using parameters defined in the NEMA-NU4 standard. Images taken during a first in-vivo rat study clearly display small anatomical structures such as the myocardium.

Overall, it has been shown that SAFIR-II meets the specifications outlined at the project's start. It is functional and capable of operating at high measurement activities while providing high-quality images.

5.1 SAFIR-II by Comparison

SAFIR-II was created because, at the time, no other device could fulfill the task intended for it and because it was considered unlikely that such a device would become available in the near future. To my knowledge, there is no other existing PET-MRI scanner that can acquire preclinical images as quickly as SAFIR-II. However, during its development, several manufacturers have created commercial scanners that exhibit excellent performance. While the performance focus of SAFIR-II is significantly shifted from these systems, it is nevertheless worthwhile to compare the scanners with one another.

Table 5.1 lists technical specifications and performance parameters mentioned in this thesis for SAFIR-II and three commercial scanners. The chosen scanners include the Mediso nanoScan[®] [95, 96], the Scintica SimPET[™]-XL [97], and a PET insert by Bruker [98]. They were selected because characterizations according to the NEMA NU4 standard have been published for each of them. The SAFIR-I insert [35, 24] was included as well to highlight SAFIR-II's improvements.

At the inception of the SAFIR project, the targeted spatial resolution of 2 mm was a standard value for commercial preclinical PET scanners. During SAFIR-II's development, modern scanners have improved on this, with sub-millimeter spatial resolutions becoming more commonplace. Likewise, the sensitivity of SAFIR-II does not reach that of commercial scanners, even if it outperforms its predecessor. As presented in chapter 2, this was to be expected, as the low sensitivity is a consequence of design choices made to facilitate its high target activity. Likewise, SAFIR-II's crystal size was adapted to the targeted spatial resolution. To improve the spatial resolution, smaller crystals would be needed, which would significantly increase the complexity of the electronics.

Table 5.2 compares the different scanners for the various image quality parameters evaluated using the NEMA-NU4 image quality phantom. Here, one can observe SAFIR-II's excellent image quality, showing SOR values on par with the nanoScan[®] and an unparalleled uniformity. On the other hand, the RC values of smaller rod diameters are significantly lower for both SAFIR scanners compared to their commercial counterparts. This is a direct result of the differences in spatial resolution.

SAFIR-II also exhibits excellent performance in other aspects. While the timing resolutions of the three mentioned scanners have not been published, SAFIR-II's 221 ps are remarkably close to the 214 ps of the Siemens Biograph Vision, a top-of-the-line clinical scanner using TOF information [14]. Lastly, while detailed evaluations of SAFIR-II's count rate performance are still to be done, the minimal loss of data at high activities is a significant achievement, enabling scans over a large range of injected activities.

5 Conclusion

Scanner	Axial FOV	Inner diameter	Crystal Size	Peak Sensitivity	Spatial Resolution
SAFIR-II	145 mm	114 mm	$2.0 \times 2.0 \times 13 \text{ mm}^3$	3.89 %	1.7 mm
SAFIR-I	54 mm	114 mm	$2.12 \times 2.12 \times 13 \text{ mm}^3$	2.30 %	—
nanoScan [®]	94 mm	160 mm	$1.12 \times 1.12 \times 13 \text{ mm}^3$	8.4 %	0.8 mm
SimPET-XL [™]	110 mm	76 mm	$1.2 \times 1.2 \times 10 \text{ mm}^3$	7.25 %	—
Bruker	150 mm	114 mm	$50 \times 50 \times 10 \text{ mm}^3$	11.0 %	0.7 mm

Table 5.1: Comparison of various performance parameters and technical specifications between SAFIR-II and other PET-MRI scanners. The spatial resolution was compared via Derenzo phantom measurements, which were not performed for all scanners. The Bruker insert utilizes monolithic crystals for its detection, hence the larger size.

Scanner	Uniformity	SOR		RC				
		Air	Water	5 mm	4 mm	3 mm	2 mm	1 mm
SAFIR-II	2.96 %	0.057	0.075	0.92	0.79	0.58	0.32	0.053
SAFIR-I	4.8 %	0.218	0.220	1.08	0.84	0.54	0.28	0.04
nanoScan [®]	3.52 %	0.058	0.062	1.03	0.98	0.90	0.84	0.26
SimPET-XL [™]	3.89 %	0.036	0.036	0.95	0.91	0.79	0.62	0.14
Bruker	6.5 %	0.12	0.22	0.94	0.95	0.91	0.64	0.14

Table 5.2: Comparison of NEMA image quality parameters between SAFIR-II and other PET-MR scanners. The values quoted for SAFIR-II are those acquired at a measurement activity of 500 MBq.

5.2 Outlook

With the design and construction of SAFIR-II having been completed, the scanner is now operational and performing as desired. Beyond minor technical adjustments, no further modifications or additions are planned for the detector. The minor adjustments include the software transition to STIR version 6, plans to shorten the time required to reconstruct images from several days to mere hours, and improvements to the tubing of the air cooling circuit. With the scanner now useable, several measurements and studies exist which can be performed in the future.

Concerning technical performance measurements, a complete characterization of the scanner according to the NEMA-NU4 standard might be conducted and published in the future. Some of these measurements have already been performed for this thesis, but particularly the measurement of the Noise-Equivalent Count Rate (NECR) and a point source-based spatial resolution measurement are still outstanding. Additionally, exploring the limits of SAFIR-II's capabilities might yield benefits for its performance. The scanner was designed for acquisitions at 500 MBq, but it might be able to acquire images at even higher activities. There are also several adjustable parameters for SAFIR-II that can be explored to boost SAFIR-II's performance further. While some of these parameters, such as the timing threshold of the PETA-8 front end, were already examined in this thesis, many more parameters are left unexplored. This might include image reconstruction-specific parameters, such as various filterings applied between iterations, alternative reconstruction algorithms, and adjustments to the correction methods. Room for improvement might also be found in the data analysis software, such as the usage of a coincidence timing window of variable length depending on the length of an LOR or accepting triple coincidences instead of rejecting them. The latter has the benefit of acquiring more 'true' coincidences at the drawback of also acquiring more noise through 'randoms', but depending on the performance of the utilized random correction, this might be a beneficial choice.

Lastly, preparations for a preclinical study on the Oxygen-Glucose-Index (OGI) mentioned in chapter 1 are well underway at the time of writing. This study will include measurements on the cerebral metabolic rate of glucose and oxygen in rats, both performed once for a resting state and once during focal activation via electrical stimulation of the forepaw. It will be performed using rats, though potential future studies might investigate the brain metabolism in various mouse strains. As an example, effects of high-fat or ketogenic diets on brain glucose metabolism might be assessed. While no concrete plans for studies targeting other organs exist yet, studies on cardiac perfusion using simultaneous MR and PET imaging could also benefit from the performance features of SAFIR-II.

Glossary

ADC	analog-to-digital converter
APD	avalanche photodiode
ASIC	application-specific integrated circuit
ATP	adenosine triphosphate
BGO	bismuth germanium oxide
CT	computed tomography
CBF	cerebral blood flow
CPU	central processing unit
CF	calibration factor
CRC	cyclic redundancy check
DRP	Dual Ring Prototype
DAQ	Data Acquisition
EPI	echo-planar imaging
ESR	enhanced specular reflector
FBP	filtered back projection
FDG	[¹⁸ F]Fluorodeoxyglucose
FOV	field of view
FTR	Fast Tomographic Reconstruction
FPGA	field-programmable gate array
FWHM	full width at half maximum
FSM	finite state machine
GUI	graphical user interface
IC	integrated circuit

5 Conclusion

ICSR	inter-crystal scatter recovery
JTAG	Joint Action Test Group
LDO	low-dropout
LOR	line of response
LSO	lutetium oxyorthosilicate
LYSO	lutetium-yttrium oxyorthosilicate
MPPC	multi-pixel-photon counter
MRI	magnetic resonance imaging
MWS	multi-window sorting
MLEM	maximum likelihood expectation maximization
NEMA	National Electrical Manufacturers Association
OV	overvoltage
OGI	Oxygen-Glucose-Index
OSEM	ordered subset expectation maximization
OSMAPOS	ordered subset maximum a posteriori one-step-late
PET	positron emission tomography
PCB	printed circuit board
PDE	photon detection efficiency
PMT	photomultiplier tube
PMMA	Polymethylmethacrylate
RC	recovery coefficient
ROI	region of interest
SAFIR	Small Animal Fast Insert for MRI
SAFIR-I	SAFIR detector I
SAFIR-II	SAFIR detector II
SFCM	SAFIR Fast Control Master
SBD	SAFIR Bias Distribution

SBTV3	SAFIR Bias and Temperature Version 3
SDIP3	SAFIR Digital Interface for PETA Version 3
SOR	spill-over ratio
SPPD	SAFIR Primary Power Distribution
SSPD	SAFIR Secondary Power Distribution
SFCD	SAFIR Fast Control Distribution
SiPM	silicon photomultiplier
SNR	signal-to-noise ratio
SSD	solid-state drive
STIR	Software for Tomographic Image Reconstruction
SWS	single-window sorting
TDC	time-to-digital converter
TOF	time-of-flight
TTL	transistor-transistor logic
UDP	User Datagram Protocol
VTP	valley-to-peak ratio

Bibliography

- [1] Benjamin P Abbott et al. “Observation of gravitational waves from a binary black hole merger”. In: *Physical review letters* 116.6 (2016), p. 061102.
- [2] Cms Collaboration et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *arXiv preprint arXiv:1207.7235* (2012).
- [3] Giorgio Apollinari et al. “High luminosity large hadron collider HL-LHC”. In: *arXiv preprint arXiv:1705.08830* (2017).
- [4] Joseph A Angelo. *Spacecraft for Astronomy*. Infobase Publishing, 2014.
- [5] Nick Scoville et al. “COSMOS: Hubble space telescope observations”. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 38.
- [6] Jonathan P Gardner et al. “The james webb space telescope”. In: *Space Science Reviews* 123 (2006), pp. 485–606.
- [7] Oliver Kraff et al. “MRI at 7 Tesla and above: demonstrated and potential capabilities”. In: *Journal of Magnetic Resonance Imaging* 41.1 (2015), pp. 13–33.
- [8] Brian Dellon and Yoky Matsuoka. “Prosthetics, exoskeletons, and rehabilitation [grand challenges of robotics]”. In: *IEEE robotics & automation magazine* 14.1 (2007), pp. 30–34.
- [9] Filippo Amato et al. *Artificial neural networks in medical diagnosis*. 2013.
- [10] Michel M Ter-Pogossian and Peter Herscovitch. “Radioactive oxygen-15 in the study of cerebral blood flow, blood volume, and oxygen metabolism”. In: *Seminars in nuclear medicine*. Vol. 15. 4. Elsevier. 1985, pp. 377–394.
- [11] Michael A Steiner. “PET—The History Behind the Technology”. In: *Senior Thesis Projects* (2002). URL: http://trace.tennessee.edu/utk_interstp2/107.
- [12] Ramsey D Badawi et al. “First human imaging studies with the EXPLORER total-body PET scanner”. In: *Journal of Nuclear Medicine* 60.3 (2019), pp. 299–303.
- [13] Terry Jones and David Townsend. “History and future technical innovation in positron emission tomography”. In: *Journal of Medical Imaging* 4.1 (2017), pp. 011013–011013.
- [14] Michael E Casey and Dustin R Osborne. “Siemens biograph vision 600”. In: *Advances in PET: the latest in instrumentation, technology, and clinical practice*. Springer, 2020, pp. 71–91.
- [15] David W. Townsend. “Combined Positron Emission Tomography—Computed Tomography: The Historical Perspective”. In: *Semin Ultrasound CT MR* 29.4 (2008), pp. 232–235. DOI: <https://dx.doi.org/10.1053%2Fj.sult.2008.05.006>.

- [16] Habib Zaidi et al. “Design and performance evaluation of a whole-body Ingenuity TF PET–MRI system”. In: *Physics in Medicine & Biology* 56.10 (2011), p. 3091.
- [17] Drew A Torigian et al. “PET/MR imaging: technical aspects and potential clinical applications”. In: *Radiology* 267.1 (2013), pp. 26–44.
- [18] Hans Herzog and Christoph Lerche. “Advances in clinical PET/MRI instrumentation”. In: *PET clinics* 11.2 (2016), pp. 95–103.
- [19] Peter T Fox et al. “Nonoxidative glucose consumption during focal physiologic neural activity”. In: *Science* 241.4864 (1988), pp. 462–464.
- [20] Robert Becker et al. “Monte-Carlo simulation based estimation of NECR, sensitivity, and spatial resolution of a novel preclinical PET insert for MR”. In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. IEEE. 2015, pp. 1–3.
- [21] Bruker Cooperation. *Preclinical MRI: BioSpec 70/30*. URL: <https://www.bruker.com/en/products-and-solutions/preclinical-imaging/mri/biospec/biospec-70-30-and-94-30.html> (visited on 12/20/2023).
- [22] Christian Edwin Ritzer. “Development and First Performance Tests of the SAFIR Prototype PET-MR Scanner”. PhD thesis. ETH Zurich, 2020.
- [23] Parisa Khateri. “SAFIR Prototype PET Insert: Image Reconstruction and NEMA Characterization”. PhD thesis. ETH Zurich, 2021.
- [24] Pascal Laurent Bebié. “SAFIR-I: Characterization and Optimization of a Small Animal Fast PET Insert for MRI”. PhD thesis. ETH Zurich, 2023.
- [25] Pascal Bebié et al. “SAFIR-I: Design and Performance of a High-Rate Preclinical PET Insert for MRI”. In: *Sensors* 21.21 (2021), p. 7037.
- [26] Pascal Bebié et al. “Effects of inter-crystal optical separation layers on unwanted light crosstalk and on performance parameters of the SAFIR PET/MR scanner”. In: *2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. IEEE. 2021, pp. 1–2.
- [27] William W Moses. “Fundamental limits of spatial resolution in PET”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 648 (2011), S236–S240.
- [28] Meishan Boya Advanced Materials Co. Ltd. *L(Y)SO Scintillation crystals*. URL: https://www.boya-materials.com/?products_63/52.html.
- [29] Dale L Bailey et al. *Positron emission tomography*. Vol. 2. Springer, 2005.
- [30] M.J. Berger et al. *XCOM - Photon Cross Sections Database*. URL: <https://physics.nist.gov/PhysRefData/Xcom/Text/chap4.html> (visited on 12/20/2023).
- [31] Oskar Klein and Yoshio Nishina. “Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac”. In: *Zeitschrift für Physik* 52.11 (1929), pp. 853–868.

Bibliography

- [32] RJ Gehrke, C Casey, and R Ko Murray. “Half-life of lu 176”. In: *Physical Review C* 41.6 (1990), p. 2878.
- [33] C Akapong Phunpueok et al. “Comparison of photofraction for LuYAP: Ce, LYSO: Ce and BGO crystals in gamma ray detection”. In: (2014).
- [34] NEMA: National Electrical Manufacturers Association. “Performance measurements for small animal positron emission tomographs”. In: (2008).
- [35] Pascal Bebié et al. “SAFIR-I: first NEMA NU 4-2008-based performance characterization”. In: *EJNMMI physics* 10.1 (2023), pp. 1–20.
- [36] Hamamatsu Corporation & New Jersey Institute of Technology Slawomir Piatek PhD. *What is an SiPM and how does it work?* URL: <https://hub.hamamatsu.com/us/en/technical-notes/mppc-sipms/what-is-an-SiPM-and-how-does-it-work.html> (visited on 04/19/2024).
- [37] Fischer et al. *PETA6, a multi channel Readout ASIC for SiPM Sensors*. 2015.
- [38] Shoaib Usman and Amol Patil. “Radiation detector deadtime and pile up: A review of the status of science”. In: *Nuclear Engineering and Technology* 50.7 (2018), pp. 1006–1016.
- [39] Finisar/Coherent Corp. *10GBASE-SR 400m Industrial Temperature SFP+ Optical Transceiver*. URL: <https://ii-vi.com/product/10gbase-sr-400m-industrial-temperature-sfp-optical-transceiver/>.
- [40] Christian Ritzer et al. “Compact MR-compatible DC-DC converter module”. In: *Journal of Instrumentation* 14.09 (2019), P09016.
- [41] Analog Devices Inc. *Octal, 14-Bit, SPI Voltage Output denseDAC With 5 ppm/C On-Chip Reference*. URL: https://www.analog.com/en/products/ad5648.html?doc=AD5628_5648_5668.pdf (visited on 12/20/2023).
- [42] Microchip Technology Inc. *AVR ATmega Mikrocontroller IC 8-Bit 20MHz 32KB (16K x 16) FLASH 28-PDIP*. URL: <https://www.microchip.com/en-us/product/atmega328> (visited on 12/20/2023).
- [43] Texas Instruments. *LMK00301 3-GHz 10-Output Ultra-Low Additive Jitter Differential Clock Buffer and Level Translator*. URL: <https://www.ti.com/product/LMK00301> (visited on 12/20/2023).
- [44] now Renesas Electronics Former IDT. *IDT8P34S1208/8P34S1208: 2:8 LVDS 1.8V / 2.5V Fanout Buffer for 1PPS and High-Speed Clocks*. URL: <https://www.renesas.com/us/en/products/clocks-timing/application-specific-clocks/rf-jesd204bc-timing/rf-buffers/8p34s1208-28-lvds-18v-25v-fanout-buffer-1pps-and-high-speed-clocks> (visited on 12/20/2023).
- [45] TDK Lambda. *Z+ Programmable DC Power Supplies 200W/400W/600W/800W in 2U Built-in USB*. URL: <https://www.emea.lambda.tdk.com/ch/products/z-plus>.
- [46] FPZ S.p.A. *LATERAL CHANNEL BLOWERS: SCL K05-MS-2.2kW*. URL: <https://www.fpzusa.com/>.

- [47] MLT AG. Labortechnik. *Luftkuehler Typ 311-W-RB400-LR-PP*, 1.6 kW at -10°C , 2.3 kW at -5°C . URL: <https://www.mlt.ch/umlaufkuehler/kuehlmobile-32-kw-bis-385-kw/>.
- [48] Thermal Grizzly Holding GmbH. *Ultra High Performance Thermal Grease: Kryonaut*. URL: <https://www.thermal-grizzly.com/produkte/2-kryonaut>.
- [49] Yoctopuce S.A.R.L. *USB Environmental Sensors: Yocto-Meteo-V2*. URL: <https://www.yoctopuce.com/EN/products/usb-environmental-sensors/yocto-meteo-v2>.
- [50] E. Mark Haacke et al. 1999.
- [51] D Munalli et al. "Electromagnetic shielding effectiveness of carbon fibre reinforced composites". In: *Composites Part B: Engineering* 173 (2019), p. 106906.
- [52] Bruker Cooperation. *Internal Communication. Specification for BioSpec 70/30 and RF RES 300 1H 75/40 Q TR*. URL: <https://www.bruker.com/en.html>.
- [53] MD. Allen D. Elster. *MRIQuestions.com-Questions and Answers in MRI: Echo-Planar Imaging (EPI)*. URL: <https://mriquestions.com/echo-planar-imaging.html> (visited on 12/20/2023).
- [54] Michael H Buonocore and Lisheng Gao. "Ghost artifact reduction for echo planar imaging using image phase correction". In: *Magnetic resonance in medicine* 38.1 (1997), pp. 89–100.
- [55] Jeff Barrow. "Reducing Ground Bounce in DC-to-DC Converters-Some Grounding Essentials". In: *Analog Dialogue* 41.06 (2007).
- [56] Advanced Micro Devices. *Ryzen™ Threadripper™ PRO 3975WX*. URL: <https://www.amd.com/en/product/10176> (visited on 12/20/2023).
- [57] Kingston Technology. *KTH-PL429/32G DDR4 2933MT/s ECC Registered DIMM CL21 2RX4 1.2V 288-pin 8Gbit*. URL: <https://www.kingston.com/en/memory/search/discontinuedmodels?partid=KTH-PL429/32G> (visited on 12/20/2023).
- [58] Intel® Cooperation. *Ethernet-Converged-Network-Adapter X710-DA4*. URL: <https://ark.intel.com/content/www/de/de/ark/products/83965/intel-ethernet-converged-network-adapter-x710da4.html> (visited on 12/20/2023).
- [59] Ltd. Samsung Electronics Co. *PM9A3 NVMe® U.2 7.68TB*. URL: <https://www.samsung.com/us/business/computing/memory-storage/enterprise-solid-state-drives/pm9a3-nvme-u-2-ssd-7-6tb-mz-ql27t600/> (visited on 12/20/2023).
- [60] Lenovo (Deutschland) GmbH. *ThinkPad E14 Gen 2 (Intel)*. URL: <https://www.lenovo.com/de/de/p/laptops/thinkpad/thinkpade/e14-g2/22tpe14e4n2> (visited on 12/20/2023).
- [61] Analog Devices Inc. *Quad I2C Voltage, Current and Temperature Monitor*. URL: <https://www.analog.com/en/products/ltc2990.html> (visited on 01/08/2024).

Bibliography

- [62] Josep F Oliver et al. “Comparison of coincidence identification techniques for high resolution PET”. In: *2008 IEEE Nuclear Science Symposium Conference Record*. IEEE. 2008, pp. 4732–4735.
- [63] Fred James and Matts Roos. “MINUIT: a system for function minimization and analysis of the parameter errors and corrections”. In: *Comput. Phys. Commun.* 10.CERN-DD-75-20 (1975), pp. 343–367.
- [64] Roger Fletcher. “A new approach to variable metric algorithms”. In: *The computer journal* 13.3 (1970), pp. 317–322.
- [65] Frederick James Rene Bruns. *ROOT reference guide: TMinuit.cxx source code, Implementation in C++ of the Minuit package written by Fred James*. URL: https://root.cern.ch/doc/master/TMinuit_8cxx_source.html#l03992 (visited on 01/25/2024).
- [66] Sergey Bochkhanov and Vladimir Bystritsky. *ALGLIB-a cross-platform numerical analysis and data processing library*. 2011. URL: <http://www.alglib.net>.
- [67] Christopher C Paige and Michael A Saunders. “LSQR: An algorithm for sparse linear equations and sparse least squares”. In: *ACM Transactions on Mathematical Software (TOMS)* 8.1 (1982), pp. 43–71.
- [68] Shan Tong, Adam M Alessio, and Paul E Kinahan. “Image reconstruction for PET/CT scanners: past achievements and future challenges”. In: *Imaging in medicine* 2.5 (2010), p. 529.
- [69] Kris Thielemans et al. “STIR: software for tomographic image reconstruction release 2”. In: *Physics in Medicine & Biology* 57.4 (2012), p. 867.
- [70] Adam Alessio and Paul Kinahan. “PET image reconstruction”. In: *Nuclear Medicine, 2nd Edition* (2006).
- [71] Johann Radon. “1.1 über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten”. In: *Classic papers in modern diagnostic radiology* 5.21 (2005), p. 124.
- [72] D Qiu and E Seeram. “Does iterative reconstruction improve image quality and reduce dose in computed tomography”. In: *Radiol Open J* 1.2 (2016), pp. 42–54.
- [73] Thorsten M Buzug. “Computed tomography”. In: *Springer handbook of medical technology*. Springer, 2011, pp. 311–342.
- [74] Lawrence A Shepp and Yehuda Vardi. “Maximum likelihood reconstruction for emission tomography”. In: *IEEE transactions on medical imaging* 1.2 (1982), pp. 113–122.
- [75] Artur Słomski et al. “3D PET image reconstruction based on the maximum likelihood estimation method (MLEM) algorithm”. In: *Bio-Algorithms and Med-Systems* 10.1 (2014), pp. 1–7.
- [76] Belzunce MA, Verrastro CA, Cohen IM, et al. “Cuda parallel implementation of image reconstruction algorithm for positron emission tomography”. In: *The Open Medical Imaging Journal* 6.1 (2012).

- [77] Anastasios Gaitanis et al. “PET image reconstruction: A stopping rule for the MLEM algorithm based on properties of the updating coefficients”. In: *Computerized Medical Imaging and Graphics* 34.2 (2010), pp. 131–141.
- [78] Robert L Siddon. “Fast calculation of the exact radiological path for a three-dimensional CT array”. In: *Medical physics* 12.2 (1985), pp. 252–255.
- [79] Michel Defrise, Paul E Kinahan, and Christian J Michel. “Image reconstruction algorithms in PET”. In: *Positron emission tomography: basic sciences*. Springer, 2005, pp. 63–91.
- [80] H Malcolm Hudson and Richard S Larkin. “Accelerated image reconstruction using ordered subsets of projection data”. In: *IEEE transactions on medical imaging* 13.4 (1994), pp. 601–609.
- [81] A Michael Morey and Dan J Kadrmas. “Effect of varying number of OSEM subsets on PET lesion detectability”. In: *Journal of nuclear medicine technology* 41.4 (2013), pp. 268–273.
- [82] K. Thielemans et al. *STIR: Software for Tomographic Image Reconstruction User’s Guide Version 6.0*. URL: <https://stir.sourceforge.net/documentation/STIR-UsersGuide.pdf> (visited on 03/23/2024).
- [83] M Jacobson et al. “Enhanced 3D PET OSEM reconstruction using inter-update Metz filtering”. In: *Physics in Medicine & Biology* 45.8 (2000), p. 2417.
- [84] Patrick Hallen, David Schug, and Volkmar Schulz. “Comments on the NEMA NU 4-2008 standard on performance measurement of small animal positron emission tomographs”. In: *EJNMMI physics* 7.1 (2020), pp. 1–20.
- [85] Paul E Kinahan, Bruce H Hasegawa, and Thomas Beyer. “X-ray-based attenuation correction for positron emission tomography/computed tomography scanners”. In: *Seminars in nuclear medicine*. Vol. 33. 3. Elsevier. 2003, pp. 166–179.
- [86] Rozenn Le Goff-Rougetet et al. “Segmented MR images for brain attenuation correction in PET”. In: *Medical imaging 1994: image processing*. Vol. 2167. SPIE. 1994, pp. 725–736.
- [87] RD Badawi and PK Marsden. “Developments in component-based normalization for 3D PET”. In: *Physics in medicine & biology* 44.2 (1999), p. 571.
- [88] Dale L Bailey et al. “An investigation of factors affecting detector and geometric correction in normalization of 3-D PET data”. In: *IEEE Transactions on Nuclear Science* 43.6 (1996), pp. 3300–3307.
- [89] Edward J Hoffman et al. “PET system calibrations and corrections for quantitative and spatially accurate images”. In: *IEEE Transactions on Nuclear Science* 36.1 (1989), pp. 1108–1112.
- [90] Robley Duglison Evans and RD Evans. *The atomic nucleus*. Vol. 582. McGraw-Hill New York, 1955, pp. 791–793.

Bibliography

- [91] Edward J Hoffman et al. “Quantitation in positron emission computed tomography: 4. Effect of accidental coincidences.” In: *Journal of computer assisted tomography* 5.3 (1981), pp. 391–400.
- [92] Josep F Oliver and M Rafecas. “Modelling random coincidences in positron emission tomography by using singles and prompts: a comparison study”. In: *PloS one* 11.9 (2016), e0162096.
- [93] Charles C Watson. “New, faster, image-based scatter correction for 3D PET”. In: *IEEE Transactions on Nuclear Science* 47.4 (2000), pp. 1587–1594.
- [94] C Tsoumpas et al. “Evaluation of the single scatter simulation algorithm implemented in the STIR library”. In: *IEEE Symposium Conference Record Nuclear Science 2004*. Vol. 6. IEEE. 2004, pp. 3361–3365.
- [95] Kálmán Nagy et al. “Performance evaluation of the small-animal nanoScan PET/MRI system”. In: *Journal of Nuclear Medicine* 54.10 (2013), pp. 1825–1832.
- [96] Istvan Szanda et al. “National Electrical Manufacturers Association NU-4 performance evaluation of the PET component of the NanoPET/CT preclinical PET/CT scanner”. In: *Journal of nuclear medicine* 52.11 (2011), pp. 1741–1747.
- [97] Minjee Seo et al. “Performance evaluation of SimPET-L and SimPET-XL: MRI-compatible small-animal PET systems with rat-body imaging capability”. In: *EJN-*MMI physics** 10.1 (2023), p. 16.
- [98] Willy Gsell et al. “Characterization of a preclinical PET insert in a 7 tesla MRI scanner: beyond NEMA testing”. In: *Physics in Medicine & Biology* 65.24 (2020), p. 245016.