

Medipix3:
Spectral X-ray micro-CT and hadron therapy



Navrit J. S. Bal

Medipix3: Spectral X-ray micro-CT and hadron therapy

Medipix3: Spectral X-ray micro-CT and hadron therapy

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. ir. P.P.C.C. Verbeek

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen in de Aula der Universiteit

op woensdag 8 november 2023, te 14:00 uur

door NAVRIT JOHAN SINGH BAL

geboren te London

Promotiecommissie

| | | |
|--------------------|--|--|
| <i>Promotores:</i> | prof. dr. F.L. Linde dr. M. Fransen | Universiteit van Amsterdam Nikhef NWO-I |
|--------------------|--|--|

| | | |
|-----------------------|---|--|
| <i>Overige leden:</i> | prof. dr. A.P. Colijn prof. dr. ir. P.J. de Jong dr. ing. M.G. van Beuzekom prof. dr. S.C.M. Bentvelsen dr. J.M. Sonneveld dr. J.A. Nijkamp prof. dr. R.M.A. Heeren | Universiteit van Amsterdam Universiteit van Amsterdam Nikhef NWO-I Universiteit van Amsterdam Universiteit van Amsterdam Aarhus University Maastricht University |
|-----------------------|---|--|

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Copyright © 2023 by Navrit J. S. Bal

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission from the author.

ISBN/EAN 978-94-6419-910-9



This work is part of the Netherlands Organisation for Scientific Research (NWO). It was carried out at the National Institute for Subatomic Physics (Nikhef) in Amsterdam, the Netherlands and at the European Organisation for Nuclear Research (CERN) in Geneva, Switzerland.



This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 675265 - OMA (Optimization of Medical Accelerators).

Contents

| | |
|---|------------|
| ACRONYMS & GLOSSARY | II |
| I INTRODUCTION | 13 |
| 1.1 Context | 13 |
| 1.2 X-rays | 14 |
| 1.2.1 Light-matter interaction types | 15 |
| 1.3 Proton and heavy ions | 18 |
| 1.4 Radiotherapy | 21 |
| 2 MEDIPIX3 | 25 |
| 2.1 Overview | 25 |
| 2.2 Technical specifications | 27 |
| 2.3 Pulse processing chain | 28 |
| 2.4 Charge Summing Mode (CSM) | 29 |
| 2.5 Energy resolution | 31 |
| 2.6 Figures of Merit | 35 |
| 2.6.1 Signal to noise ratio (SNR) | 35 |
| 2.6.2 Count rate linearity | 35 |
| 2.6.3 Threshold-energy calibration | 36 |
| 2.6.4 System stability | 38 |
| 2.6.5 Summary | 56 |
| 3 PROTONS AND HEAVY IONS | 57 |
| 3.1 Ion identification and the volcano effect at HIMAC, Japan | 57 |
| 3.2 Beam profile measurements at Clatterbridge Cancer Centre - Wirral NHS (National Health Service), UK | 61 |
| 3.3 Clinical fluence measurements at MedAustron, Austria | 85 |
| 4 X-RAY COMPUTED TOMOGRAPHY | 107 |
| 4.1 Introduction | 107 |
| 4.2 Experimental setup | 114 |
| 4.3 Sinograms | 121 |
| 4.4 Artefacts | 123 |
| 4.4.1 Beam hardening | 123 |
| 4.4.2 Streaking artefacts | 123 |
| 4.4.3 Cupping artefacts | 124 |
| 4.4.4 Ring artefacts | 125 |
| 4.4.5 Other artefacts | 129 |
| 5 SPECTRAL X-RAY IMAGING | 131 |
| 5.1 Linear attenuation coefficient over energy | 131 |

| | | |
|-------|--|-----|
| 5.2 | Soft tissue x-ray contrast | 132 |
| 5.3 | Energy calibration | 136 |
| 5.4 | XRF Experiments | 146 |
| 5.4.1 | Silicon Drift Detector | 146 |
| 5.4.2 | Si 110 μm pixel pitch | 147 |
| 5.4.3 | GaAs 55 μm pixel pitch | 148 |
| 5.5 | Energy calibrated images (ECI) | 149 |
| 5.5.1 | Introduction | 149 |
| 5.5.2 | Methods | 150 |
| 5.5.3 | Measurement overview | 153 |
| 5.5.4 | Results and discussion | 153 |
| 5.5.5 | Summary | 166 |
| 5.5.6 | Existing efforts | 166 |
| 5.6 | Summary | 167 |
| 6 | SPECTRAL X-RAY CT APPLIED: RE-EXCISION AFTER POSITIVE MARGIN | 169 |
| 6.1 | Current clinical practice | 169 |
| 6.2 | Human breast tumour - 2018 | 171 |
| 6.3 | Mouse breast tumour - 2021 | 172 |
| 6.3.1 | Measurements and results | 173 |
| 6.4 | Summary | 181 |
| 6.5 | Outlook | 181 |
| 7 | SUMMARY AND OUTLOOK | 183 |
| 7.1 | Executive summary | 183 |
| 7.2 | Summary | 184 |
| 7.3 | Outlook | 186 |
| | BIBLIOGRAPHY | 196 |
| | APPENDIX A DEXTER - READOUT AND CONTROL SOFTWARE | 197 |
| | APPENDIX B DRIFT RATES PER PIXEL | 201 |
| | APPENDIX C MEDAUSTRON MEASUREMENT OVERVIEW | 205 |
| | APPENDIX D EQUALISATION | 209 |
| | APPENDIX E COLUMN BASED FILTERING ALGORITHM | 213 |
| | SCIENTIFIC OUTPUT | 215 |
| | SUMMARY | 219 |
| | SAMENVATTING | 221 |
| | ACKNOWLEDGEMENTS | 224 |

I have a lot of problems, but few of them are mine.

Marjolein L. E. Heidotting, 23:00 09 August 2019

ACRONYMS

ASIC: Application-specific integrated circuit
CSM: Charge summing mode
DAC: Digital to analogue
CT: Computed tomography
ECI: Energy calibrated images
ECI_FFC: Energy calibrated and then flat field corrected images
FFC: Flat field corrected
GaAs: Gallium arsenide
keV: Kilo electron volt
kVp: Peak x-ray acceleration voltage
LET: Linear energy transfer
Mo: Molybdenum
Si: Silicon
SNR: Signal-to-noise ratio
SPM: Single pixel mode
XRF: X-ray fluoroscopy

GLOSSARY

55 μ m pixels: Single pixels. There are one or two thresholds.
110 μ m pixels: Four pixels combined to read out one 110 μ m detector element. There are four or eight thresholds.
Chip: An ASIC.
Detector: The complete device; a hybrid pixel detector (e.g. Medipix3) combined with a sensor, such as Si or GaAs.
Hybrid pixel detector: An electronic device where the sensor layer is separated from the pixel electronics (hybrid), has multiple pixels (pixel) and the detects radiation (detector).
Medipix3: The ASIC used in this work which records the electronic signal from the sensor.
Sensor: The material used to convert ionising radiation to electronic signals which are measured by the connected ASIC. For example, it could be made of Si, GaAs and others.
Threshold: A configurable voltage which determines if an electronic signal is counted or not. This voltage is proportional to the measured signal.

1

Introduction

1.1 CONTEXT

Cancer is still a leading cause of death globally at a rate of 17% in 2017 despite all of the research and development (1). In 2020, global oncology drug spending totalled \$164 billion, an increase of almost a factor of two since just 2015 (2). Improvements are constantly being made in both diagnosis and treatment in terms of speed, cost and ultimately patient outcomes. Technological developments are a significant component of these improvements.

Novel radiation imaging and radiation treatment technology used in medical contexts is partly driven by HEP (High Energy Physics) experiments. This includes CERN (European Council for Nuclear Research) and in turn, the LHC (Large Hadron Collider). In the field of radiation detection, these experiments have the most demanding requirements which forces them to innovate directly. These requirements include the need to detect single particles, at lower energies and at higher rates than ever before. The technologies that result from this can be commercialised, potentially leading to off-the-shelf products. These products can then be integrated into clinical workflows and improve them accordingly.

The first particle detectors, photographic film and scintillation screens, were developed more than 100 years ago. The first electronic particle detector was the Geiger-Müller tube. This developed into strip and gas detectors such as VELO (Vertex Locator) sub-detector and GEMpix. R&D (research and development) in HEP experiments has directly resulted in an entire family of single particle counting hybrid pixel chips such as the Medipix and Timepix series and a dedicated development group at CERN, the Medipix group.

1. INTRODUCTION

1.2 X-RAYS

I

Photons are fundamental particles, they cannot be subdivided into other sub-particles. They are the quanta of the electromagnetic field, they carry the electric and magnetic fields, they bind electrons to atomic nuclei, form chemical bonds and they cover everything from radio waves to visible light and up to gamma radiation. X-rays are relatively high energy photons and are defined as having energies between approximately 10^2 to 10^5 eV with corresponding wavelengths of approximately 10^{-9} to 10^{-12} m.

X-rays are well known in both the scientific community and the general public. X-rays were accidentally discovered in 1895 by Wilhelm Röntgen. His seminal paper 'On a New Kind of Rays' was published in the following year (3). He detailed various properties of x-rays such as the 'absorption by various bodies must depend chiefly on their density'. He produced the now famous radiographs (x-ray images) of both his wife's hand wearing a ring and the lesser known image of the compass; the original figures with captions are in figure 1.1.



FIG. 1.—Photograph of the bones in the fingers of a living human hand.
The third finger has a ring upon it.




FIG. 2.—Photograph of a compass card and needle completely
enclosed in a metal case.

Figure 1.1: The first published radiographs made by Wilhelm Röntgen in 1895.

He started observing x-rays with fluorescence screens which convert x-ray photons to visible photons. Fluorescence screens only emit visible light for the duration of x-ray irradiation and so cannot be used to produce a permanent x-ray image. He then moved onto photographic dry plates to produce permanent x-ray images, such as the compass. These and subsequent imaging methods all produced greyscale images, intensity distributions.

X-rays have highly energy dependent absorption (cross section). Typical x-ray sources use accelerated electrons which are focused onto a metal target; the electrons interact with the charges (nuclei and bound electrons) of the metal target with a net result of the incident electron slowing down, changing direction and emitting a photon. Due to the conservation of momentum, the en-

ergy of that photon is equal to the energy lost by the electron being deflected. This is known as bremsstrahlung (braking radiation in German). X-ray sources utilising this effect produce a broad energy distribution with a maximum x-ray energy matching the acceleration voltage of the incoming electrons, typically 50 to 150 keV. This maximum photon energy is referred to as the peak voltage, kVp.

Therefore, the primary method used to change contrast in an x-ray image was to alter the x-ray spectrum (energy distribution). This can be done in many ways, including using thin metal filters and changing the electron acceleration voltage. Using metal foils reduces the intensity of the incoming beam significantly; more than 90% intensity reduction is common. Changing the electron acceleration voltage and therefore the peak x-ray energy also alters the shape of the spectrum, resulting in a change of contrast in a given image. These techniques are necessary because the methods for x-ray imaging are not energy sensitive, this is equivalent to black and white photography where the colour information is reduced to shades of grey.

1.2.1 LIGHT-MATTER INTERACTION TYPES

There are many possible interactions between photons and matter as can be found in textbooks and in 'Light-matter interactions with photonic quasi-particles' (4). The relevant interactions for x-rays of the energies used in this work (1—100 keV) are photoelectric, inelastic (Compton) scattering and elastic (Thomson / Rayleigh / coherent) scattering. A refresher of these interactions follows.

The photoelectric interaction is where a photon transfers all of its energy to a bound electron resulting in an ejected photoelectron. The photon must have at least the minimum amount of energy required to induce photo-emission of electrons for that material, this is referred to as the work function. The photoelectric effect dominates at low energies.

Inelastic scattering, also known as Compton scattering, is where the incident photon transfers part of its energy to an electron and is scattered in a different direction with a lower energy (longer wavelength). The three dimensional angular distribution is dependent on energy as described by the Klein–Nishina equation. Inelastic scattering becomes dominant above $\simeq 100$ keV.

Elastic scattering is where the photon and electron energy remain unchanged after the scattering. The incident photon transfers its energy to an electron, which is briefly excited and then relaxes and re-emits the photon with the same energy as before the interaction with an angular distribution equal over all angles. Elastic scattering is most significant at lower energies but does not dominate at any energy in the range of interest for this work. This is shown for copper in figure 1.2.




Figure 1.2: The photon cross section for copper from 0.1 to 500 keV. The vertical lines are various transition energies, also known as absorption edges and a subset of these are plotted. Absorption edges are where the probability of photoelectric interactions (photoionisation) increases sharply at an energy level just beyond the binding energy of a specific electron shell (e.g. K, L, M) of the absorbing atom. These energies are plotted as vertical dashed lines at 0.93, 0.95, 8.04 and 8.91 keV, the first two are L-edges and the second two are K-edges. The solid lines are listed in the legend in terms of descending average cross section. The library ‘xraylib’ was used for the data (5) via the Python bindings.

For this example with copper in figure 1.2, a photon of at least the K edge energy (8.91 keV) is required to interact via the photoelectric effect to eject a K shell (principal quantum number, $n = 1$) electron. Below this characteristic minimum energy, the photon will be absorbed and re-emitted.

The photoelectron proceeds by one of the following three methods:

- The electron absorbs all of the photon energy. It is completely liberated from the atom and becomes a free electron. Any energy above the binding energy contributes towards the kinetic energy of the photoelectron;
- ‘Auger effect’: An electron from a higher energy state can occupy the vacancy left from the photon, which knocks out an outer shell electron, this is an ‘Auger electron’ (6);
- ‘Coster-Kronig’ transition: A special case of the Auger effect where a vacancy is filled by an higher energy electron within the same shell (same principal quantum number), for example from a L_3 to L_1 state. The orbital angular momentum (ℓ), magnetic moment (m_ℓ) and electron spin (m_s) quantum numbers can be different.

The probability for which of these three methods actually occurs varies with the electronic configuration of the atom in question. These are the fluorescence, Auger and Coster-Kronig yields and contribute to the relative intensities of the fluorescence lines as observed with XRF (X-ray Fluorescence) measurements.

Since pair production (nuclear interactions) is only relevant at twice the rest mass energy of an electron ($2 \times 511\text{keV}$) and above, this is not relevant for this work. At most, 90 keV x-ray photons are used.

As a rule of thumb, the photon cross section is proportional to the fourth to fifth power of the atomic number, Z , also known as the proton number, divided by the energy to the power of 3.5, E . $\sigma \propto \frac{Z^{4 \text{ to } 5}}{E^{3.5}}$.

Transmission x-ray imaging is the dominant x-ray imaging method used due its simplicity and highly efficient use of the x-ray radiation. In this case, the photoelectric interaction is the most dominant interaction. It switches to inelastic scattering after a certain energy. This is widely used for medical CT (Computed Tomography) scanning in hospitals and for non-destructive inspection within objects. It is typically limited to $\simeq 150\text{ keV}$ due to the inelastic scattering becoming dominant around that energy; this scattering effectively introduces blurring to images. Anti-scatter grids help with this phenomenon to some extent. The dominant process depends on the atomic number, Z , and the photon energy. This is shown in figure 1.3.




Figure 1.3: The dominant interaction mechanisms between photons and matter for a wide range of energies and atomic numbers, Z , of the matter.

The relative insignificance of scattering enables very high spatial resolution with transmission imaging. With modern synchrotrons the spatial resolution achievable can reach the nanometre scale; even enabling nano-tomography (7). Typically, the energy dependent information is effectively discarded. The notable exception is dual energy CT, where the peak voltage of the x-ray tube

is varied and the alternating images are subtracted from each other. This work explores the information available in the energy space and how it propagates to x-ray micro computed tomography.

As a side note, dark field and phase contrast imaging both use the wave properties of photons to reveal areas of high attenuation variation, such as the interface between lung tissue and the air within them (8; 9). This is an area of active research for medical applications.

Yet another method to extract information with the photoelectric interaction is x-ray fluorescence measurements. This is where an incident photon beam interacts with atoms whose electrons have characteristic binding energies which are material specific, resulting in re-emitted photons of these specific energies in all directions. This measurement method has low quantum efficiency. In this work, these very specific energies are used for determining the response of every pixel to one very narrow energy distribution, less than one eV width. This is referred to as a threshold-energy calibration within this work. In terms of other methods to take advantage of this behaviour, scanning a micro-focused x-ray beam across a sample combined with a high energy resolution detector allows reconstructions of objects with high spatial and energy resolution. This technique is effectively limited to synchrotrons due to the excellent spatial resolutions achievable combined with among the highest available x-ray flux.

Compton back-scattering imaging takes advantage of the inelastic (Compton) scattering allowing the imaging of highly attenuating and thick samples including steel rebar and concrete and even through a thick and dense wall (8 mm of steel) (10; 11). The typical range for this method is $\simeq 100$ — 1022 keV. It typically uses collimators and anti-scatter grids and has become used in non-destructive screening in seaports for cargo containers and for explosives, weapons and personnel screening at airports (12).

1.3 PROTON AND HEAVY IONS

Hadrons are composite particles of at least two quarks bound by the strong interaction, this includes protons and ions. Ions are atoms or molecules with a net charge. They are then stripped of their electrons, resulting in protons and nuclei respectively. These are then accelerated to higher energies and used for various scientific and medical purposes such as cancer treatment.

Charged particles such as protons and heavy ions interact in many ways with matter. How do protons and heavy ions deposit energy into a material? The Bethe-Bloch equation answers this question; it gives the mean rate of energy loss (stopping power) of a heavy charged particle as described by Bethe and Bloch in 1932 (13):

$$-\left\langle \frac{dE}{dx} \right\rangle = 2\pi N_A r_e^2 m_e c^2 \frac{Z}{A} \frac{z^2}{\beta^2} \left\{ \ln \left(\frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} \right) - 2\beta^2 - \delta - \frac{2C}{Z} \right\} \quad (1.1)$$

The definitions of the many terms are:

- $N_A = 6.022 \times 10^{23} \rightarrow$ Avogadro's number;

- $r_e = \frac{e^2}{4\pi\epsilon_0 m_e c^2} = 2.8 \text{ fm} \rightarrow$ electron radius;
- $m_e = 511 \text{ keV} \rightarrow$ electron rest mass;
- $Z \rightarrow$ atomic number (proton number) of the absorbing material;
- $A \rightarrow$ atomic mass number of the absorbing material;
- $z \rightarrow$ charge of an incident particle in units of electrons, e;
- $\beta = \frac{v}{c} \rightarrow$ The velocity of the particle divided by the speed of light;
- $\gamma = \frac{1}{\sqrt{1-\beta^2}} \rightarrow$ The Lorentz factor
- $\delta \rightarrow$ density correction - relevant for high energies;
- $C \rightarrow$ shell correction - relevant for low energies;
- $W_{max} \rightarrow$ maximum energy transfer in a single collision;
- $I \rightarrow$ mean excitation energy of the absorbing material.

An example relevant to this work follows. Using the NIST (The U.S. Department of Commerce's National Institute of Standards and Technology) PSTAR (Stopping Power and Range Tables for Protons) database, we can plot the mean rate of energy loss of a 60 MeV proton in silicon using the Bethe-Bloch equation, see figure 1.4. This shows how the stopping power due to electronic interactions are at least two orders of magnitude higher than nuclear interactions over the energy range plotted. The total stopping power is the sum of the electronic and nuclear stopping powers.




Figure 1.4: The mean rate of energy loss of protons in silicon using the Bethe-Bloch equation. The total energy loss is dominated by the electronic contributions over the displayed energy scale.

As a hadron passes through material, it transfers energy via the electromagnetic and nuclear interactions to the atoms it passes by. The total interaction cross section (probability) increases as the remaining hadron energy approaches zero. This results in a sudden increase in the rate of energy deposition per unit length (stopping power, dE/dx) after a relatively constant dose. This is known as the ‘Bragg peak’, it is the critical feature of how protons interact with matter that makes it useful for cancer treatment. The dose deposition of x-rays and gamma photons exponentially decreases as a function of depth. The dose depositions for protons, carbon ions and x-ray photons can be seen in figure 1.5.




Figure 1.5: The relative energy depositions of x-rays (photons), protons and carbon ions. Photons exponentially lose energy whereas protons and ions deposit most of their energy at an energy dependent range resulting in the ‘Bragg peak’. Protons and ions can therefore spare surrounding healthy tissue much better than photons. Note that the typical energy range of the x-rays used for radiotherapy are in the range of 300 keV to 1.5 MeV. Image from (14).

Due to the nuclear interactions, hadrons can cause radiation damage to radiation detectors via various mechanisms, which cause microscopic defects. This is an area of active research; it is particularly relevant for high energy physics, radiation detection in space, industrial and medical contexts (15–19). As transistors shrink, they become increasingly susceptible to radiation induced disruption or damage. This affects the hybrid pixel detectors, the type used in this work, for both the chips themselves and their sensors.

The Bragg peak allows radiation oncologists to make treatment plans that spare surrounding healthy tissue far better than with x-ray and gamma photons while still being able to treat tumours.

Linear energy transfer (LET) is a related important quantity, it is the average energy deposition per unit length along the path of an ionising particle; the units are typically keV/ μm . LET is proportional to the particle charge squared and inversely proportional to the particle’s kinetic energy for non-relativistic particles. It is a relevant quantity for estimating the biological damage per unit volume.

1.4 RADIOTHERAPY

Radiation therapy, also commonly known as radiotherapy, utilises ionising radiation and is one of many cancer treatment methods. Typically high energy photons, x-rays or gamma photons, are

used because they are relatively easy to generate and it is possible to deposit dose anywhere within a human body without invasive procedures. Dose is defined as the absorption of one joule of radiation energy per kilogram of matter; the unit is 'Gray'. The most common method to generate x-rays and gamma photons is with an accelerated and focused electron beam striking a high Z target resulting in photon emission via bremsstrahlung.

In clinical practice, treatment plans must be made with the dose depositions of the particle used in mind. For photons, there is the exponential decay and for hadrons there is the Bragg peak. The ideal treatment would only affect the tumour and spare the healthy surrounding tissue entirely. Actual treatment plans usually target the photon beams at the tumour from multiple angles, this is known as intensity-modulated radiation therapy (IMRT). While the ratio of dose of tumour to healthy tissue is favourable, IMRT does deposit dose in healthy tissue. Figure 1.6 shows the differences between IMPT (intensity modulated proton therapy) and conventional IMRT and the differences between them, the 'unnecessary radiation'. The unnecessary radiation visually indicates how much healthy surrounding tissue can be spared of irradiation. Note how the IMPT (Proton) treatment plan has minimal dose to the brain, brain stem, eyes, nose and spinal column. This is one major benefit of hadron therapy over conventional radiotherapy using photons.

Depending on the rate of cell division, human tissues have variable levels of radiosensitivity. The least sensitive are nerves and the most sensitive include bone marrow, lymphatic tissue, sperm cells and fetal tissue. This results in variable radiation induced side effects for the patient, anything from manageable to clinically unacceptable.




Figure 1.6: Two treatment plans using intensity modulated proton therapy (IMPT Protons) versus intensity modulated radiotherapy (IMRT Photons) and the differences between the two. The colour scale goes from red to blue where red is the maximum dose and blue is the minimum. This particular tumour is in the upper part of the throat behind the nose, the ‘nasopharynx’. This shows how conventional IMRT deposits significantly more unnecessary dose to surrounding healthy tissue compared to IMPT. Image from (20), accessed on 07-02-2022.

All methods using high energy photons suffer this effect, often this results in secondary cancers, cardiovascular disease and fertility complications among others (21). Therefore radiation oncologists are never able to exclusively irradiate a tumour. This is particularly an issue when it is necessary to irradiate tissue that is or is near very radiosensitive tissue that cannot be operated on otherwise.

Other types of ionising radiation are used in radiation oncology. These include alpha and beta emitters, hadrons (protons & heavy ions) and even focused very high energy electrons (VHEE). They are complementary treatment methods due to their different dose deposition properties. Hadron therapy is a growing modality for cancer treatment using hadrons, which in this context only refers to protons and heavy ions. Hadrons are composed of protons and neutrons and therefore the interaction mechanisms are not the same as x-rays and gamma photons.

Proton and heavy ion radiotherapy centres are the direct result of high energy particle accelerator research. The requirements for medical accelerators are quite different to those in high en-

ergy physics research. The precision required is lower, the accelerators must be smaller, the reliability and predictability must be higher and the cost has to be much lower. There is a general drive to reduce the size, complexity and cost of medical accelerators to ramp up their availability to the general population. There are three clinical centres relevant for this work, Clatterbridge Cancer Centre (CCC) (United Kingdom), Heavy Ion Medical Accelerator in Chiba (HIMAC) (Japan) and MedAustron (Austria). These facilities vary significantly in many metrics such as age, particles available, energy ranges and clinical focus. As of September 2021 according to PTCOG (Particle Therapy Co-Operative Group), there are 112 hadron therapy centres in operation, 41 under construction and 15 in the planning stage. As of 2019, over 220,000 patients have been treated with protons and nearly 40,000 using carbon ions.

MedAustron is the direct output of a CERN project, PIMMS (Proton Ion Medical Machine Study) (22) and is a CE (European conformity) certified medical product according to the EU (European Union) Medical Device Directive, '93/42/EEC'. It was realised at a cost of approximately € 200 million. An image of the particle accelerator is shown in figure 1.7.



Figure 1.7: A picture of the hall containing the medical particle accelerator used to treat patients with proton and carbon ions at MedAustron, Austria. Credits for the picture go to Dr. Thomas Kaestenbauer, 2013.

Navrit: I'm looking for fresh bugs
Kevin: I've got some in my code here.

Navrit Bal & Kevin Heijhoff, 13:37 25-02-2020

2

2

Medipix3

2.1 OVERVIEW

We now move onto the tool which can be used to measure x-ray photons, protons and heavy ions, Medipix3, a hybrid pixel detector chip.

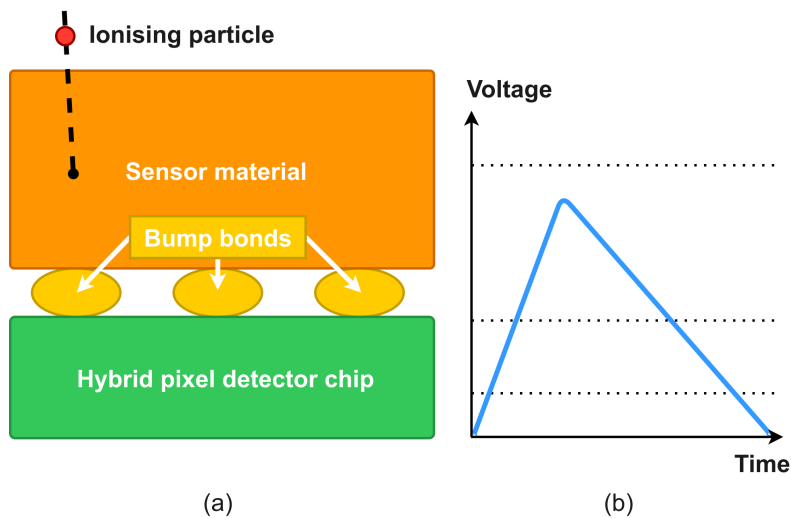


Figure 2.1: The most basic hybrid pixel detector overview (a) and a plot showing the signal (voltage) shape over time with the horizontal dashed lines visualising thresholds (b).

A hybrid pixel detector at the most basic level can be summarised as a device with a sensor material coupled to a hybrid pixel detector chip. Ionising particles interact within sensor material via various mechanisms. The sensor is often coupled to the chip with bump bonds as shown in figure 2.1a. The particle interactions release charge that is collected by the pixel electronics. Inside a pixel, the electronics convert this charge into a signal of a voltage over time as shown in figure 2.1b. It is counted if it crosses a configurable threshold, this threshold is illustrated as the horizontal dotted lines in figure 2.1b. If there are multiple threshold counters, it is possible for a signal to cross the lower threshold but not a higher threshold. In the example from 2.1b, the first and second thresholds would count once each and the third would not count. This strategy enables binning (discretisation) of a continuous energy spectrum into a number of bins at relatively high count rates with precisely the same input signal and noise. It is important to note that noise in the electronics limits the measurable precision of a given signal amplitude. The same noise also sets a limit to how low the threshold can be set before random hits due to noise are being measured.

Medipix₃ is a hybrid pixel detector chip which can be used with several sensor materials. Charge summing mode (CSM) is the main new feature introduced with this chip, it attempts to address a fundamental issue with this type of detector. Namely, that the induced charge is often deposited over more than one pixel. The pixel pitch, number of energy thresholds, readout modes and so on are all configurable depending on the specific user requirements. In many ways, Medipix₃ is like a Swiss army knife. It is designed to be flexible to address many different situations and is not required to excel in any specific application. This has made it a valuable research chip, which is well placed to inform spin-off chips for more specific tasks.

The detector used in this work consisted of a two by two layout of Medipix₃ chips, with a total of four chips. An image of the detector without a silicon sensor is shown in figure 2.2. The chip numbering scheme is shown in figure 2.3, this becomes relevant in section 5.3.




Figure 2.2: A Medipix3 based detector in a two by two layout without a silicon sensor. Each square is a single Medipix3 chip. They are mounted on a ‘chip board’ - the PCB (polychlorinated biphenyl) board with a red silkscreen and a gold coloured rim.

| | |
|---|---|
| 0 | 1 |
| 3 | 2 |

Figure 2.3: The chip numbers of a Medipix3 based detector in a two by two layout as decoded and saved by the Dexter readout and control software.

2.2 TECHNICAL SPECIFICATIONS

An overview of the technical specifications follows:

- Number of pixels: 256×256 (fine pitch mode, FPM) or $4 \times 128 \times 128$ (colour mode, CM) = 65 536 per chip;
- Pixel pitch: 55 μm (FPM) or 110 μm (CM);
- Energy thresholds: one or two for FPM and four or eight for CM;
- Electronic noise (design specification): 70—150 electrons, 253—542 eV with a Si sensor at room temperature;

- Typical operational temperature: $+10 - +60\text{ }^{\circ}\text{C}$;
- Readout modes: Sequential R/W (Read/Write) and Continuous R/W. Sequential R/W takes time to readout every frame but can use both counters, so two or eight thresholds are usable. Continuous R/W uses both counters in 1, 6 and 12 bit modes to enable zero dead time readout and so one or four thresholds are usable;
- Bit depth: 1, 6, 12 and 24 bits. 24 bit readout mode uses both counters and so one 24 bit frame can be readout in that mode.

2.3 PULSE PROCESSING CHAIN

When an x-ray photon hits the sensor and travels through the sensor bulk, if it interacts via the photoelectric effect, it deposits all of its energy into an electron, part of that energy is used to overcome the binding energy (as mentioned earlier) with the remaining energy left for further ionisations. The high energy electron excites electron-hole pairs, which drift to the front and backside of the sensor due to an applied electric field. The moving charge induces a current (Shockley-Ramo theorem) through the pixel pads (collection electrodes), which are reverse biased, depleted ‘pn’ junctions. The measured charge is proportional to the energy deposition of the particle. This is the start of the pulse processing chain within one pixel.

A member of the Medipix group, Rafael Ballabriga-Suñé goes into great detail on this subject in his thesis (23), it contains the most detailed and complete technical information including circuit block diagrams.

Depending on the chip configuration, the following chain varies. This example goes through the simplest case: Fine Pitch Mode (FPM, $55\text{ }\mu\text{m}$ pixel pitch) with Single Pixel Mode (SPM). When in SPM, every pixel is independent of all surrounding pixels. SPM involves no inter-pixel communication unlike Charge Summing Mode, CSM.

The following paragraph is a summary of section 3.3 in (23). The charge pulse enters the preamplifier (charge sensitive amplifier, CSA). It outputs a voltage pulse (V) with a magnitude equal to the charge (Q) divided by the pixel capacitance (C); this comes from rearranging the equation $Q = CV$. The preamplifier output is then processed by a first order semi-Gaussian shaper, which means that it takes that voltage pulse and outputs an exponentially decaying voltage pulse whose magnitude is proportional to the charge collected in that pixel.

This is then fed to two discriminators (low and high), which have configurable thresholds, which are set by their respective DACs (digital to analogue converters). Once a signal pulse rises above the set threshold, the discriminator outputs a digital voltage pulse to the pulse processing logic block. The implementation of this block is not publicly disclosed; it acts as an ADC (analogue to digital converter). This leads to two counters (low and high) which can increase by zero or one per count depending on whether the signal amplitude is above the threshold set for the low and high counters.

One can select an energy window by setting the high threshold above the low threshold and subtracting them from each other. In Charge Summing Mode (CSM) and/or Colour Mode (110 μm pixel pitch), there is inter-pixel communication within this chain. Four of these signal processing chains, from four pixels, are connected to the one active pixel input amplifier, resulting in a large pixel with eight thresholds. For the readers who seek electronic block diagrams, figure 2.4 helps to digest this section.

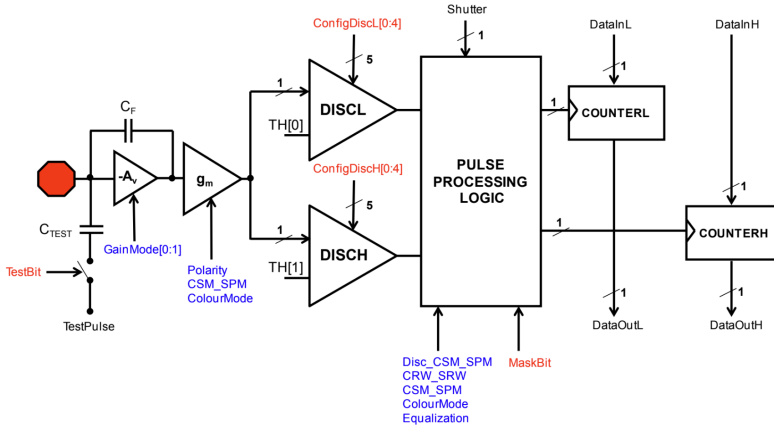


Figure 2.4: Figure seven from the Medipix3 manual (24). It is the circuit block diagram of the Medipix3 front-end electronics when the chip is in SPM (Single Pixel Mode) and Fine Pitch Mode (FPM, 55 μm pixel pitch).

2.4 CHARGE SUMMING MODE (CSM)

An overview of Charge Summing Mode (CSM) follows. This includes an explanation of the advantages and disadvantages of CSM. After that, there is a discussion of the specific conditions under which CSM is preferable over SPM for different sensor materials and energies.

CSM is one of the main design improvements over Medipix2. CSM is an on-chip solution to reconstruct and assign charge to the correct pixel over a 2×2 pixel area.

When an ionising particle enters a sensor material (e.g. silicon), electron-hole pairs are produced and charge drifts towards the pixel implants on the sensor which flows through the pixels in the chip. Focusing on photons now, if the photon interacts via the photoelectric effect, the photon is absorbed resulting in an electron with the photons' energy minus its binding energy. Most of the binding energy is released as auger electron (local energy deposition) or another fluorescence photon. The latter may be absorbed at a larger distance (contributing to charge sharing) or even escape. If the photon interacts via Compton scattering, not all of the photon's initial energy would have been deposited in the sensor and so the information about the photon's energy before the interaction is lost. Depending on the electric field either electrons or holes will drift towards the

chip. It is likely to deposit charge over more than one pixel. This particularly affects thicker sensors, high Z sensors and high LET particles such as electrons, protons, and ions. This is shown in Figure 2.5.

CSM aims to improve key imaging quality metrics (MTF (Modulation Transfer Function) and DQE (Detective Quantum Efficiency)) and the measured energy spectrum at the expense of increased power, increased front-end noise, lower gain and higher minimum detectable charge.




Figure 2.5: An overview of how ionising particles can interact with the sensor material via the photoelectric effect in hybrid pixel detectors. The electrons and holes drift to the cathode and anode depending on the polarity of the induced electric field.

CSM is implemented within the pulse processing logic using an arbitration network between adjacent pixels. The algorithm is described in the Medipix₃RX manual v1.4 as follows (24):

“

- The lower threshold in every pixel (or cluster) is compared to the locally deposited signal;
- The arbitration circuitry suppresses the pixels with a lower signal, identifying the pixel (or cluster) with the largest energy deposit. The information of the hit pixel after the allocation algorithm is stored in the lower counter (this information is referred to as “Single Pixel Mode Arbitrated” because the information is based on the local charge deposited in the pixel);
- In parallel, the charge has been reconstructed and compared with the highest threshold in the four adjacent summing nodes with respect to the pixel;
- The pixel with the highest charge checks the adjacent summing circuits to see if at least one of them exceeds the highest threshold.

CSM can be used with sequential RW readout mode only, not continuous RW mode. This is because both counters are needed for full event reconstruction for CSM due to a deficiency in the chip design.

CSM gain is 50% of SPM (Single Pixel Mode) gain. In this work, gain is defined as the number of threshold DAC units per energy unit, keV. This is because in SPM and SPMa, two current branches are routed to their respective comparators, ‘threshold low’ (THL) and ‘threshold high’ (THH). In CSM, four additional current branches are routed to the four summing nodes in the corners for every pixel.

‘Single pixel arbitrated’ means that for a given hit, the counter associated to THL increases if the signal is above THL and the signal is the largest in its neighbourhood. This is different from the traditional Single Pixel Mode whereby the counter increases only if the signal is above THL.

This is demonstrated directly in section 5.4.

The minimum detectable charge is also higher than SPM because the front-end noise is higher due to the noise components from four pixels adding in quadrature. For detecting x-rays with a silicon sensor in CSM, the minimum detectable photon energy is reported to be approximately 6–7 keV (25).

Count rate linearity is also affected with CSM since it takes time for the arbitration network to process events, whilst processing, it is effectively paralysed. Therefore, the counting rate for a given number of photons becomes non-linear earlier than it would for SPM. For CSM, it was found that the upper linear limit is approximately 8×10^6 counts/sec/mm² (26). This is completely independent of the frame rate.

For CSM and FPM, the pixel spacing during equalisation should be at least three not to interfere with the CSM algorithm which also operates on two by two pixel grids.

The differences between SPM and CSM have been studied with a GaAs sensor in detail (27).

2.5 ENERGY RESOLUTION

The following subsection is intended to refresh the more experienced reader and to give a brief introduction to the rest.

When an ionising particle such as an x-ray photon enters matter, it can interact by a number of mechanisms, as explained in section 1.2.1. The interacting particle’s energy is at least partially transferred to electrons, holes and phonons (quasiparticles / lattice vibrations) within the sensor material. The resulting electron-hole pair cloud starts to diffuse in all directions and since there is an electric field applied (bias voltage), the cloud splits into two components. The holes drift in one direction and the electrons drift in the other direction, along the applied electric field. Charge recombination of the holes and electrons is not a significant factor in solid state sensor materials until very high radiation levels due to their high charge carrier concentrations.

In the case of purely photoelectric interactions where all energy of the incoming particle is transferred to the fully depleted sensor, a certain number of electron hole pairs, N , are produced. The number of electron-hole pairs created is the ratio of the incident particle's energy to the mean energy required to produce one electron-hole pair.

$$N = \frac{E}{E_{pair}} \quad (2.1)$$

E_{pair} is approximately proportional to the material dependent band gap energy, $E_{band\ gap}$ (28).

$$E_{pair} \approx 2.8 \times E_{band\ gap} + 0.6\ eV \quad (2.2)$$

The Fano factor is defined as the ratio of the variance to the mean as described by Ugo Fano in 1947 (29). μ is the mean and σ is the standard deviation of the statistics of N .

$$F_N = \frac{\sigma_N^2}{\mu_N} \quad (2.3)$$

This factor appears in particle detection due the fact that the energy loss process is not purely statistical due the discrete electron shell energies and therefore there are a finite number of possible transitions between those states. A better energy resolution is therefore measured compared to the predicted value using Poisson statistics exclusively. Poisson statistics alone would predict unity. It is important to note that the Fano factor for particle detection can be quoted with a dependence on temperature, this is because the mean energy required to generate one electron-hole pair is a function of temperature because the band gap changes. For Si, the Fano factor was measured to be 0.118 ± 0.004 with an electron-hole pair energy of $3.73 \pm 0.09\ eV$ with a gradient of $-0.0131 \pm 0.0004\ \% K^{-1}$ between 80 and 270 K (30). In summary, due to the discreteness, the statistics become more favourable and which is summarised in the Fano factor; it is often much less than 1. This leads to a much more narrow distribution (better sigma/mean ratio).

The intrinsic energy resolution is related to the Fano Factor as follows as defined in (28), this definition will be used throughout this work. It becomes relevant from section 5.1. Where the intrinsic energy resolution is E_{FWHM} , FWHM is the Full Width at Half Maximum, E_{pair} is the material and temperature dependent mean energy required to produce one electron-hole pair and E is the mean measured energy deposition (28).

$$E_{FWHM} = 2.355 \times \sqrt{F_N \cdot E \cdot E_{pair}} \quad (2.4)$$

The intrinsic energy resolution, E_{FWHM} , is shown in figure 2.6 for a silicon sensor at 300 K; this is the lowest achievable energy resolution as a function of photon energy. At a relevant energy for this work, 10 keV, the intrinsic energy resolution is 154 eV; the relative energy resolution is therefore 1.54 %. As mentioned in subsection 2.2, the design value of the Medipix3 electronic noise is 253—542 eV with a Si sensor at room temperature. Under these conditions, the electronic noise is multiples

of the intrinsic energy resolution. In this configuration, the electronic noise therefore dominates the effective energy resolution.





Figure 2.6: The intrinsic energy resolution, E_{FWHM} , as a function of photon energy in keV. This assumes room temperature (300 K), a silicon sensor, with a Fano factor (F) = 0.118 and mean electron-hole pair generation energy, $E_i = 3.62$ eV.

In general, to optimise energy resolution one should minimise the Fano factor, use the semiconductor with the smallest band gap energy and minimise the incident particle energy. In reality, there are further considerations to be made; there are few semiconductors in existence with high quality and widespread manufacturing. There are material R&D efforts in progress to improve this situation. CVD (chemical vapour deposition) diamond research has been taking place in the RD42 collaboration for the HL-LHC (High-Luminosity LHC) (31) and other developments using Timepix3 (32). In Russia, GaAs:Cr (Chromium compensated Gallium Arsenide) has been progressing with multiple ASICs, the Medipix2 (33), HEXITEC (34), and a 1000 μm thick sensor bonded to Medipix3 chips, which I tested at Nikhef. This detector suffered from poor bump bonding but was better than the previous attempts.

An overview for the process of measuring the energy resolution with Medipix3 is shown in figure 2.7, it shows where terms derive from such as energy resolution, threshold dispersion and energy resolution variation. Ultimately, it reduces down to measuring a single energy input with a threshold scan, ideally this is an error function. One then numerically differentiates this distribution and calculates the FWHM. The FWHM ($\text{FWHM} = 2.355 \times \sigma$, where σ = standard deviation) encapsulates all effects mentioned and all other effects, such as electronic noise, charge sharing and background x-rays from any non-target material and so on. This is the number that is referred to as the energy resolution.



A threshold scan needs a radiation source, which is ideally monochromatic and has no energy variation. A threshold scan is a series of measurements where only the energy threshold is scanned over some range.

This results in a number of images, **N**. These are stored as **N** TIFF (Tag Image File Format) files. The resulting data is of the shape (**N**, **detector_columns**, **detector_rows**). "**detector_columns**" and "**detector_rows**" are equal and multiples of 128, this comes from the number of pixels per chip.

A region of interest (ROI) is chosen. Depending on the desired granularity and execution speed, one could select the whole detector, per chip or per pixel. The execution speed varies significantly between these options due to the parallelised operations used for large array calculations.

For the selected ROI(s), the counts as a function of threshold in DAC units can be plotted. Depending on the radiation source, a lower energy component may be present. If so, this will present as increasing counts after the main rise. This will also happen due to charge sharing between pixels and no active charge summing algorithm.

Alternatively, there could be no charge sharing between pixels and there is an active charge summing algorithm. In Medipix3, this is the Charge Summing Mode (CSM). Then, there will be a plateau from low energy (low threshold) up to the transition threshold (energy). This is the energy reference which corresponds to the mean energy of the ideally monochromatic radiation source. This is the vertical dashed yellow line.

This is then numerically differentiated and multiplied by -1 (vertically flipped) to produce the next plot.

A normal (Gaussian) distribution can be fit to this data. If there is a low energy tail, the data used for the fit can be restricted to the mean plus/minus a value. The mean and standard deviation of this fit are then extracted.

This process is repeated for the number of ROIs. This results in a number (**M**) of means and standard deviations of these fits, **M** \times (**μ, σ**).

This is split into two 1 dimensional arrays, one for μ and the other for σ . Further metrics can be extracted as follows with units typically in keV or electrons.

Energy reference = $\mu(\mu(\text{Th}))$
 Energy resolution = $\mu(\sigma(\text{Th}))$
 Threshold dispersion = $\sigma(\mu(\text{Th}))$
 Energy resolution variation = $\sigma(\sigma(\text{Th}))$

Figure 2.7: An overview of energy calibration measurements from the source to the final terms.

2.6 FIGURES OF MERIT

2.6.1 SIGNAL TO NOISE RATIO (SNR)

The signal to noise ratio has many definitions depending on the field. Since this work is focused on imaging performance, the definition used is the ratio of the mean signal squared to the standard deviation of the signal squared:

$$SNR = \frac{\mu^2}{\sigma^2} \quad (2.5)$$

In imaging, the signal would be pixel values within an area of constant input. For example, with x-rays this could be imaging a thin foil of equal thickness and composition perpendicular to the detector and source.

One may ponder at this point - how high does the SNR need to be for a ‘good’ image? One interpretation of that question is: how high does the SNR need to be for an object to be detectable? The answer to this question is the Rose criterion, which states that the SNR must be at least 3–5 to reliably detect an object (35). While it is a limited metric, the Rose criterion is used in the medical imaging field.

2.6.2 COUNT RATE LINEARITY

An ideal detector always counts exactly once per particle. In real photon counting detectors, this is not the case due to how the pixels in counting detectors have dead-time where they cannot count.

Suppose a single particle deposits all of its energy into one pixel and the deposited energy is above the minimum threshold. After a very short time, that one pixel counts once.

As the particle flux is increased, the pulses created by single particles begin to overlap each other. This overlapping both distorts the energy spectrum (pulse pile-up) and changes the relationship between incident particles and measured counts to be increasingly non-linear. After a certain flux, pixels will actually count less than the number of incident particles. The fraction that can be tolerated to be lost due to pile up sets the upper count rate limit for counting detectors. Paying close attention to count rate linearity is therefore critical to measuring an undistorted energy spectrum and a linear relationship between incident particle flux and counts.

This effect is ideally measured by using a radiation source which knows precisely how many particles it emits, can be focused down to a single pixel and has a wide flux range. Fortunately, synchrotrons are quite close to the described perfect source. X-ray tubes can also be used to estimate this effect, the main issues with using them is the lack of precise current measurements for the tube, untrustworthy and low precision current readings through the x-ray control software and they cannot be focused to single pixels. These deficiencies could of course be addressed by using an appropriate secondary x-ray detector as a cross check.

All counting detectors are affected by count rate linearity limitations, including the nuclear detection gold standard, HPGe (High Purity Germanium) detectors. Medipix3 offers a software configurable method to change the front-end response which can extend the upper linear count rate limit at the expense of noise and therefore energy resolution. For example, this is very useful if one is measuring a high flux of particles and the energy resolution is not relevant for that measurement. The method is a DAC for the Krummenacher current, I_{Krum} , which is proportional to the fall time of the signal at the expense of higher noise and controls the leakage current compensation. This is part of the Charge Sensitive Amplifier (CSA) as described in the Medipix3 manual (24).

This has been studied in some detail using x-rays at synchrotrons (36; 37) and with standard x-ray tubes, see section ‘6.7 Count Rate Linearity of the Medipix3 Detector’ (26). These studies yield numbers in the range of $\times 10^5$ — $\times 10^6$ counts per second per pixel and are consistent with each other. An example figure showing measurements of this effect comparing 55 and 110 μm pixel pitches with SPM and CSM is shown in figure 2.8 from (36).




Figure 2.8: A ‘comparison between the 55 μm and 110 μm detector’ for count rate linearities with various configurations: different dead times, I_{Krum} DAC settings, pixel pitches, with SPM and CSM. Figure 3 from (36).

2.6.3 THRESHOLD-ENERGY CALIBRATION

For any energy dependent study, it is important to know the mapping between the thresholds set in DAC units and energy. Medipix3 was designed to have a linear relation between threshold and energy. The threshold is a DAC set for each chip and depending on the chip configuration,

there are either one, two, four or eight thresholds per chip. Since pixels have different gains (slopes [keV/DAC]) and noise baselines (intercepts [keV]) as functions of threshold, they can also be calibrated.

For each individual pixel and each threshold a straight line calibration from threshold (DAC units) to energy (keV) can be determined. With that set of calibrations, one arrives at a relatively high number of straight line fits. The number depends on the chip configuration (colour mode or fine pitch mode) and the number of thresholds which are read. These combinations are shown in table 2.1.

After equalising, the basic threshold-energy calibration method is to only calibrate the chips, this tells at what mean energy the threshold(s) are set at. The pixel response distribution is therefore not known beyond the mean value.

Table 2.1: The number of straight line pixel calibrations given all possible combinations of the Medipix3 chip configuration

| Pixels | Thresholds | | | |
|-----------------------|------------|--------|-------|--------|
| | 1 | 2 | 4 | 8 |
| CM, 128×128 | 16384 | 32768 | 65536 | 131072 |
| FPM, 256×256 | 65536 | 131072 | N/A | N/A |

Since it is not possible to change the individual gains of each pixel with Medipix3 but just the noise baseline, it is not possible to improve the detector response uniformity at this point. The measurements to get to this point are relatively trivial and this level of knowledge is often sufficient, therefore this is often the end of the threshold-energy calibration story.

If one wishes to further improve the detector energy response uniformity, a per-pixel energy calibration is necessary. The requirements for such per-pixel calibration measurements are higher; the number of points to be calibrated on is about three orders of magnitude higher than for a chip. Instead of averaging the response of $128 \times 128 = 16384$ or $256 \times 256 = 65536$ pixels, each pixel is treated as an independent counting detector. This per pixel energy calibration is a post-acquisition step and is discussed in detail later in chapter 5.5. Such calibration measurements allow for significantly more information to be extracted per pixel, such as threshold dispersion, energy resolution, energy resolution variation and so on.

In either case, the calibration source requirements are the same. The ideal source can inject single energies into each pixel with zero bandwidth (the Dirac delta distribution), be easily varied over energy, have very high count rates, can provide uniform input and be stable over time. Yet again synchrotrons are close to this ideal source, except that they are not easily accessible. Other options are x-ray fluorescence (XRF), gamma emitting radioactive sources and test pulses from the pixel front-end. XRF measurements are used for calibration in this work exclusively because they are the best source overall given equipment availability during this work. This is summarised in table 2.2. A

method using the tube peak voltage (38) is not included in this overview due to the lack of accuracy and information acquired from this method and so it is not comparable to the other methods in Table 2.2.

Table 2.2: Qualitative comparison of sources for energy calibrations. A ✓ indicates that a source is sufficient in the metric displayed, × indicates insufficiency and ~ indicates a state somewhere between insufficient and sufficient.

| | Sources | | | |
|-----------------------------|---------|------------------|-------------|-------------|
| | XRF | γ sources | Test pulses | Synchrotron |
| Energy bandwidth | ✓ | ✓ | ~ | ✓ |
| Energy distribution quality | ~ | ✓ | ✓ | ✓ |
| High count rates | ✓ | ~ | ✓ | ✓ |
| Stability | ✓ | ✓ | ✓ | ✓ |
| Variable energy | ✓ | ~ | ✓ | ✓ |
| Uniform irradiation | ✓ | ~ | ✓ | ✓ |
| Sensor effects | ✓ | ✓ | × | ✓ |
| Ease of access | ✓ | ✓ | ✓ | × |

This is further studied with experiments in section 5.3.

2.6.4 SYSTEM STABILITY

Any device measuring a constant input and does not output a constant response means that the measurement device is not absolutely stable. All systems exhibit this to some degree, the relevant points are how unstable the device is and can it be compensated for. Compensating for such effects could be attempted with various levels of complexity; several approaches used for hybrid pixel detectors follow. Basic approaches include averaging multiple measurements and using separate devices to independently cross-check the primary device (39). Finally, more advanced options include chip and pixel characterisation and calibration, using prior knowledge to model the system response and deconvolute signals (40) and various pixel and chip level corrections (39).

There is a compromise to be made between the duration of a measurement and photon statistics. The longer the measurement is, the more photons can be measured which means that the relative error from photon counting statistics decreases over time, with the relation $\frac{\sqrt{N}}{N}$, where N is the number of photons. This relation ($\frac{\sqrt{N}}{N}$) comes from how photon counting is well described by Poissonian statistics at large numbers of counts (N). The Poisson function has a variance equal to the mean and for a large mean, the Poisson distribution approximates a Gaussian distribution. Both

the mean and variance of this Gaussian distribution are equal to N (and thus with a sigma of \sqrt{N}). Therefore the fractional error is simply $\frac{\sqrt{N}}{N}$.

Imagine a drift in count rate over time and that it is linear with some white (flat frequency distribution) noise added. If this effect was detector-wide then it is likely that it can be corrected for as well. If the individual pixels are changing response over time, that poses a ‘moving target’ problem. It is not feasible to correct for something that is constantly changing without additional information and/or an additional control loop. This means that pixel response stability relative to measurement times is critical for the maximum achievable precision with a given detector.




Figure 2.9: The relative error in percent as a function of time with various realistic values for the count rates per pixel and various example drift rates. The dark grey shaded area is excluded due to the maximum count rate per second per pixel for any Medipix3 chip. The x axis (time) is from 0.1 to 24 hours and the y axis (relative error) is from 0.0001 to 1%. 3.6×10^8 counts per hour per pixel (100 kHz per pixel) is the approximate upper limit for the linear counting regime for single Medipix3 pixels. No drift was assumed for the various count rates plotted for the solid lines.

The relative error as a function of time is illustrated with figure 2.9 with realistic values of the count rates per hour for a single pixel and various drift rates. The dark grey shaded region shows the lower limit of precision that one could attain based on the approximate maximum count rate that a Medipix3 pixel can handle while still having good count rate linearity. In other words, count rate linearity at design specification in SPM and CSM (36).

Count rate linearity is discussed in section 2.6.2 and the upper limit of approximately 10^5 counts per second per pixel (100 kHz per pixel) has been validated experimentally. The points at which the dashed lines (drift rates per hour) meet the solid lines ($\frac{\sqrt{N}}{N}$, relative error) indicate how long it takes to achieve the ultimate precision for a specific count rate and drift rate combination. After this time, the drift exceeds the relative error and the measurement will only degrade beyond that time.

The drift model presented of a linear increase with white noise overlaid is correct to the first order for the Medipix₃. There are second order effects, such as an additional low frequency component and how the direction of the drift oscillates over a long timescale. Modelling beyond first order effects is not within the scope of this analysis.

METHOD

For this analysis a Medipix₃ in CM (colour mode) is used; the pixels are $110 \times 110 \mu\text{m}$ squares. There are eight thresholds available for operating in Single Pixel Mode (SPM). The lowest threshold was set just above the pixel base-line noise (approximately minimum 4 keV in x-ray energy). Four to five DAC units correspond to approximately 1 keV in energy. For this measurement, no energy calibration is used and energy levels are expressed in units of DAC units.

After a perfect equalisation procedure, all pixels would have the same response to the same energy input. For this analysis we are imaging an open beam so all pixels receive the same energy spectrum at the same intensity. As a result, ideally, the only differences in counts between pixels would be due to photon count statistics. Photon counting statistics follow a Poisson distribution. For large numbers ($N > 100$) the Poisson distribution approximates to a Gaussian distribution. With an average of N photons the variance also is N , resulting in a sigma (σ) = \sqrt{N} . This is presented as noise in an image and sets a limit to the best contrast to noise ratio that can be achieved.

As discussed, inter-pixel count variation should ideally be photon counting fluctuations, \sqrt{N} only. However, due to the limited number of adjustment bits, the precision of the result is limited, so there will be a dispersion between pixels. Ideally, the solution would be to simply add more bits. This is not feasible because the adjustment bits are implemented per physical pixel. In the $55 \times 55 \mu\text{m}^2$ area available, everything that cannot be in the chip periphery must be in this small $3025 \mu\text{m}^2$ area. The number of adjustment bits per pixel is therefore a direct compromise between equalisation precision and space for all other digital, analogue and routing components required per pixel.

The limited number of adjustment bits presents as differences in inter-pixel sensitivities and therefore in systematic differences in the counted number of photons. Since there is a variation in the values of N , the distribution width will be larger than \sqrt{N} , which adds in quadrature to the photon counting statistics. If the differences in sensitivities between pixels are constant, then the corresponding differences in number of counts could be compensated for with a combination of addition and/or multiplication. In that case, the statistical fluctuations in photon counts will scale with this multiplication factor, causing it to deviate from $\frac{\sqrt{N}}{N}$. For deviations of N with less than

\sqrt{N} this effect is considered to be negligible, this is because the contribution will be less than the intrinsic uncertainty of \sqrt{N} .

Photon count rates are expected to vary over time due to drift in various subsystem components. There can be fluctuations in the pixel electronics that propagate as fluctuations in threshold voltages, resulting in the average count rate changing over time. Temperature changes may cause baseline levels in analogue circuitry to drift, radiation induced effects can influence the noise baselines and random telegraphic noise (RTN) may cause sudden jumps in signal levels. X-ray tube output can also change over time due to drift in tube voltage and/or tube current. Such voltage and current drift results in count rate changes common to all pixels. To first order, this will only affect the average number of counts for each pixel by the same amount. This only affects the signal to noise ratio through the $\frac{\sqrt{N}}{N}$. In this case of common drift across the detector, the signal to noise ratio would improve by increasing x-ray intensity and/or acquisition time.

If there are fluctuations in the count rate that are not shared by all pixels, there is yet another source of uncertainty which adds to the statistical lower limit. This will be visible in differences between images as an excess to the expected statistical contribution. This depends on the drift rate relative to the frame rate. For example, if the drift rate is consistently much slower than the frame rate then the drift would become negligible between consecutive frames. This therefore sets a lower detectable limit for drift rate.

MEASUREMENTS

Several measurements were run at low, medium and high frame rates and acquisition lengths to characterise system stability as described in section 2.6.4. The various frame rates were chosen to sample the frequency space at several time scales and to keep the overall measurement size manageable at less than 100 Gb per measurement. Measurement details are summarised in table 2.3.

Table 2.3: An overview of the system stability measurements of all the key parameters.

| Run | Duration | Number of images | Shutter open time (ms) | Shutter closed time (ms) | Tube peak voltage (keV) | Tube current (μA) |
|-----|----------|------------------|------------------------|--------------------------|-------------------------|--------------------------------|
| 1 | 82 s | 4635 | 16 | 1.6 | 25 | 200 |
| 2 | 2 hr | 7000 | 16 | 1000 | 25 | 200 |
| 3 | 2.7 days | 23693 | 500 | 9500 | 40 | 10 |
| 4 | 2 hr | 28076 | 0.01 | 200 | 40 | 1 |

ANALYSIS 1/2 - MEAN DRIFT RATES AND RANDOM TELEGRAPHIC NOISE (RTN)

System stability can be evaluated at multiple levels. This analysis has two angles of attack. The first is split based on logical units of the detector system - the whole detector (multiple chips), single chips and finally pixels. This breakdown was chosen to tell a more complete story about the detector-wide drift rate, focusing on how random telegraphic noise (RTN) impacts each logical unit. The second is a statistics driven approach. It gives insight into the optimum acquisition time for photon counting detectors in terms of count rate and drift rate.

An ideal measurement consists of a constant signal input with identical pixels, a perfectly equalised and stable detector. If one plots the mean counts of a such a measurement, this would result in a normal distribution with a sigma equal to the statistical uncertainty (\sqrt{N} , where N is the number of counts).

In reality, detectors will not be perfectly equalised, pixels have different gains and noise baselines and the response of pixels will not be absolutely stable over time. All of these effects result in deviations from the expected normal distribution, both in terms of the distribution mean and standard deviation. These deviations give us information about how the detector actually responds.

For example, the average counts could change over time, indicating a net drift effect over all pixels. It is possible to attempt to indirectly infer the cause of the variation from its properties. An example of this for run three (see table 2.3) is displayed in figure 2.10.

Figure 2.10 contains a 2D histogram without error bars due to the relatively high density of data points in the x axis (time). For many measurements, there will be fluctuations in the number of counts between measurements corresponding to photon counting statistics. In this case, the number of counts are the total number of counts of all pixels, divided by the number of pixels, with sigma corresponding to the chip wide total number of counts. The width of the distribution matches this for short time scales. In this case N is the total number of counts of all pixels in one frame divided by the number of pixels. As a result, $\frac{\sqrt{N}}{N}$ is the ratio you would expect for the sum of all counts of all pixels. As shown in the green pop-out rectangle, there are 359 points/hour.

The drift values for all runs are shown in table 2.4. Figure 2.10 shows an average (mean) drift of 0.01% per hour over the entire measurement time of 2.7 days (64.8 hours) for run three. Clearly, this metric does not tell the full story. Local variations can be very different in magnitude and direction.

Table 2.4: The full detector drift rates for all runs over the entire measurement time. Note that run one only lasted 82 seconds. The values are displayed to two decimal places.

| Run | Drift rate (%/hr) |
|-----|-------------------|
| 1 | 0.00 |
| 2 | 0.15 |
| 3 | 0.01 |
| 4 | -0.08 |




Figure 2.10: The upper plot shows how the average counts per frame for all pixels for run three data is relatively stable over time in absolute terms. The run numbers can be found in table 2.3. The lower plot is a 2D histogram of the same data as the upper plot, the y axis is zoomed into the data range. There is a straight line fit in blue showing the long term drift of 0.317 counts/hour and 0.0123 %/hour. Each bin contains a number of points in both axes, this number is shown using colour in the lower plot and the colour bar in the lower right. Zooming into a one hour time window, we see the underlying data in the highlighted green rectangle between 30 and 40 hours. No RTN (Random Telegraphic Noise) is visible here.

We now move onto the the next logical unit of this system, the Medipix₃ chips. Using the same analysis as for the whole detector, it can be observed in table 2.5 that the chips always drift in the same direction and there is no correlation between runs for the same chip. Since the signs of the chip and detector-wide values are always equal, it is expected that the underlying cause of the drift is some combination of detector electronics variation and the x-ray tube current variation.

Table 2.5: The chip-wide drift rates for all runs over the entire measurement time. Note that run one only lasted 82 seconds. The values are displayed to two decimal places.

| Run | Drift rates (%/hr) | | | |
|-----|--------------------|--------|--------|--------|
| | Chip 0 | Chip 1 | Chip 2 | Chip 3 |
| 1 | 0.00 | 0.00 | 0.00 | 0.00 |
| 2 | 0.12 | 0.17 | 0.20 | 0.99 |
| 3 | 0.01 | 0.01 | 0.01 | 0.01 |
| 4 | -0.05 | -0.13 | -0.03 | -0.09 |

All of the underlying plots for table 2.5 look very similar to 2.10 and so are not included, except the following in figure 2.11.

Figure 2.11 shows the outlier, which is run two, chip zero. In this figure, it is observed that there is a square wave with a seemingly random period and similar amplitude overlaying the signal. This is referred to as Random Telegraphic Noise (RTN), Random Telegraphic Signal (RTS) or burst noise'. RTN is a noise source which randomly jumps between two or more discrete levels and is found increasingly more in Si transistors as they shrink in size. It is most commonly caused by charge trapping and liberation events at the Si/gate insulator interface; this is a quantum process, hence the discrete nature of the signal jumping.

An investigation by the chip designers at CERN determined that 'the origin of all of these problems was found to be related to degradation of the integrity of the thin gate oxide during fabrication resulting from an absence of tie-down diodes at Metal-1 on the analog matched pairs. In fact, the layout was fully consistent with the antenna design rules provided by the foundry but did not follow a recommendation related to tie downs for matched pairs. This effect became visible due to the downscaling of the gate oxide thickness from ~ 5 nm (for Medipix₂ and Timepix) to ~ 2 nm and also the large number of pixels in the chip.' (41).




Figure 2.11: The upper plot shows how the average counts per frame for chip zero for run two data is relatively stable over time in absolute terms. The run numbers can be found in table 2.3. The lower plot is a 2D histogram of the same data as the upper plot, the y axis is zoomed into the data range. There is a straight line fit in blue showing the long term drift of 2.66 counts/hour and 0.117 %/hour. Each bin contains a number of points in both axes, this number is shown using colour in the lower plot and the colour bar in the lower right. Zooming into a one minute time window, we see the underlying data in the highlighted green rectangle between 1 and 1.5 hours. RTN (Random Telegraphic Noise) is visible here. It manifests itself as the sudden increases and decreases in the y axis.

Finally, we discuss the behaviour of individual pixels.

Pixels are observed to be behaving differently to the chips or the detector as a whole. As seen in table 2.6, four arbitrarily selected neighbouring pixels in chip zero have both positive and negative drift rates. The magnitude of the variation also varies significantly more than at chip or detector level.

Table 2.6: The drift rates for all pixels for all runs. σ is the standard deviation of the drift distribution.

| Run | Pixel drift rates (%/hr) | |
|-----|--------------------------|--------------------|
| | Mean (%/hour) | 3σ (%/hour) |
| 1 | -2.79 | 219.5 |
| 2 | 0.14 | 2.3 |
| 3 | 0.01 | 0.1 |
| 4 | -0.09 | 35.3 |


For example in run two, the absolute ratio of the least to most stable pixel drift rates is $|\frac{-2.32}{0.0626}| = 37.1$. Evidently, some pixels are significantly more stable than others when subject to the same irradiation.

The stability variation over time can also be seen between runs in table 2.6. The variation between runs for the same pixels is mostly within the same order of magnitude, but not always. Run 1 had 90 pixels beyond the mean $\pm 3\sigma$ range, these significantly increase the standard deviation of the distribution to 219.5 %/hour. The cross pixels have slightly lower drift rates for run 2. Run 3 had high drift rates for central pixels, negative drift for the other cross pixels and elevated drift of neighbouring to cross pixels. Finally, run 4 had very low drift rates for some central pixels.


For example, the drift rates per pixel and the accompanying histogram are shown for run two in figure 2.12. The plots for the other runs are included in appendix B.

No obvious correlation between pixels was observed and the drift rates of individual pixels were approximately normally distributed. These data-sets show that the mean drift rate decreases as the measurement duration increases. This is consistent with the random walk hypothesis. Further detailed studies would be required to make an attempt to isolate the exact causes of the observed variations and make further conclusions.

In this analysis, the determined drift rate was effectively the mean drift rate. This is due to the simple linear regression used which uses all data points, clearly it is not a good fit to the data. RTN introduces significant short-term variation in the count rate for what is assumed to be a constant input. The more relevant metric may be the drift rate over a typical single measurement time or collection of measurements. For microCT (micro Computed Tomography) scans, a single projection is typically between one second and one minute with approximately as many projections as pixels (> 256 projections); resulting in a total scan time of between 10 minutes and 10 hours.



(a) A 2D visualisation of the mean pixel drift rates in %/hour. Note the slightly lower drift rates in the ‘cross’ pixels. Green indicates dead pixels.



(b) The distribution of mean pixel drifts as a histogram with the mean plus/minus three standard deviations marked with the red vertical lines. Note the log scale on the y axis.

Figure 2.12: The mean drift rates of all pixels for run two in %/hour.

Qualitatively, the conclusion is that the optimum measurement time & the drift rates limit the maximum accuracy. It is possible to determine the optimum measurement time using a very simple protocol for each individual experimental setup. This protocol simply involves measuring many frames in the same conditions and similar measurement times as the specific experimental setup demands. If a shorter measurement time is required, the maximum achievable precision decreases for a fixed count rate. The maximum achievable precision for a given measurement time can be quantified with this approach. This directly translates into what the system can resolve and thus the system can be optimised given the trade-off between time, particle flux and required precision.

ANALYSIS 2/2 - STATISTICS AND MEASUREMENTS

The second part of the analysis approaches system stability from a different angle where the stability per pixel is characterised.

This is a statistical analysis supported by three data-sets. The three measurements were taken on different days using typical microCT conditions in the Nikhef cabinet. Unless otherwise specified, the data-set used is from 4 November 2021. The same equalisation and configuration was used as for the rest of this thesis.

The relevant parameters were:

- Number of frames = 40,000;
- Number of thresholds recorded = 8;
- Number of images = $40,000 \times 8 = 320,000$;
- Image type = 16 bit unsigned integer, TIFF (Tagged Image File Format) file;
- Data-set size = 42.0 Gb;
- Threshold mean energies, from threshold zero to seven = 5.0, 9.0, 10.5, 13.2, 16.1, 19.0, 22.0, 24.0 keV;
- Shutter open time = 88.4 ms;
- Shutter closed time (readout time) = 1.6 ms;
- Total measurement time = 1 hour;
- Tube peak voltage = 30 kV;
- Tube current = 200 μ A;
- Source to detector distance = 385 ± 5 mm;
- Detector = Medipix₃ 110 μ m pixel pitch, 500 μ m silicon sensor biased to +150 V, leakage current $\simeq 1-10$ μ A.

The following analysis uses a one hour measurement - typical microCT measurement time.

The difference between two consecutive frames of a set of pixels with approximately the same count rate, should be zero on average. The difference between two consecutive frames is now referred to as a 'delta frame'. The standard deviation of a delta frame, $\sigma_{f_n - f_{n-1}}$, should be the quadratic sum of the standard deviation of the two frames, equation 2.6, where N is the number of counts.

$$\sigma_{f_n - f_{n-1}} = \sqrt{\sigma_{frame_n}^2 + \sigma_{frame_n-1}^2} = \sqrt{2N} \quad (2.6)$$

This is illustrated in the lower half of figure 2.13. The mean is very close to zero at -0.58 counts. The expected value for the standard deviation in this case is: $\sqrt{2183.894 + 2183.313} = 66.08$ counts.

The measured value is slightly lower at 65.1 counts as also seen in figure 2.14. This indicates that there is a correlation between counts which means that the counting rate is slightly too high. This is very close to the expected value.




Figure 2.13: Two histograms, the upper subplot is the histogram of the count values of the first frame. The lower subplot is the standard deviation of the first delta frame, the first frame (f_0) minus the second frame (f_1). The mean and standard deviation values are shown in the legend for both subplots. Due to variations in pixel gains, pixel offsets and systematic offsets between the four chips, the upper subplot is not expected to be normally distributed. Therefore the standard deviation of the upper plot should not be equal to the square root of the mean.

A distribution width more than $\sqrt{2N}$ indicates an uncertainty in addition to photon counting fluctuations; this is effectively the sum of all non-photon counting statistics. A width of less than $\sqrt{2N}$ indicates a correlation between photon counts which is likely due to using a count rate that is too high (pile-up effects), this is observed in figure 2.14.




Figure 2.14: The standard deviation of delta frames over time. The mean value of the data is slightly less than the expected value ($\sqrt{2N}$ of f_0 , where $f_0 = 2183.89$ counts) indicated by the horizontal solid blue line at 66.08 counts.

For the sake of reasoning we assume all pixels in consideration to have the same mean number of counts N , unless otherwise noted. The standard deviation of a delta frame is:

$$\sigma_{f_n - f_{n-1}} = \sqrt{2N + \sigma_{others}^2} \quad (2.7)$$

σ_{others} is the extra uncertainty from all non-photon counting related sources. Of course this is ideally zero, however in reality this is not the case as seen in figure 2.15. The relative uncertainty including both photon counting and all other sources uncertainties will be characterised:

$$\frac{\sigma}{N} = \frac{\sqrt{2N + \sigma_{others}^2}}{N} \quad (2.8)$$




Figure 2.15: The standard deviation of the difference to the first frame, $\sigma(f_n - f_0)$, is shown on the top. This is the photon counting statistics combined with all other contributions. The standard deviation of the difference with the first frame, $\sigma_{others} = \sqrt{\sigma(f_n - f_0)^2 - \sqrt{2N}}$, is shown on the bottom. This is just non-photon counting contributions.

Figure 2.16 shows the fit parameters in the upper half and the fit residuals in the lower half. The fit residuals indicate how the model appears to fit the data well because the slope (a) and the intercept (b) parameters are very close to zero. The first bin in time is consistently lower than the fit would suggest, given that this feature is observed over multiple data-sets, this likely indicates a

deficiency in the fit model and not RTN.

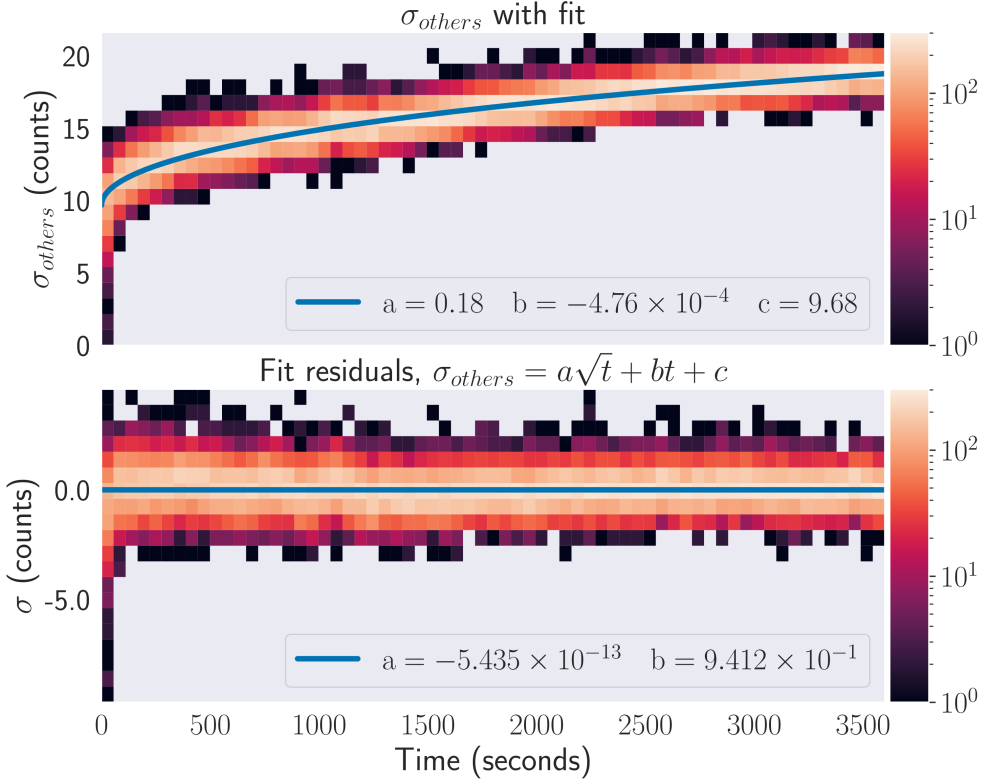


Figure 2.16: The standard deviation of the differences to the first frame minus photon counting statistics contribution with a fitted function with the form of $a\sqrt{t} + bt + c$ on top. The residuals are plotted below.

As stated in equation 2.7, the standard deviation of a delta frame is:

$$\sigma_{f_n - f_{n-1}} = \sqrt{2N + \sigma_{others}^2}$$

There appears to be no knowledge about the underlying mechanisms of pixel drifting behaviour. It happens that for different data-sets with significantly varying timescales, frame rates and count rates, this equation was found to fit relatively well for all cases. Further studies are required to understand the mechanisms and optimise the model. This is a first approach to quantify the additional uncertainty on top of photon counting statistics.

It is therefore hypothesised that based on the behaviour of σ_{others} in figures 2.15 and 2.16, the model in equation 2.9 is a reasonable approximation; where a , b and c are fit parameters and t is

time.

$$\sigma_{others} = \sqrt{\sigma_{total}^2 - 2N} \simeq a\sqrt{t} + bt + c \quad (2.9)$$

At this point, one may wonder about the physical meaning of these fit parameters (a , b & c). It was expected that σ_{others} would vary due to a random walk type effect. Therefore it should scale with the square root of time and indeed we see that there is a $a\sqrt{t}$ term. It was also expected that there might be a linear drift component which is encapsulated by the second parameter, b . This covers the bt term. The final parameter, c , would ideally pass through the coordinates (0, 0) because at $t = 0$, the drift should also be precisely 0. All data-sets indicate that a small offset is actually present - there is a relatively sharper rise in σ_{others} in the first bin consistently. Without this offset, the predicted value of σ_{others} at $t = 0$ would be too high consistently.

We do not have detailed knowledge of the implementation and exact manufacturing process used to fabricate these Medipix3 chips. Therefore, we also lack the knowledge regarding the underlying cause of the drifting observed. The exclusive aims of this analysis were to reliably quantify these effects and to optimise the acquisition protocol. Fixing the drift problem would require significant attention from the chip designers and semiconductor fabrication plant staff; it is therefore out of scope of this work. This drift behaviour is not expected to be the same with Medipix4.

As a side note, the SNR and change in SNR over time for all pixels as plotted in figure 2.17. Firstly, it shows yet another example of RTN. More importantly though, it can be observed that the change in SNR between adjacent frames over time is relatively stable and flat over one hour. The horizontal lines plotted are the mean $= -2.8 \times 10^{-5} \pm 3\sigma$ where $3\sigma = 0.614$. This is consistent with figure 2.15.




Figure 2.17: The SNR (signal to noise ratio) and the change in SNR for all pixels over a one hour measurement. The blue horizontal line is the straight line fit with the two parameters (a and b) listed in the legend. The two red horizontal lines are the mean value $\pm 3\sigma$ to indicate the small number of points outside that range. Note the logarithmic colour scale.

2.6.5 SUMMARY

In summary, this is not a perfectly stable system and is therefore not only limited by photon counting statistics. This system instability is quantified at system, chip and pixel levels. Therefore, there are optimum measurement times which will depend on count rate and drift rate. The count rate can be varied with tube current settings, the distance from source to detector and filters.

Part one discusses the breakdown of the drift rate analysis to the three logical components (detector, chips and pixels), resulting in a more complete picture. While the detector as a whole may drift, the individual chips also drift in the same direction by similar amounts. The chips do sometimes show RTN which affects the short term variations greatly but not the longer term drift rates so much. The pixels have significant variations in drift rates and are also affected by RTN. Since both single pixels and chips are affected by RTN, there is evidence that RTN likely has two causes in Medipix₃, this is consistent with Ballabriga et al. (41).

Moving onto part two, the photon counting drift contributions are shown to be dominant for the count rates we work with. The non photon counting drift contributions in figure 2.16 only contribute to less than three counts over one hour. Given that we typically have a flux between 1—100 kHz per pixel, the non photon counting drift is six to eight orders of magnitude lower than our typical flux. For a low flux, precision measurement, this presents a lower achievable relative error. MicroCT measurements are relatively very high flux and are limited by photon counting statistics, not pixel drift.

The minimum relative error is limited by the drift rates and their associated uncertainties. This property is shared between photon counting detectors. In astronomy for example, where the goal is often to measure the faintest objects, the measurement time can be extended to match the required sensitivity level partly because they are able to achieve very high system stability.

3

Protons and heavy ions

This chapter explores the possibility of using Medipix3 within hadron therapy and involves significant collaboration with colleagues.

A related work not included in this thesis was submitted for the proceedings of the International Particle Accelerator Conference (IPAC) 2021. The title was ‘Tracking and LET Measurements with the MiniPIX-TimePIX Detector for 60 MeV Clinical Protons’. It was included in Jacinta Yap’s PhD thesis (42). We measured the LET (linear energy transfer) of 60 MeV protons using a Timepix based detector at several angles. LET distributions have been challenging to measure until now due to detector limitations. These LET distributions are useful for hadron therapy treatment planning and for modelling radiobiological effects (43). Ideally these distributions would have been measured in water or water equivalent material instead of silicon. These distributions could be integrated into Monte Carlo frameworks used for treatment planning such as FLUKA (44; 45). My contributions included data analysis, experimental planning and execution, interpretation of data, data acquisition and detector operation.

3.1 ION IDENTIFICATION AND THE VOLCANO EFFECT AT HIMAC, JAPAN

This section determines the upper limit in terms of energy per particle that could be measured, John Idarraga did the measurements and I produced the text and figures.

The Medipix3 chip has been tested in a highly ionising environment with protons, 4-He, 40-Ar and 58-Fe at HIMAC (Heavy Ion Medical Accelerator in Chiba, Japan) (46) in collaboration with Dr John Idarrága-Muñoz. A silicon sensor was used with a thickness of 500 μm , bump bonded to four Medipix3 chips. An overview of the particle species and energies and other relevant information can be found in table 3.1.

Table 3.1: Quantities involving the particle species and their respective energies and LET (linear energy transfer)

| Particle | Proton | Helium | Argon | Iron |
|---|--------|--------|--------|-------|
| Nucleon number, A (nucleons) | 1 | 4 | 40 | 56 |
| Energy per nucleon, E (MeV/A) | 100 | 100 | 650 | 250 |
| Linear Energy Transfer, LET (MeV/ μm) | 0.0012 | 0.0053 | 0.1466 | 0.492 |
| LET in 500 μm Si (MeV) | 0.6 | 2.65 | 73.3 | 246 |

In these measurements, the amount of energy deposited in a single $55 \times 55 \mu\text{m}$ pixel can go up to 10 MeV (corresponding to 2.76×10^6 electrons in silicon), this is more than four orders of magnitude above the minimum detection energy in Si of 4 keV (approximately 1100 electrons). The so-called ‘volcano effect’ (47; 48) is observed in the Medipix3 chip as in other Medipix detectors at approximately 7 MeV. This finding informs chip designers and users of the upper limit of the amount of energy which can be deposited in the Medipix3 front-end before encountering the ‘volcano effect’.

Single particle counting detectors are intended to count once per particle. For the radiation species and energies that this is designed for, that is mostly accurate. However, since this is using relatively high energy protons and heavy ions, this assumption breaks down. This is because the energy deposition of 100 MeV scale protons is three orders of magnitude higher than keV x-ray photons. With such a high input charge, the pixel front-end electronics are far out of design specification and do not produce an accurate output. It is hypothesised that this effect is due to under-damped oscillations in the front-end electronics which produce multiple counts where only one should have been registered*.

The ‘volcano effect’ refers to the phenomenon that occurs in an image where the input charge is high enough to cause an artefact. This can be observed when looking at a cross section of a group of pixels that detected the charge cloud. The number of counts steadily rises until they peak, then subsequently rapidly falls to zero (or near zero) and then rises up to the previous peak and steadily decreases to zero; see figure 3.1.

*This was discussed with chip designers, mainly Rafael Ballabriga-Suñé




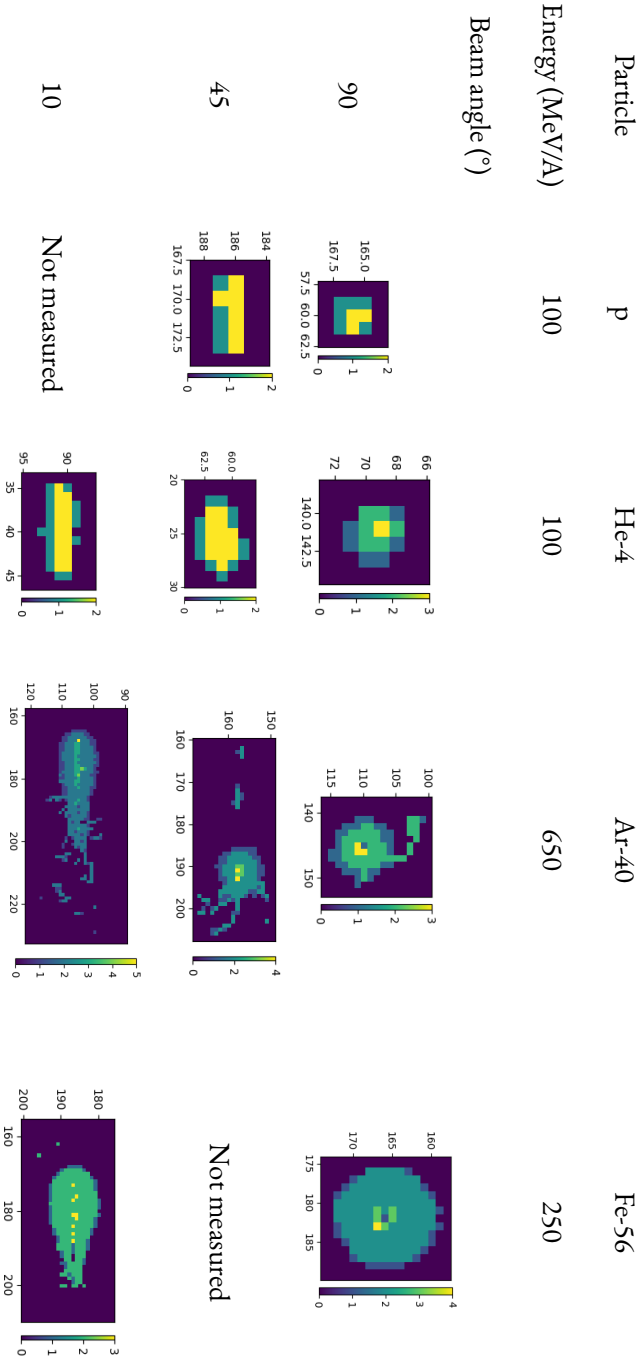
Figure 3.1: Example profile of the volcano effect.

This effect places the upper practical limit on the maximum LET (linear energy transfer) particles that can be used with the Medipix3 chip in the current configuration and without any compensation for this effect, which constrains the possible applications. Attempts to compensate for this effect have been made (47). Within design parameters, the charge collected in a given pixel is linearly proportional to the output of the front-end of that pixel within approximately five percent. There is an upper limit to this linearity which is far beyond the design specification. Above this limit, the front-end response of the chip is very non-linear, it suddenly drops to zero or near zero output.

This does not mean quantitative measurements are not possible, one could use Medipix3 for heavy ion measurements even though the chip is operating far beyond nominal design conditions. More measurements would be necessary to investigate the stability of these responses and how counts are related to the amount of deposited charge or energy. The focus here is the usability for ion identification measurements.

Table 3.2 shows single particle hits with 100 MeV protons, 100 MeV/A He-4 ions, 650 MeV/A Ar-40 ions and 250 MeV/A Fe-56 ions at 90, 45 and 10 degrees between the beam and the detector plane. Note the non-constant scaling of the images; each pixel is a square. The scale of each subplot varies significantly between the 100 MeV proton at zero degrees (top left) and 250 MeV/A Fe-56 ions (bottom right); the top left is only showing a cluster of nine pixels whereas the bottom right is showing more than 200 pixels.

Table 3.2: Each image is a single particle event from a 1 ms exposure time. Different angles are chosen to change the energy deposition in the sensor layer; higher angles mean higher energy deposition in the sensor. The Bragg peak is within the sensor for the 45 and 80 degree events. The volcano effect is visible from argon and iron as the darkening in the centre of the clusters.



The volcano effect is observed starting at Ar-40 ions. The related Timepix (49) family of chips observe the same effect at a similar LET of 2.2 MeV (50). This is unexpected because Timepix chips are very different to Medipix, even though they have been designed by the same team using the same techniques and processes. The volcano effect is not observed for light ions and protons. Charge sharing proportional to the deposited charge between pixels is observed as expected.

These data were measured during the same beam-time as a complementary study into heavy ion spectroscopy using Timepix3 which found that the volcano effect was observed at around 500 keV for Timepix3 (51).

These measurements show the upper limit of LET that Medipix3 can handle without post-measurement compensation. Additionally, it shows that one could use Medipix3 for ion identification through analysing the various cluster properties. This has been done for Timepix previously for NASA (National Aeronautics and Space Administration) on the ISS (International Space Station) (17; 19; 52).

3.2 BEAM PROFILE MEASUREMENTS AT CLATTERBRIDGE CANCER CENTRE - WIRRAL NHS (NATIONAL HEALTH SERVICE), UK

The beam profile measurements at Clatterbridge Cancer Centre demonstrate the first use of the Medipix3 detector using a clinical proton therapy cyclotron. Jacinta Yap and I had an approximately equal contribution towards this section. I contributed all Medipix3 content, Jacinta contributed all of the EBT3 film content and Andrzej Kacperek with his team at the CCC facilitated this work with the cyclotron operations and access. Jacinta and I contributed equally towards the remaining sections. Javier Resta López & Carsten Welsch were Jacinta's supervisors.

TECHNICAL REPORT

3

Medipix3 for dosimetry and real-time beam monitoring: first tests at a 60 MeV proton therapy facility

J.S.L. Yap,^{a,b,c,*} N.J.S. Bal,^d A. Kacperek,^e J. Resta López^{b,c,f} and C.P. Welsch^{b,c}

^a*School of Physics, University of Melbourne,
Swanston Street & Tin Alley, Melbourne 3010, Australia*

^b*Department of Physics, University of Liverpool,
Oxford Street, Liverpool L69 7ZE, United Kingdom*

^c*Cockcroft Institute,
Keckwick Lane, Warrington WA4 4AD, United Kingdom*

^d*Nikhef,
Science Park 105, 1098 XG Amsterdam, Netherlands*

^e*Clatterbridge Cancer Centre, NHS Foundation Trust,
Clatterbridge Road, Wirral CH63 4JY, United Kingdom*

^f*Institute of Materials Science (ICMUV), University of Valencia,
Carrer del Catedr tic Jos  Beltr n Mart nez, 2, Val ncia 46980, Spain*

E-mail: jacinta.yap@unimelb.edu.au

ABSTRACT: Charged particle therapy (CPT) is an advanced modality of radiation therapy which has grown rapidly worldwide, driven by recent developments in technology and methods of delivery. To ensure safe and high quality treatments, various instruments are used for a range of different measurements such as for quality assurance, monitoring and dosimetry purposes. With the emergence of new and enhanced delivery techniques, systems with improved capabilities are needed to exceed existing performance limitations of conventional tools. The Medipix3 is a hybrid pixel detector able to count individual protons with millisecond time resolution at clinical flux with near instant readout and count rate linearity. The system has previously demonstrated use in medical and other applications, showing wide versatility and potential for particle therapy. In this work we present measurements of the Medipix3 detector in the 60 MeV ocular proton therapy beamline at the Clatterbridge Cancer Centre, U.K. The beam current and lateral beam profiles were evaluated at multiple positions in the treatment line and compared with EBT3 Gafchromic film. The recorded count rate linearity and temporal analysis of the beam structure was measured with Medipix3 across the full range of available beam intensities, up to 3.12×10^{10} protons/s. We explore the capacity of Medipix3 to provide non-reference measurements and its applicability as a tool for dosimetry and

*Corresponding author.



beam monitoring for CPT. This is the first known time the performance of the Medipix3 detector technology has been tested within a clinical, high proton flux environment.

KEYWORDS: Hybrid detectors; Instrumentation for hadron therapy; Instrumentation for particle-beam therapy; Beam-line instrumentation (beam position and profile monitors, beam-intensity monitors, bunch length monitors)

ARXIV EPRINT: [2107.02369v3](https://arxiv.org/abs/2107.02369v3)

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | Materials and methods | 2 |
| 2.1 | Clatterbridge proton beam | 2 |
| 2.2 | Medipix3 detector | 3 |
| 2.3 | Experimental setup | 5 |
| 2.4 | EBT3 film | 7 |
| 2.5 | Film calibration and image analysis | 8 |
| 2.6 | Detector activation & temporal analysis | 9 |
| 3 | Measurements and results | 10 |
| 3.1 | Beam profiles | 10 |
| 3.2 | Detector response | 13 |
| 3.3 | Detector activation | 14 |
| 3.4 | Temporal analysis of the beam | 15 |
| 4 | Conclusions and outlook | 18 |

1 Introduction

The use of particle beams for radiotherapy is expanding worldwide, led by developments in related technologies, growing clinical practice and reported improvements in treatment outcomes. The physical advantages offered by ions as demonstrated by the characteristic ‘Bragg Peak’, translate to a higher possibility of delivering a more precise amount of radiation with greater radiobiological effectiveness. The exploitation of these benefits have progressed significantly in recent years: advanced delivery techniques and modalities such as proton beam therapy (PBT) are well established.

To ensure the safe and effective delivery of charged particle therapy, it is essential that the characteristics of the beam are accurately and reliably measured. Various systems are used for different measurements for beam monitoring, quality assurance as well as dosimetry under reference or non-reference conditions. Procedures may vary across facilities as they depend on accelerator type, delivery method, vendor, regulations and other parameters. Dosimetry practices are well defined as the determination of dose must be precise and reproducible. The recommended protocol for absolute, relative, reference and non-reference measurements adopted for CPT are detailed in reports [1, 2]; the general standards for existing detectors are discussed extensively in literature and summarised in [3–7]. Each detector has certain advantages and disadvantages given the measurement application and no single tool can characterise all the necessary beam parameters: multiple detectors for different measurements are required in the workflow. For dosimetry and beam monitoring, the fundamental technology of conventional tools have remained relatively unchanged for the past

10–20 years [6]. In contrast, limitations are becoming more pronounced as new techniques are emerging, particularly with the evolution of accelerators and the shift to higher ‘FLASH’ dose rates [8, 9]. These performance requirements surpass current instruments [10], driving research and development into detector systems with improved capabilities.

Silicon based detectors have been utilised mainly due to their high spatial resolution and sensitivity but the typically used diodes and commercially available systems exhibit issues with energy dependence and a poor response at higher dose rates, along with other disadvantages [11]. Nevertheless, there are many recognised beneficial properties and solid state devices have been widely investigated, resulting in developments into several novel instruments [12–16] for a range of applications in radiotherapy. Similarly, the use of hybrid pixel detectors has demonstrated great versatility and has been successfully applied for a broad range of areas such as astrophysics, dosimetry, electron microscopy, life sciences, etc. The Medipix detector technology was originally developed for particle tracking at the LHC and then used for radiation imaging and x-ray detection [17]. The newer generations have further progressed its specific use for medical applications and suitability for CPT [18, 19].

The Medipix3 is a hybrid pixel detector which comprises a single quantum counting chip and a SPIDR (Speedy Pixel Detector Readout) readout system, enabling direct measurements by individual particle counting with relatively high count rate and good radiation hardness [20]. The chip is capable of high flux operation, thus making it appropriate for the characterisation and measurement of ion beams. In this paper we present test results of the system within a clinical PBT environment for the first time, to examine the potential and applicability of the Medipix3 detector as a tool for dosimetry and monitoring for CPT beams.

2 Materials and methods

Measurements were performed with the 60 MeV clinical proton therapy beam at the Clatterbridge Cancer Centre (CCC), U.K. The Medipix3 system was placed at three different locations in the treatment beamline and irradiated simultaneously with GAFchromic™ EBT3 film [21] using beam currents ranging from 0.012–1.97 nA. The transverse beam profiles obtained with film were determined and compared with the hit distribution recorded by the detector. The average count rate for each run was also examined to determine the count rate linearity of the response of the detector.

2.1 Clatterbridge proton beam

The Clatterbridge facility is the world’s first hospital based particle therapy centre and has been successfully treating patients for ocular cancers since 1989 [22]. The accelerator and beamline was originally built and commissioned for fast neutron therapy trials but was converted shortly after into a proton therapy service, specifically for the treatment of eye tumours. Given its pioneering and historical operation, many of the treatment line components and quality assurance (QA) equipment were constructed in-house. Several modifications to the accelerator and transport line were necessary and performed over time, resulting in differences to characteristics and properties of the beam reported initially at commissioning. The current state of the facility has been studied in detail in [23], general characteristics and treatment beam parameters are mentioned; for this study, dose rates below and within the clinical range were used.

As the facility operates a busy patient schedule, the delivery system is optimised and designed in such a way to consistently deliver a beam constrained by clinical requirements. Basic beam measurements and dose monitoring are performed within the treatment beamline itself: the integrated charge is provided by the electrometer attached to the second scattering foil, a pair of ionisation chamber dose monitors (F) give the converted dose in monitor units and a set of tungsten cross-wires (G) are used for patient positioning (imaged with x-ray panels). An X-Y diode scanner (not shown), placed after the treatment nozzle, is used to check the dose uniformity in both directions in the transverse plane before each treatment fraction. As shown below in figure 1, a resulting conformal dose distribution is delivered by several beam shaping elements in the passive scattering system: (A) two tungsten foils and a brass stopper moderate the shape and fluence of the beam, (B & C) range shifters vary the depth, (E) multiple beamline anti-scatter collimators restrict the transverse spread and a patient specific collimator constructed to match the tumour shape in the axial direction, can be placed within the (H) brass treatment nozzle.




Figure 1. Schematic of the CCC treatment beamline with diagnostic components (electrometer, dose monitors and cross-wires) and delivery elements (A-H; scattering foils, modulators, nozzle etc.) labelled.

The CCC treatment line provides unique clinical conditions which are not common by modern standards; the passive delivery system as well as resulting uncertainties associated with the beam parameters and quality presented both a challenging and promising environment for testing with the Medipix3 system. In addition, these measurements also provide further information about the behaviour and parameters of the Clatterbridge beam. It is noted that there are ongoing experiments and simulation studies being performed on the beamline [24–26]. Information on the transverse beam profiles, beam divergence and lateral spread are also useful for model verification and validation.

2.2 Medipix3 detector

The Medipix3 is a versatile and fast hybrid pixel application specific integrated circuit (ASIC). Among its features are configurable pixel pitch between $55\text{ }\mu\text{m}$ and $110\text{ }\mu\text{m}$, relatively high count rates of up to $\sim 100\text{ kHz}$ per pixel ($\sim 6.5\text{ GHz}$ over one chip if evenly distributed) and it can use many different sensor types of various thicknesses. Common sensor options include high-resistivity silicon (Si), gallium arsenide (GaAs), cadmium telluride (CdTe), CVD (chemical vapor deposition) diamond and even certain gases.

As illustrated in figure 2, when a 60 MeV proton impinges on the sensor and travels through the sensor bulk, it produces many electron-hole pairs by electronic and nuclear interactions where the mean rate of energy loss is described by the Bethe-Bloch equation. The generated electron-hole pairs drift to the front and back-side of the sensor due to an applied electric field. The moving




Figure 2. An overview of a hybrid pixel detector when a high energy proton enters the sensor bulk and transfers energy mainly by ionising sensor atoms producing electron-hole pairs which drift to the cathode and anode depending on the polarity of the induced electric field.

charge induces a current (Shockley-Ramo theorem) through the pixel implants (collection electrodes) which are reverse biased, depleted p-n junctions. This starts the pulse processing chain within one pixel. The measured charge is proportional to the energy deposition of the particle. Since the mean range of a 60 MeV proton in Si is $\frac{\text{CSDA range}}{\rho_{\text{Si}}} = \frac{3.94 \text{ g/cm}^2}{2.33 \text{ g/cm}^3} = 16.93 \text{ mm}$ and the detector is 500 μm thick, the Bragg peak does not occur within the silicon bulk of the detector if the detector is perpendicular to the beam.¹ If the detector is positioned parallel to the beam, the Bragg peak for 60 MeV protons would then be within the width of the detector; the feasibility to perform depth-dose profile measurements also presents a further avenue to explore in future.

Nevertheless, at this proton energy, it is likely that the deposited charge will be collected by more than one pixel. This effect is called ‘charge sharing’ [28]. It results in the detector counting an average of more than once per 60 MeV proton. A detector mode exists to address charge sharing called the ‘charge summing mode’ which sums the collected charge over an arbitrary 2×2 pixel grid. The penalty is a reduction of approximately one order of magnitude in count rate along with double the electronic noise.

In addition to simple geometric effects, it is likely that the charge cloud generated by the protons will also not be centred on the pixels and therefore cause an increased count in comparison to the number of protons which actually pass through the sensor. Charge summing mode was not used in this work due to the significant count rate penalty.

¹CSDA proton range in Si is from the PSTAR database, NIST [27].

A detailed description of the Medipix3 chip can be found in [20]. For this study, the ASIC was bump-bonded to a 500 μm silicon sensor with a pixel pitch of 55 μm . This detector is made of four $55 \times 55 \mu\text{m}^2$ chips in a 2×2 arrangement, where each chip consists of $256 \times 256 = 65,536$ pixels. Therefore with 4 chips, the total number of square pixels is 262,144 and the active area covers $28 \times 28 \text{ mm}^2$.

These measurements were performed with a Medipix3 based detector using a SPIDR readout system [29, 30] from Nikhef, with the experimental layout as shown in figure 4. The detector was biased with +100 V and configured to collect holes. The sensitive area of the detector was held in an aluminium cooling block and connected to the readout system by fibre cables. A small assembly was designed to securely hold a piece of film 3.5 cm in front of and parallel to the sensor and was clamped onto the edges of the block (figure 3).




Figure 3. Medipix3 chipboard affixed to an aluminium cooling block with the red SPIDR readout system behind. The 10 GBit/s fibre cable takes the data from the detector to the readout computer and a power cable is also attached to the SPIDR (not in view).

2.3 Experimental setup

The Medipix3 detector was placed at three locations throughout the treatment beamline and irradiated under varying beam conditions (table 1). In the integration zone, downstream of the scattering foils and modulation box (aluminium drift pipe removed) and two positions downstream of the treatment nozzle. Sections of EBT3 Gafchromic film were also positioned in front of the detector (figure 4) and irradiated simultaneously for direct comparison.

Prior to the experiment, the estimated dose and fluence were calculated to determine the range of measurements and expected number of registered events. The Medipix3 chip can tolerate relatively high amounts of radiation (still functional after irradiation by 1 MeV neutron equivalent fluence of $5 \times 10^{14} \text{ cm}^{-2}$ [31]) and hence was also situated in the beginning of the integration zone. At this position, the beam FWHM was determined to be smaller than the sensor sensitive area however




Figure 4. Experimental setup with measurements performed at three different locations (1.8 cm in the integration zone, 6 cm and 26.5 cm after the treatment nozzle) with the Medipix3 detector and EBT3 film simultaneously. A clamp assembly holds the film 3.5 cm directly upstream of the detector face.

Table 1. Overview measurement parameters of detector position, acquisition time and beam currents.

| Run (#) | Beam current (nA) | Time (s) | Distance from nozzle (cm) |
|---------|-------------------|----------|---------------------------|
| 2 | 0.012 | 97.2 | 9.5 |
| 3 | 0.052 | 99.8 | 9.5 |
| 4 | 0.35 | 49.8 | 9.5 |
| 5 | 0.69 | 44.9 | 9.5 |
| 6 | 0.27 | 32.6 | 9.5 |
| 7 | 0.27 | 29.5 | 30.0 |
| 8 | 1.35 | 75.0 | 30.0 |
| 9 | 1.35 | 68.9 | 9.5 |
| 10 | 1.97 | 66.4 | 9.5 |
| 11 | 0 | 9.0 | 9.5 |
| 14 | 2.2 | 103.0 | Integration zone |
| 15 | 2.1 | 285.3 | Integration zone |

diverges significantly along the beamline. A 20 mm collimator was placed within the treatment nozzle for the downstream measurements, to ensure complete beam coverage across the sensor.

A minimum stable beam current of 0.012 nA was attained on the day and then ramped up to higher currents for subsequent measurements. This is measured by the electrometer connected to the second scattering foil and is listed in table 1 for each run for consistency: the dose monitors could not provide readings whilst the detector was placed within the integration zone. The foil currents have a linear relationship with the ion chamber monitor units.

2.4 EBT3 film

EBT3 GAFchromic film is a standard radiochromic film dosimeter commonly used for quality assurance in radiation therapy to analyse the geometrical beam characteristics (i.e. uniformity, shape and size) and measure the 2D or 3D dose distribution [32, 33]. For patient specific or machine verification, it is essential to be able to perform checks with high accuracy and reliability, prior to delivering a course of treatment. In proton beam therapy, the use of Gafchromic EBT3 for film dosimetry to determine beam performance and quality is well established [34–36]. EBT3 Gafchromic film is made of a 28 μm layer of Lucite, enclosed by 125 μm of polyester substrate on each side (figure 5).




Figure 5. Sketch of the composite layers in EBT3 film, a 28 μm active region is surrounded by two thicker substrate layers.

Exposure to ionising radiation results in polymerisation of free radicals within the active layer, inducing the film to darken [37]. EBT3 self develops and the dark colouring or optical density (OD) is proportional to the extent of irradiation, increasing with absorbed dose. The geometrical beam distribution is provided as a function of the dose, evaluated by converting the grey value of each pixel to an OD value. The high spatial resolution and low energy dependence enables measurements of the transverse dose profiles. However, the use of EBT3 film in CPT is limited due to significant quenching effects and saturation at points of high linear energy transfer, such as at the Bragg peak [32, 38]. These effects were not experienced as these were transmission measurements. It is also noted that there are multiple considerations and associated sources of uncertainties [39]. Calibration measurements must be performed under specific conditions: there are established methods and standard protocol for the OD to dose conversion process as widely reported in literature [32, 35, 40, 41]. These measurements were carried out in accordance with these methods; the OD value for each pixel results in a corresponding dose (Gy) and is determined by establishing a correlation with known quantities of radiation. This is done by exposing a calibration set of film (figure 6) to well defined quantities of radiation under standard conditions, with reference to the dose measured by the ion chambers.

The calibration set of film was placed downstream of the nozzle and irradiated individually with doses ranging from 4–25 Gy to generate a calibration curve. For the film attached to the detector system (figure 4), a single piece was cut into several equivalent segments and labelled to ensure that the direction and orientation remained consistent. Following complete development (> 24 hours), the irradiated film pieces were scanned using an EPSON 750 scanner and saved as 48-bit TIFF (Tag Image File Format) images with no colour corrections at 150 dots per inch (DPI). All film analysis



Figure 6. Calibration set of film irradiation with 4–25 Gy.

was done using the image processing software ImageJ [42]. A circular region of interest (ROI) was selected such that it was encased inside each beam spot and duplicated for each film. ImageJ was used to generate the ROI intensity metrics for each colour channel to obtain a calibration curve, providing a correlation between the OD and dose.

2.5 Film calibration and image analysis

The calibration curve is obtained by evaluating the net OD values [41] across the full dose range, where the film response to the measured dose is expressed as the difference between transmission intensities:

$$\text{net OD} = OD_{\text{exp}} - OD_{\text{unexp}} = \log_{10} \left(\frac{I_{\text{unexp}} - I_{\text{bckg}}}{I_{\text{exp}} - I_{\text{bckg}}} \right), \quad (2.1)$$

where *exp* refers to whether the film was irradiated (*unexp*, unexposed) and *bckg* is the zero-light transmission quantity. This is the pixel value related to the white light value of the scanner used. *I* is the respective intensity value and is taken across each colour channel (red, green, blue). The uncertainties and possible errors can be calculated as similarly found in [41], where:

$$\sigma_{\text{net OD}} = \frac{1}{\ln(10)} \sqrt{\left(\frac{\sigma_{\text{unexp}}^2 + \sigma_{\text{bckg}}^2}{(I_{\text{unexp}} - I_{\text{bckg}})^2} + \frac{\sigma_{\text{exp}}^2 + \sigma_{\text{bckg}}^2}{(I_{\text{exp}} - I_{\text{bckg}})^2} \right)}. \quad (2.2)$$

The net OD values are plotted against the corresponding irradiated doses to determine the response curve as shown in figure 7. A calibration curve for each colour channel is obtained and given standard protocol, only the red channel is considered for this case. A curve fit was applied which enables the grey values from the irradiated films to be converted to dose and plotted against position to obtain the transverse beam distributions. Several scripts were developed in Matlab to automate this conversion process which are documented and accessible from [43].

The Medipix3 images were obtained by integrating the counts over all frames for each run and also evaluated using ImageJ. Minor artefacts in the central and cross pixels are observed from the interface between the four chips resulting in a larger than average effective pixel size for the cross pixels. Changes for the different grey value range and conversion of pixel size to mm were also made. Minimal post-processing was performed for the Medipix3 data, outlying pixels (noisy and dead) were interpolated from neighbouring pixels, the images summed and then a Gaussian blur of sigma equal to 1 pixel is applied in order to reduce pixel-to-pixel gain variations. The pixel values correspond to the number of times a pixel received charge above a low threshold of 5 keV. Due to charge sharing over multiple pixels, each proton produces more than one count on average.




Figure 7. Calibration curve for conversion of net OD (optical density) to dose (Gy).

For simplicity, we directly converted the counts to dose by scaling the pixel values to the calibrated film values in Gy by radial distance across the transverse plane, for each section irradiated at the same experimental location as the detector face. This preserves the linearity of the grey values as the grey pixel intensities correspond to numbers of hits and also correlate the magnitude of hits to a determined quantity: dose. These were scaled to the dose measured by the film and not directly calculated from counts or the resolved beam current recorded by Medipix3 due to the uncertainties with the electrometer. These were further perturbed by beam instabilities, particularly at the low currents during moments where there was a complete loss of beam. This was presumed to be caused by a dropout in the RF supplied to either one of the ‘dees’ of the cyclotron, resulting in a loss of the accelerating field between the two electrodes and therefore a disruption to the beam. Additionally, this may be related to the deflector which has deteriorated with use or from changes to the ion source when the facility concluded neutron therapy trials and was repurposed as a proton therapy clinic [23]. Maintenance cleaning of the cyclotron tank can also influence beam operation, resulting in changes to the beam characteristics.

2.6 Detector activation & temporal analysis

Activation is the process whereby incident particles transfer energy to target materials via nuclear interactions which results in a different nuclear energy state to before the interaction. These new nuclear states are typically unstable and go on to transition to a lower energy state by emitting energy in the form of gamma photons or can even fission into multiple fragments.

During irradiation with protons and heavy ions, any detector will become activated to some extent. Depending on the detector technology, flux and radiation type, this can result in radiation damage to the detector. This could be observed as variations in gain or noise baselines or the number

of unresponsive pixels. It may be possible for the detector to partially or fully recover from such changes with software compensation, by simply leaving the detector for a long time, annealing the sensor material at a higher temperature or other hardware level interventions.

The mean count rates for each run and between successive runs were calculated by averaging counts over frames where the beam was not active, to get an indication of the levels of activation within the high flux environment. Furthermore, due to the time resolved nature of the Medipix measurements as opposed to integrated over an entire measurement with film, it was possible to assess the beam current stability across a single measurement. This was further investigated by computing the Fourier transform in order to determine the frequency components of the mean counts over time. If the mean count rate purely varied randomly over time, the frequency component distribution would be flat. However, the distinct narrow peaks demonstrate that the mean counts and therefore the beam current oscillated non-randomly over time.

3 Measurements and results

3.1 Beam profiles

Beam profiles for each run were obtained and comparisons between both methods are shown for runs 6, 7, 14 and 15. Figure 8 shows the transverse distributions measured with the film (blue) and Medipix3 (red) within the integration zone for a standard (run 14, figure 8a) and at a much higher level (run 15, figure 8b) of accumulated dose.

At this location in the beamline, the beam distribution exhibits a double peak. This is caused by the beam stopper attached to the second scattering foil (figure 1) which attenuates a large proportion of the fluence in order to produce a flat, uniform beam at isocentre. For these runs, there is observable correspondence between both methods however the asymmetry of the distributions is distinct. Irradiating the film to a higher accumulated dose (figure 8b) appears to dilute the film beam distribution toward the directly measured profile by Medipix3. This may be attributed to a smaller range and less variation between the pixel intensity ranges. A corresponding smoothing effect can be seen for the Medipix3 profile.

The tilt seen in both plots suggests that the whole system was not precisely aligned at the centre of the beam axis but slightly translated. This results in skewed profiles and marked differences between the maximum dose at each apex. For both cases events were detected right through to the edges of the Medipix3 ($28 \times 28 \text{ mm}^2$) sensor, reflected by the clipped tails along the orthogonal plane. The comparative variations between the two different methods are reasoned to be a consequence of the differences in the position, analysis procedure and fundamental detection processes. In these plots, the Medipix3 measures wider profiles due the 3.5 cm gap between the film and detector. At this midway point in the beamline, the beam has yet to be significantly collimated and therefore is largely divergent: this is reflected in the differences between the penumbra regions.

Measurement runs 6 and 7 with the system placed at 9.5 cm and 30 cm after the treatment nozzle are shown in figure 9. A slightly wider profile was attained at 30 cm as the beam spreads out further downstream. Better agreement between the distributions was achieved for these cases, particularly at the lateral penumbra. A tilt in the same direction is also observed here, although the resulting profiles are less imbalanced. A flat, rectangular distribution is expected here but evidently, even a slight angular deviation results in an asymmetrical profile which is easily detectable by both methods.

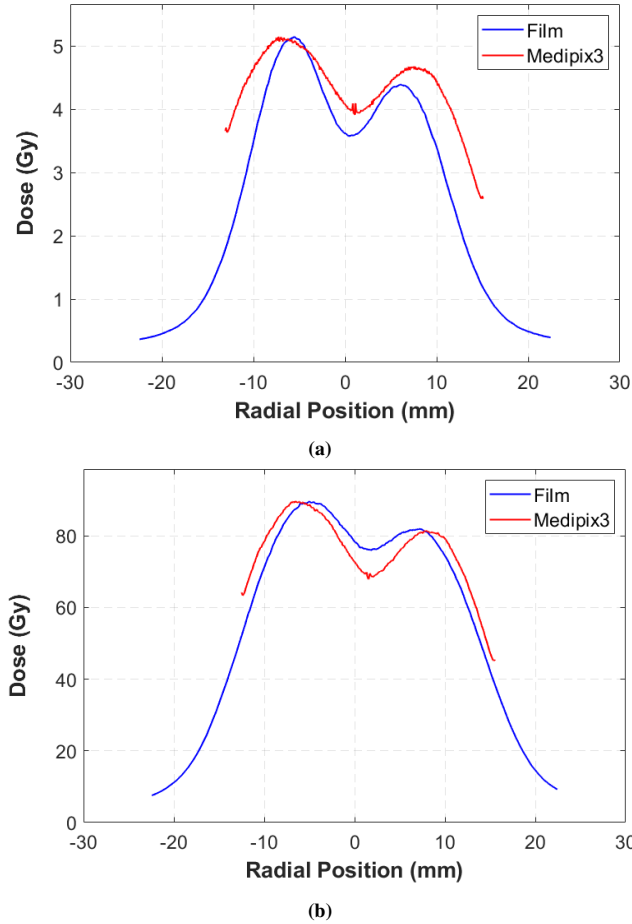


Figure 8. Beam profiles obtained with film and Medipix3 for (a) run 14 and (b) run 15 in the integration zone.

There are several considerations for these irregularities which are related to the image analysis, beam stability and quality on the day, as well as unknowns within the delivery system. It was found that the selection of the ROI in ImageJ was important and can affect the shape and resolution of the profile. As the beam distribution at the first position differs considerably to measurements after the nozzle, a rectangle ROI was chosen. This was specified such that there was total horizontal beam coverage and sufficient height to generate a smooth profile: the profile smears out if the ROI is too large. These ROI settings were kept the same for all images.

For all of the film profiles (figures 8–9) the tail regions do not extend to 0 Gy, this is due to the chosen fit for the calibration curves. We did not correct for this as primarily the central regions and highest points of dose were considered for comparison with the Medipix3 detector. Furthermore, at these lower levels of dose (< 5 Gy), there is a small margin of difference given the exponential fit derived from the calibration curve (figure 7).

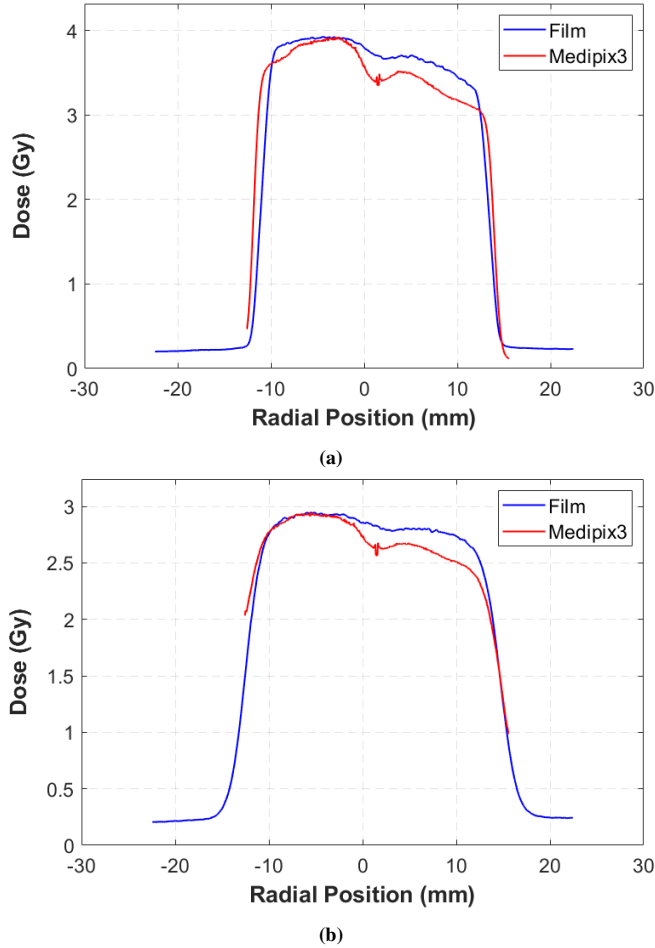


Figure 9. Beam profiles obtained with film and Medipix3 at (a) 9.5 cm and (b) 30 cm after the treatment nozzle.

Nonetheless, the film provided an intermediary means to correlate dose at all these positions to the ionisation chambers: these are indicative doses rather than precise measurements due to operational uncertainties described in section 2.5. The limitations surrounding the analysis and conversion process are indicated by the error bars in the calibration curves. A way to mitigate these would be to proceed with a more robust process, by strictly monitoring the film protocol to maintain uniformity with the film irradiation, development, scanning and analysis; or to attain profiles using well established or commercial methods. For example, using commercial software and hardware to determine the beam profiles from the irradiated film or by direct measurements with standard profile monitors or ionisation chambers. The possibility of determining the dose from counts is further described in section 3.2.

In addition to any misalignment of the system, it is also possible that small beam offsets at the start of the treatment line or aberrations with the shape or positioning of the beam stopper could impact the resulting beam distributions. This is less of an issue for the actual delivery of treatment as the beam is well conditioned to operate within the stable clinical currents with acceptable parameters, maintained by the QA diagnostics. Furthermore, due to the design of the delivery system and drift distance (~ 1.5 m) to the nozzle, further fluctuations are moderated by operational changes with the transport line optics. For these measurements, several issues were encountered as the beam currents requested were lower than the clinical range and the beam itself was not regulated as required for treatment.

3.2 Detector response

The comparative profiles demonstrate similar capabilities to provide spatial measurements, particularly at the penumbra regions. However, these measurements could be considered to be a ‘worst case’ scenario for the detector due to the relatively thick silicon sensor and the charge sharing effect contributing to a guaranteed increase in count rate. Both effects increase the difference between expected and measured count rate.

An estimation of the magnitude of effective charge sharing for single protons was made by using the ImageJ Shape Filter plugin which uses the IJBlob library [44]. This identifies clusters and returns many cluster properties such as area, aspect ratio, circularity etc. At the beginning of the measurements (between runs 2–3) using 967 identified clusters, the mean cluster area was 7.6 pixels, the mode area was 4 pixels with 187 counts and the minimum area was 3 pixels. The central part of the image was excluded for this first analysis because the beam was approximately centred on the detector which would have skewed the analysis. By the end of the measurement (run 11) using 1874 identified clusters, the mean cluster area increased to 9.3 pixels, the mode area decreased to 3 pixels with 348 counts and the minimum area stayed at 3 pixels.

This is evidence that the measured count rate by the Medipix3 could be corrected for by dividing the count rate by a factor between 3–9.3. In order to arrive at a more precise and reliable correction factor, dedicated measurements with very low flux of protons with precisely known energy would be necessary.

The minor detector artefacts observed in the central and cross pixels are due to a larger than average effective pixel size for those pixels, approximately a factor of 3 times larger. This varies depending on the exact manufacturing process which improves over time by shrinking this gap. In order to compensate for the larger pixel size with x-rays, a first order correction is to divide the raw counts by 2.8. However, the response of cross and central pixels cannot be very well corrected for by a linear scaling since there are non-linear effects including the count rate linearity. Therefore, only a basic linear scaling was applied to approximately match the counts of neighbouring pixels.

Figure 10 shows that the detector has a linear response across the entire tested range of beam currents from 0.012 to 1.97 nA. There is relatively large uncertainty of the average count rate due to the electrometer beam current measurements, dominated by the infrequent, manual readings. This is included in the residuals of the data and the variation in beam current which are both approximately in the order of 10%. Further measurements with improved beam current control would significantly reduce the count rate uncertainty and thus provide a reliable demonstration of count rate linearity with the Medipix3 detector. Nevertheless, the ultra low beam currents from runs 2 and 3 would not be possible to measure with other commonly used instruments with the precision of single protons;




Figure 10. Count rate linearity over all active pixels recorded at 10 FPS (frames per second) for 6 beam currents between $0.012^{+13\%}_{-9\%}$ and $1.97^{+4\%}_{-7\%}$ nA.

this detector therefore enables semi-destructive beam current measurements from single events to 10^{10} protons per second with the possible temporal resolution down to 0.5 ms.

The Medipix3 detector appears to have sufficient count rate linearity and sensitivity for beam characterisation and quality assurance measurements. Once cross-checked with another detector for accurate verification of the beam current, the impact of the ‘charge sharing’ effect on absolute dosimetry could be quantified accurately. A first order correction would be to simply divide the counts by the average cluster size, for these measurements this value would be 3–9.3 as mentioned in section 3.2. It is also anticipated that the Medipix3 chip configuration could be optimised significantly for high flux, clinical energy protons from the default x-ray (4–30 keV) settings. For future tests, if the count rates are low enough, it would be relevant to use the ‘charge summing mode’ on Medipix3 for furthering detector characterisation and possible mitigation of charge sharing. This has been shown to work as intended with relatively low energy x-rays [28], however the energy deposition from these 60 MeV protons is significantly higher, and so needs verification. Further work in this area could determine the energy deposited in the sensor and obtain definitive conversion factors to determine absorbed dose directly from the induced charge. This would also enable direct comparison with the dose distributions obtained by the film.

3.3 Detector activation

An overview of five background measurements during and between measurement runs is shown in table 2. These quantify how much the detector was activated over time, including the standard deviation, start and end count rates. The main count rate with the beam off is shown to increase over the course of the measurement day from 0 counts per second before the beam irradiation commenced to 1.3 counts per second at the end. The final measurement (run 15) was taken during a proton beam

interruption lasting 94 ms and so is different to the other measurements. The average count rate of 1.3 counts per second for this run is therefore indicative of the instantaneous detector activation between full beam intensity at 2.1 nA. This is relevant as a number of excited states have short half lives in the millisecond to second range which would not be evident in the other 4 measurements shown in this table.

Table 2. Overview measurement parameters of detector activation in counts per second with acquisition times and run ranges or numbers. The rows showing ranges of run numbers indicate the measurement was between run numbers.

| Run (#) | Time (s) | Mean (counts/s) | Standard deviation (counts/s) | Start (counts/s) | End (counts/s) |
|---------|----------|----------------------|-------------------------------|----------------------|----------------------|
| 2–3 | 50 | 7.2×10^{-4} | 1.3×10^{-3} | 9.5×10^{-4} | 6.7×10^{-4} |
| 5–6 | 20 | 4.4×10^{-2} | 2.8×10^{-2} | 4.7×10^{-2} | 4.0×10^{-2} |
| 8–9 | 10 | 6.4×10^{-2} | 3.8×10^{-2} | 6.7×10^{-2} | 6.1×10^{-2} |
| 11 | 8.3 | 9.8×10^{-2} | 6.4×10^{-2} | 9.9×10^{-2} | 9.7×10^{-2} |
| 15 | 0.094 | 1.3 | 0.63 | 1.5 | 1.2 |

3.4 Temporal analysis of the beam

A linear response was observed across the current range and the beam performance was reliable at higher clinical dose rates. Investigating the stability of the beam flux is a novelty in itself as it is not possible to measure the beam profile in real-time with millisecond resolution with other methods such as EBT3 film, beam current monitors or wire scanners. The beam current monitors could be cross-checked with these data if it was possible to log their outputs at high speed, the control system did not allow such control. These measurements verified that the beam profile was stable over time with a variable flux during a measurement. However this does not significantly affect the clinical treatment as CCC uses passive modulation with patient specific collimators for dose distribution control rather than a scanning pencil beam.

Figure 11 shows the mean counts integrated over all pixels over the typical measurement times for this work (a few minutes), at 100 FPS. The mean counts per frame are directly proportional to the instantaneous dose for a 10 ms period. The variations in the mean counts per frame show that the measured counts are not totally stable over time, the source of the instability could be due to the detector or the accelerator. The mean counts per frame varied over a measurement from the start to end by approximately 50 counts out of between 600 and 660 counts, which is $\sim 8\%$. However, recurring jitter were seen in the data where variations in the count rate over time appeared as small amplitude waves at low (few ms) and higher (of the order of hundreds of ms) frequencies. The mean counts per frame were qualitatively observed to oscillate at least one higher frequency during the measurements. These peculiarities were quantified with a frequency decomposition analysis.

In figure 12, the outliers are indicated by the three labelled frequency peaks. The 50.0 Hz component is assumed to be related to the U.K. AC mains electricity as it exactly matches the frequency. The 48.2 Hz component was not experienced during previous detector tests with x-rays or seen with any other measurements at the beamline. It is unclear if it is associated with the accelerator or detector: further measurements are necessary to identify a candidate. It is hypothesised that




Figure 11. Temporal stability of the beam over all pixels over nearly 5 minutes, recorded at 100 FPS, 26.3 Gy/minute in the integration zone.

the 1.80 Hz component is related to the ion source. As mentioned in section 2.5, the change to a proton therapy service required the nominal arc current (of the order of mA) of the ion source to be limited such that the beam currents produced were practical for treatment (nA range). Only the 50 Hz component was observed in later measurements performed at MedAustron [45] which is a significantly different facility and operates a synchrotron. As such, it is presumed that these other frequency peaks are specifically associated with the CCC cyclotron or may be due to interactions with beamline elements.

The significance of these rapid oscillations may be meaningful as they demonstrate something otherwise unforeseen which may have an impact on beam performance. For example, in passive delivery at CCC and the beam is modulated with a rotating wheel to generate a spread out Bragg peak, the graduations determine the longitudinal distribution of dose. Variations in the beam current could cause asynchronicity with the rotational frequency of the wheel and aberrations with the dose delivered for each step.

If the oscillations are present in regular operation, these measurements may reveal an underlying aspect of the cyclotron. An irregularity in the structure or accelerating process can influence the beam dynamics (orbit, stability, focusing effects from the magnetic field etc.) of the circulating protons. This may in turn impact the extracted beam quality such as the energy spread, emittance and dispersion. The jitter is presumed to be caused by a combination of different factors including the design, adjustments with the ion source and aged components. It is unknown if or how the beam is adjusted upon extraction. If there is a collimator at the exit then some properties of the beam would be regulated, such as the beam distribution in the transverse plane.

This analysis poses interesting avenues to study the beam structure and dynamics of accelerated CPT beams with the Medipix3 detector, for routine measurements or commissioning purposes.




Figure 12. Frequency components of the beam intensity as recorded by all pixels over nearly 5 minutes recorded at 100 FPS from figure 11. The 50.0 and 48.2 Hz components are the strongest followed by a 1.80 Hz component.

The timing capabilities would enable the possibility of examining the beam intensities, quality and stability of both continuous and pulsed beams. This would be also applicable for active delivery schemes such as pencil beam scanning.

Subsequent studies could investigate track projection over multiple pixels and the charge sharing as a function of beam angle relative to the detector. Efforts could be made to minimise the material in the beam path so that the detector becomes less activated over time; for this work, the large aluminium heatsink would have been the majority of the activated material and could be significantly reduced. Solutions could include a high RPM fan blowing air or dry ice over the detector or transferring the heat generated via vapor chambers, heat pipes or pyrolytic graphite to a heatsink just outside of the main beam path.

The detector settings could be significantly improved upon by tuning for high energy protons instead of the default, low energy x-rays. There are various software configurable DACs (digital-to-analogue converters) to control parameters which control front-end signal rising and falling times, signal baseline levels, thresholds among many others some of which are specifically designed with high count rates in mind like ‘Pole Zero Cancellation’ (R_{PZ}). In the same direction, ideally the detector would have a very thin sensor in order to induce much less charge sharing and still measure a very high amplitude signal. This would also reduce the detector activation as discussed in section 3.3.

4 Conclusions and outlook

Significant innovations in technology have resulted in greater accessibility and benefits of CPT worldwide. However, rapid developments in recent years have highlighted the need for improved diagnostic systems and performance capabilities, to fulfil the requirements of emerging, advanced delivery techniques. There is ongoing interest in the application of silicon detectors and in this work, we examined the potential of the Medipix3 hybrid pixel detector for CPT.

Experimental measurements were performed at the CCC 60 MeV clinical proton therapy beamline and compared with standard film dosimetry methods. Simultaneous irradiation of the film and detector placed at multiple locations in the delivery system allowed a direct comparison of the transverse beam distributions. There was general agreement between both methods, particularly at the lateral edges of the beam although several uncertainties resulted in experimental irregularities, also influenced by beam quality on the day.

Medipix3 has the capacity to measure the beam current by the detection of individual protons with millisecond scale temporal resolution and almost instantaneous readout time. This enabled the possibility to resolve otherwise unknown information about the CCC beam and accelerator, undetectable with typical instruments.

To facilitate its progression toward clinical implementation, further testing is recommended to characterise the cluster properties, signal uniformity, sensitivity across the detector, activation levels, dose rate thresholds, energy dependence, stability, spatial resolution and dosimetric calibration factors. These should be performed at different facilities, across the full treatment energy range and with different particle types.

We present the first set of tests using the Medipix3 detector technology in a clinical proton beam environment. This work explores the promising capabilities and versatility of Medipix3 for CPT, suggesting its possibility as a fast and efficient, future tool for routine dosimetry, commissioning and beam monitoring.

Acknowledgments

This project was supported by funding from the European Union FP7 grant agreement 215080, H2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 675265, OMA — Optimization of Medical Accelerators and the Cockcroft Institute core grant STGA00076-01. This work is part of the research programme of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organisation for Scientific Research (NWO). It was carried out at the National Institute for Subatomic Physics (Nikhef) in Amsterdam, the Netherlands.

References

- [1] P. Andreo, D. Burns, K. Hohlfeld, M. Huq, T. Kanai, F. Laitano et al., *Absorbed Dose Determination in External Beam Radiotherapy: An International Code of Practice for Dosimetry based on Standards of Absorbed Dose to Water*, Tech. Rep., International Atomic Energy Agency (IAEA), 2006 [http://www-naweb.iaea.org/nahu/DMRP/documents/CoP_V12_2006-06-05.pdf].

- [2] P.M. DeLuca, A. Wambersie and G. Whitmore, *Prescribing, recording, and reporting proton-beam therapy*, *J. ICRU* **7** (2007) 1.
- [3] J. Seco, B. Clasié and M. Partridge, *Review on the characteristics of radiation detectors for dosimetry and imaging*, *Phys. Med. Biol.* **59** (2014) R303.
- [4] S. Giordanengo, L. Manganaro and A. Vignati, *Review of technologies and procedures of clinical dosimetry for scanned ion beam radiotherapy*, *Phys. Med.* **43** (2017) 79.
- [5] C. P. Karger, O. Jäkel, H. Palmans and T. Kanai, *Dosimetry for ion beam radiotherapy*, *Physics in Medicine and Biology* **55** (2010) 193–234.
- [6] S. Giordanengo and H. Palmans, *Dose detectors, sensors, and their applications*, *Med. Phys.* **45** (2018) e1051.
- [7] B. Arjomandy, P. Taylor, C. Ainsley, S. Safai, N. Sahoo, M. Pankuch et al., *AAPM task group 224: Comprehensive proton therapy machine quality assurance*, *Med. Phys.* **46** (2019) e678.
- [8] V. Favaudon, L. Caplier, V. Monceau, F. Pouzoulet, M. Sayarath, C. Fouillade et al., *Ultrahigh dose-rate FLASH irradiation increases the differential response between normal and tumor tissue in mice*, *Sci. Transl. Med.* **6** (2014) 1.
- [9] S. Jolly, H. Owen, M. Schippers and C. Welsch, *Technical challenges for FLASH proton therapy*, *Phys. Med.* **78** (2020) 71.
- [10] K.P. Nesteruk and S. Psoroulas, *FLASH Irradiation with Proton Beams: Beam Characteristics and Their Implications for Beam Diagnostics*, *Appl. Sci.* **11** (2021) 2170.
- [11] E. Grusell and J. Medin, *General characteristics of the use of silicon diode detectors for clinical dosimetry in proton beams*, *Phys. Med. Biol.* **45** (2000) 2573.
- [12] M. Bruzzi, *Novel silicon devices for radiation therapy monitoring*, *Nucl. Instrum. Meth. A* **809** (2016) 105.
- [13] A.B. Rosenfeld, *Novel detectors for silicon based microdosimetry, their concepts and applications*, *Nucl. Instrum. Meth. A* **809** (2016) 156.
- [14] R. Schnuerer, J. Yap, H. Zhang, T. Cybulski, T. Smith, G. Haefeli et al., *Development of the LHCb VELO Detector Modules into a Standalone, Non-Invasive Online Beam Monitor for Medical Accelerators*, *Instruments* **3** (2018) 12.
- [15] A. Vignati, V. Monaco, A. Attili, N. Cartiglia, M. Donetti, M.F. Mazinani et al., *Innovative thin silicon detectors for monitoring of therapeutic proton beams: Preliminary beam tests*, *2017 JINST* **12** C12056.
- [16] J.T. Taylor, C. Waltham, T. Price, N.M. Allinson, P.P. Allport, G.L. Casse et al., *A new silicon tracker for proton imaging and dosimetry*, *Nucl. Instrum. Meth. A* **831** (2016) 362.
- [17] R. Ballabriga, M. Campbell and X. Llopart, *An introduction to the Medipix family ASICs*, *Radiat. Meas.* (2020) 106271.
- [18] M.G. Bisogni, G.A. Cirrone, G. Cuttone, A. Del Guerra, P. Lojacono, M.A. Piliero et al., *Medipix2 as a tool for proton beam characterization*, *Nucl. Instrum. Meth. A* **607** (2009) 48–50.
- [19] A. Rosenfeld, S. Alnaghy, M. Petasecca, D. Cutajar, M. Lerch, S. Pospisil et al., *Medipix detectors in radiation therapy for advanced quality-assurance*, *Radiat. Meas.* **130** (2020) 106211.
- [20] R. Ballabriga, M. Campbell, E. Heijne, X. Llopart, L. Tlustos and W. Wong, *Medipix3: A 64k pixel detector readout chip working in single photon counting mode with improved spectrometric performance*, *Nucl. Instrum. Meth. A* **633** (2011) S15.

- [21] Ashland GAFchromicTM Radiology, *GAFCHROMIC EBT*.
<http://www.gafchromic.com/gafchromic-film/radiotherapy-films/EBT/index.asp>.
- [22] A. Kacperek, *Proton therapy of eye tumours in the U.K.: A review of treatment at Clatterbridge*, *Appl. Radiat. Isot.* **67** (2009) 378.
- [23] J. Yap, J. Resta-López, A. Kacperek, R. Schnuerer, S. Jolly, S. Boogert et al., *Beam characterisation studies of the 62 MeV proton therapy beamline at the Clatterbridge Cancer Centre*, *Phys. Med.* **77** (2020) 108.
- [24] J. Yap, *TOPAS model of the 60 MeV Ocular Proton Therapy Beamline at the Clatterbridge Cancer Centre (CCC), Wirral, United Kingdom*, <https://github.com/jacyap/ClatterbridgeTreatmentLine>.
- [25] UCL, *TOPAS simulation model of the 62.5 MeV ocular proton therapy beamline at the Clatterbridge Cancer Centre*, <http://www.hep.ucl.ac.uk/pbt/wiki/Clatterbridge/TOPAS>.
- [26] E.T. Vitti, A. Kacperek and J.L. Parsons, *Targeting DNA Double-Strand Break Repair Enhances Radiosensitivity of HPV-Positive and HPV-Negative Head and Neck Squamous Cell Carcinoma to Photons and Protons*, *Cancers* **12** (2020) 1490.
- [27] M.J. Berger, J.S. Coursey, M.A. Zucker, J. Chang, S.M. Seltzer and P.M. Bergstrom, *NIST Stopping-Power & Range Tables for Electrons, Protons, and Helium Ions*, <https://www.nist.gov/pml/stopping-power-range-tables-electrons-protons-and-helium-ions>.
- [28] E.N. Gimenez, R. Ballabriga, M. Campbell, I. Horswell, X. Llopart, J. Marchal et al., *Study of charge-sharing in MEDIPIX3 using a micro-focused synchrotron beam*, *2011 JINST* **6** C01031.
- [29] J. Visser, M.v. Beuzekom, H. Boterenbrood, B.v.d. Heijden, J.I. Muñoz, S. Kulis et al., *SPIDR: a read-out system for Medipix3 & Timepix3*, *2015 JINST* **10** C12028.
- [30] B. van der Heijden, J. Visser, M. van Beuzekom, H. Boterenbrood, S. Kulis, B. Munneke et al., *SPIDR, a general-purpose readout system for pixel ASICs*, *2017 JINST* **12** C02040.
- [31] K. Akiba et al., *Characterisation of Medipix3 Silicon Detectors in a Charged-Particle Beam*, *2016 JINST* **11** P01011 [[arXiv:1509.02387](https://arxiv.org/abs/1509.02387)].
- [32] R. Castriconi, M. Ciocca, A. Mirandola, C. Sini, S. Broggi, M. Schwarz et al., *Dose–response of EBT3 radiochromic films to proton and carbon ion clinical beams*, *Phys. Med. Biol.* **62** (2017) 377–393.
- [33] S. Troja, E. Egger, P. Francescon, A. Gueli, A. Kacperek, M. Coco et al., *2D and 3D dose distribution determination in proton beam radiotherapy with GafChromic(TM) film detectors*, *Technol. Health Care* **8** (2000) 155.
- [34] S. Reinhardt, M. Hillbrand, J.J. Wilkens and W. Assmann, *Comparison of Gafchromic EBT2 and EBT3 films for clinical photon and proton beams*, *Med. Phys.* **39** (2012) 5257.
- [35] J. Sorriaux, A. Kacperek, S. Rossomme, J. A. Lee, D. Bertrand, S. Vynckier et al., *Evaluation of Gafchromic EBT3 films characteristics in therapy photon, electron and proton beams*, *Phys. Med.* **29** (2013) 599.
- [36] S. Devic, N. Tomic and D. Lewis, *Reference radiochromic film dosimetry: Review of technical aspects*, *Phys. Med.* **32** (2016) 541.
- [37] M.C. Battaglia, D. Schardt, J.M. Espino, M.I. Gallardo, M.A. Cortés-Giraldo, J.M. Quesada et al., *Dosimetric response of radiochromic films to protons of low energies in the Bragg peak region*, *Phys. Rev. Accel. Beams* **19** (2016) 064701.

- [38] S. Yonai, C. Arai, K. Shimoyama and N. Fournier-Bidoz, *Experimental Evaluation of Dosimetric Characterization of Gafchromic Ebt3 and Ebt-Xd Films for Clinical Carbon Ion Beams*, *Radiat. Prot. Dosim.* **180** (2018) 314–318.
- [39] E.Y. León Marroquin, J.A. Herrera González, M.A. Camacho López, J.E. Villarreal Barajas and O.A. García-Garduño, *Evaluation of the uncertainty in an EBT3 film dosimetry system utilizing net optical density*, *J. Appl. Clin. Med. Phys.* **17** (2016) 466.
- [40] M. Vadrucchi, G. Esposito, C. Ronsivalle, R. Cherubini, F. Marracino, R.M. Montereali et al., *Calibration of GafChromic EBT3 for absorbed dose measurements in 5 MeV proton beam and 60Co γ -rays*, *Med. Phys.* **42** (2015) 4678.
- [41] S. Devic, J. Seuntjens, G. Hegyi, E.B. Podgorsak, C.G. Soares, A.S. Kirov et al., *Dosimetric properties of improved GafChromic films for seven different digitizers*, *Med. Phys.* **31** (2004) 2392.
- [42] W. Rasband, *ImageJ. Image Processing and Analysis in Java*, <https://imagej.nih.gov/ij/>.
- [43] J. Yap, *FilmProtocol*. <https://github.com/jacyap/FilmProtocol>.
- [44] T. Wagner and H.-G. Lipinski, *IJBlob: An ImageJ Library for Connected Component Analysis and Shape Analysis*, *J. Open Res. Softw.* **1** (2013) e6.
- [45] N.J.S. Bal, C.S. Schmitzer, A. De Franco and S. Enke, *Medipix3 proton and carbon ion measurements across full energy ranges and at clinical flux rates in MedAustron IRI*, in proceedings of the 22nd International Workshop on Radiation Imaging Detectors, Ghent, Belgium, 27 June–1 July 2021 [[arXiv:2108.01908](https://arxiv.org/abs/2108.01908)].

3.3. CLINICAL FLUENCE MEASUREMENTS AT MEDAUSTRON, AUSTRIA

3.3 CLINICAL FLUENCE MEASUREMENTS AT MEDAUSTRON, AUSTRIA

This section explores the first use of this detector with protons and carbon ions at various energies and up to the maximum available beam intensity. Claus Schmitzer and Sascha Enke were critical for the accelerator control and experimental setup. Andrea De Franco reviewed the paper and gave crucial feedback for technical details which impacted the conclusions. I was responsible for everything else.

22ND INTERNATIONAL WORKSHOP ON RADIATION IMAGING DETECTORS
JUNE 27 – JULY 1, 2021
GHENT, BELGIUM

Medipix3 proton and carbon ion measurements across full energy ranges and at clinical flux rates in MedAustron IR1

N.J.S. Bal,^{a,b,*} C.S. Schmitzer,^c A. De Franco^c and S. Enke^c

^aNikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands

^bCERN, CH-1211 Geneva 23, Switzerland

^cEBG GmbH MedAustron, Marie-Curie-Straße 5,
2700 Wiener Neustadt, Austria

E-mail: navritb@nikhef.nl

ABSTRACT: The Medipix3, a hybrid pixel detector with a silicon sensor, has been evaluated as a beam instrumentation device with proton and carbon ion measurements in the non-clinical research room (IR1) of MedAustron Ion Therapy Center. Protons energies are varied from 62.4 to 800 MeV with 10^4 to 10^8 protons per second impinging on the detector surface. For carbon ions, energies are varied from 120 to 400 MeV/amu with 10^7 to 10^8 carbon ions per second. Measurements include simultaneous high resolution, beam profile and beam intensity with various beam parameters at up to 1000 FPS (frames per second), count rate linearity and an assessment of radiation damage after the measurement day using an x-ray tube to provide a homogeneous radiation measurement. The count rate linearity is found to be linear within the uncertainties (dominated by accelerator related sources due to special setup) for the measurements without degraders. Various frequency components are identified within the beam intensity over time firstly including 49.98 Hz with standard deviation, $\sigma = 0.29$, secondly 30.55 Hz $\sigma = 0.55$ and thirdly 252.51 Hz $\sigma = 0.83$. A direct correlation between the number of zero counting and noisy pixels is observed in the measurements with the highest flux. No conclusive evidence of long term radiation damage was found as a result of these measurements over one day.

*Corresponding author.

KEYWORDS: Beam-line instrumentation (beam position and profile monitors, beam-intensity monitors, bunch length monitors); Instrumentation for hadron therapy; Instrumentation for particle-beam therapy; Radiation damage to detector materials (solid state)

ARXIV EPRINT: [2108.01908](https://arxiv.org/abs/2108.01908)

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Experimental details | 2 |
| 3 | Results and discussion | 4 |
| 3.1 | Count rate linearity | 4 |
| 3.2 | Temporal beam intensity variation with frequency decomposition | 10 |
| 3.3 | Dead/unresponsive and noisy pixels over time | 10 |
| 3.4 | Carbon ion energy dependency on total counts | 11 |
| 3.5 | Radiation damage | 12 |
| 3.5.1 | During proton and carbon ion measurements | 12 |
| 3.5.2 | X-ray imaging 37 days after the proton and carbon ion measurements | 12 |
| 4 | Conclusions | 14 |
| A | Measurement overview | 15 |

1 Introduction

The Medipix3 chip is a hybrid pixel detector from the family of chips developed by the Medipix group at CERN [1] and have found many applications from electron microscopy [2] to spectral x-ray microCT [3]. The Medipix3 manual [4] contains detailed information on how the chip works.

Previous count rate linearity measurements of the Medipix3 with x-rays at a synchrotron [5], show the count rate to be linear up to the order of 10^8 photons/mm²/second. This is significantly more than the maximum expected particle flux rate in these measurements of approximately 10^4 to 10^6 particles/mm²/second. Therefore, assuming the front-end behaves similarly enough, the count rate linearity is expected to be comparable between proton, carbon ion and x-ray measurements. Given that the Medipix3 is designed for relatively low energy x-ray detection (<100 keV) and we are using 60+ MeV particles, this assumption is significant and has been verified with the measurements presented in this work.

This work follows a first measurement with protons in a clinical environment at the Clatterbridge Cancer Centre (CCC), U.K. [6]. It was demonstrated that the count rate was linear within the uncertainties from the beam variation. The CCC beam current measurements have large uncertainties due to recording the beam current only once per measurement while it was varying in the order of 10%.

2 Experimental details

Measurements took place in IR1 (irradiation room 1) which is a non-clinical research room, this is not a standard treatment room but solely for research purposes. The active detector area was $28 \times 28 \text{ mm}^2$, consisting of 4 Medipix3 chips in a 2×2 grid with a gap of approximately $220 \mu\text{m}$ between each chip. The sensor material was high resistivity silicon, with a thickness of $500 \mu\text{m}$ and pixel pitch (spatial resolution) of $55 \mu\text{m}$. The Medipix3 was primarily designed for x-rays and electron detection, however it detects any ionising radiation depositing $>5 \text{ keV}$ within a single pixel. The possible frame rates are 0–2000 FPS for 12 bit pixel counter depth (0 to 4095 counts per pixel per frame) in continuous R/W (Read/Write) mode (0 dead time). Other pixel counter depths and readout modes have different frame rate maximums which depend on the readout frequency of the chip, this is running at 200 MHz. Higher frame rates are possible with lower counter bit depths, a frame rate of up to 24 kHz is possible with a 1 bit counter depth; this mode results in a binary hit map. The detector frame rates used in this work were 50, 100, 1000 FPS in continuous R/W mode (0 dead time). Ideally one would use the maximum frame rate of 2000 FPS, however the readout computer used was not fast enough to reliably readout the data and so an upper limit of 1000 FPS was set. 50 and 100 FPS were used when the flux rate was low and the higher frame rate was not necessary in order to save disk space.

The high level experimental setup overview is shown in figure 1.




Figure 1. The equipment layout from left to right included a beam nozzle from which the particle beam is emitted, the detector was positioned in the beam isocentre using alignment line lasers. The beam isocentre is the central axis in space aligned to lasers where the centre of the target (within the patient) is positioned. The Medipix3 detector ($300 \times 100 \times 100 \text{ mm}^3$) contains a detector assembly where the radiation is measured and SPIDR v3.5 (Speedy Pixel Detector Readout) system. The SPIDR was connected to a readout computer via a 10 Gbit/s optical fibre where the data was stored. Finally the beam enters the water bottles of approximately 2 m depth so the Bragg peak was always within the bottles, these were used for shielding the robotic arm.

Due to driver limitations at the time of measurements, it was only possible to operate with SPM (Single Pixel Mode) and not CSM (Charge Summing Mode). This means that the pixels were individually counting, the charge was not summed over a 2×2 grid. At low flux, it is very likely that CSM would improve the PSF (Point Spread Function). The benefit would be maximised when the charge cloud covered an average of 4 pixels in a 2×2 grid. CSM requires inter-pixel

communication which takes time and therefore for the same front-end configuration, the count rate linearity for CSM degrades at lower count rates than SPM [5].

Only the first threshold was used for all measurements in this work. Table 1 shows the relevant front-end DACs (Digital to Analogue Converter), none were changed during the measurements. With this configuration, the threshold was set at approximately 5 keV. The threshold was left at the minimum but just above noise value which negatively impacts the PSF compared to setting the threshold at a much higher level given the energy deposition of the particle species and energies investigated in this work.

Table 1. Relevant front-end DACs used and kept constant during all measurements. DAC_DiscL has one value per chip in order. All other rows are common between all four chips.

| | Value | Units |
|-------------|----------------|-------|
| Threshold 0 | 42 | DAC |
| Gain Mode | High | N/A |
| DAC_DiscL | 71, 77, 75, 69 | DAC |
| FBK | 173 | DAC |
| GND | 121 | DAC |
| IKrum | 10 | DAC |
| Preamp | 150 | DAC |
| Shaper | 150 | DAC |

This detector system is highly optimised for x-ray detection between 4–30 keV. 500 μm of silicon limits the energy range up to approximately 30 keV because it is increasingly transparent as x-ray energy increases, see figure 9 [7]. The particle flux detected was 10^3 – 10^9 particles per second over the active area of the detector. The lower limit of detection is single particles. The upper limit for protons and carbon ions is a topic of investigation in this work. The maximum flux that the synchrotron can deliver is 10^{10} protons per second.

The proton energies used were 62.4, 148, 252, 800 MeV, and for carbon ions, 120, 260, 400 MeV/amu (atomic mass unit) ions were used. The motivation for using 800 MeV protons is for proton CT which is being investigated at MedAustron in order to measure the proton relative stopping power (RSP) with respect to water. The current method uses x-ray CT which gives the x-ray attenuation in HU (Hounsfield unit) which are then empirically mapped to relative proton stopping power, see Wayne D Newhauser et al. 2008 *Phys. Med. Biol.* **53** 2327 figure 3: “Relative linear stopping power for protons ($dE/dx|_w^x$) as a function of the scaled Hounsfield unit value (H_x , in units of HU_{sc}) in kVCT, where x denotes a material of interest and w denotes water” [8]. This figure contains one line plot consisting of three fitted straight lines with different slopes and intercepts with significant outliers; it is not a simple and clean linear relationship. This is one of the most important calibrations in this context because it feeds into every single dose distribution calculation of every patient. This conversion introduces one of the main sources of uncertainty in proton therapy treatment planning¹ and is an area of active research [9].

The degrader plates used in this work were: 10, 20, 50 & 100%. The degrader percentage is the nominal hole to surface ratio, ideally this would translate to a given percentage of incoming particles

¹ This is common knowledge in the Medical Physics community.

being transmitted relative to the total incoming number, while minimally affecting the energy. In reality the effective transmission is different. Some of the reasons behind the discrepancy are the self induced small charge forces and the slightly different accelerator optics used to compensate for such effects. As these effects vary with beam species and energy, a priori the flux reduction might differ from the nominal value according to the beam delivered. The quantitative effectiveness of the flux reduction of each degrader, which is a function of energy and species, is obtained with a series of monitors at each stage of the accelerator and in the irradiation room also. A degrader plate is a passive device, it is simply a steel plate with various numbers of holes of various sizes, this is referred to as a ‘pepper pot’ design. Ideally this produces an identical distribution of particles as the incoming beam. An example diagram of a pepper-pot degrader can be seen in figure 2, further technical details such as the specific geometry of the MedAustron degraders are not publicly available. After the degrader plates, the beam is shaped with beam optics which focus it at the isocentre and therefore are still compatible with pencil beam scanning. The degraders are one of the possible ways to regulate the particle fluence. Some measurements used degrader 100% (no degrader) and others used degrader plates (10%, 20% and 50%).




Figure 2. An example diagram of a pepper pot degrader plate where black is steel, the white circles are through-holes and the plate has a thickness in the axis of the diagram. There is a higher hole density in the centre because the beams are approximately 2D normal distributions and the beam primary axis is aligned with the centre of the degrader plates. The plate is in the order of 30 cm diameter and the holes are in the order between mm and cm.

3 Results and discussion

3.1 Count rate linearity

Accuracy and precision are both hard to determine because there were no more accurate or precise detectors or methods available to measure flux at this range.

For example, gas ionisation chambers which are commonly used in medical accelerators, saturate at relatively low count rates due to the relatively low charge carrier density in gas. They also suffer from not being able to detect low count rates because the signal induced by single particles is not detectable. Solid state detectors have a very high charge carrier density compared

to gas which means that the saturation limit is much higher. Gas based detectors benefit from being radiation hard because there is no crystalline structure to damage and the detection medium can be easily replaced.

Count rate linearity is important for accelerator calibration purposes because it makes calibration much simpler with lower associated fitting errors and maintains count rate precision over the entire range. Typically count rate linearity is assessed over either chip or pixel level. The linearity is assessed for all 4 chips for the simplicity of analysis and simplicity of the expected relationship between expected and measured count rate. In addition, the way these measurements were performed is the intended use for the detector.

The RayStation TPS (treatment planning system) from RaySearch Laboratories AB [19] was used to request 196 spots in order to irradiate the whole detector with a relatively homogeneous field. The FWHM (full width half maximum) spot size varied from 7 to 21 mm for protons and from 6.5 to 9.5 mm for carbon beams [20]. The requested spot weight was varied between 5×10^6 and 1×10^9 particles.

The summed count on the Medipix3 is the sum of every count on every frame. It is expected that this should be very linear if the relative uncertainties on the expected proton fluence are sufficiently low. This is measured over the largest possible range of proton fluence, using degrader plates 10, 20 and 50 for the lowest three fluences. The lowest three fluences were chosen due to limited measurement time and the desire to probe the lower count rate region where the detector was expected to have a better count rate linearity. Spot weights were determined based on the available pre-configured options in the control system, they are clinically relevant and are expected to be well within reasonable limits of all relevant systems.

Uncertainties in this measurement are described as follows. Firstly, not all the spots are entirely on the detector, given the relatively large FWHM of the beam at isocentre; some of the beam will not hit the detector. This effect could be quantified with information from the TPS. A brief analysis of a ‘single’ spot over 20 ms shows that this is likely not the case given the decreased linearity.

Secondly, it is observed at ultra low proton counts that 62.4 MeV protons produce a cluster of approximately 4–5 pixels. Ideally it would count once per proton or carbon ion. The detector will therefore count 4–5 times per 62.4 MeV proton as a result of this effect. As the energy increases, the cluster size decreases as less charge is deposited over the sensor depth which causes fewer pixels to count a hit. As the intensity increases, the events overlap within a single frame (minimum of 1 ms) and so cluster size per particle cannot be determined. This happens at the lowest clinical flux rates as they are high relative to the frame time. If one were able to increase the frame rate to infinity then the cluster size for high intensity measurements should not vary compared to low intensity measurements.

Cluster analysis could be done of the ultra low flux measurements to quantify this effect. This analysis was not done because the outcome would only give information about cluster sizes, their average, minimum and maximum values and would be affected by background measurements from both radioactive decay products from the activated sensor and surroundings and much less so, cosmic rays. It was therefore not expected to yield useful information as the focus is on count rate linearity rather than an absolute particle count. In order to measure an absolute particle count accurately, single cluster analysis would be necessary and is not achievable with this frame rate and flux.

Finally, there are shot-to-shot variations in the actual particle flux, as seen in table 2 for 800 MeV protons. Other proton energies and carbon ion energies were not scanned over with different degraders due to time constraints. It is assumed that this is due to variations in the extraction process or beam current from the synchrotron.

Table 2. The mean integrated counts over all frames with percentage uncertainties in the shot-to-shot intensity variation using 800 MeV protons.

| Degrader (%) | All measurements |
|--------------|--|
| 10 | 1.93×10^9 $^{+80\%}_{-84\%}$ |
| 20 | 6.85×10^9 $^{+8\%}_{-26\%}$ |
| 50 | 1.03×10^{10} $^{+3\%}_{-8\%}$ |
| 100 | 1.41×10^{10} $^{+3\%}_{-2\%}$ |

Figure 3 contains the data points with uncertainties in the expected number of particles only with two fits; one with the degrader measurements and the other without. $y = x$ is plotted as a reference to the naive expected relationship between the summed counts and the expected number of particles detected on the Medipix3. Summed counts means that all hits over the measurement are summed together, this is not counting clusters. Similar count rate linearity measurements were not performed at other proton energies or with carbon ions due to limited beam time. First order corrections to this naive expectation would include a simple geometric correction and a measurement dedicated to measuring the average cluster size.

In frames with a very low number of hits, a basic visual inspection of multiple single frames shows that the average cluster size for 62.4 MeV protons is in the 3–5 pixel range. Due to the clinical fluence, almost all of the clusters are overlapping in all measurements and so clusters cannot be reliably counted.

A treatment planning system is designed to deliver dose in the 3D distribution as programmed. The way it accomplishes this is very machine dependent. This has progressed from basic methods such as rotating a radioactive sample around a patient to state-of-the-art automated systems integrated with control systems. One such modern implementation of a TPS in a particle therapy context is RayStation®, which offers a solution for PBS (pencil-beam scan) as used at MedAustron. The output of this TPS is a raster scanned pencil beam whose profile is approximately a 2D normal distribution.

Suppose one would like to uniformly irradiate a given area with a raster scanned pencil beam and have negligible dose outside of the designated irradiation area, it is clear that one would need to modulate the intensity of the beam over time. If we assume that the intensity of the beam can be instantaneously ramped up to the maximum and down to zero, one would expect to see a cumulative dose distribution as seen in figure 4(A). As one increases the number of spots and reduces the inter-spot spacing, this grid of 2D normal distributions would tend towards the distribution in figure 4(B) which shows a decreased dose at the edges of the designated irradiation area. In this case, the irradiation zone was specified to be larger than the detector in order to approximate a uniform irradiation. This was possible because this is a non-clinical measurement which did not require minimising dose outside of the detector. A visualisation of this is shown in figure 4 as the dashed inner black squares representing the detectors are positioned arbitrarily




Figure 3. The count rate linearity of all pixels integrated over all frames for the measurement using 62.4 MeV protons. Fits are shown both with degraders ($a = 1.20502$, b is fixed at zero) and without degrader measurements ($a = 1.14307$, b is fixed at zero) with the associated relative percentage residuals from the fits. The reduced chi squared value for the fit with all data points is $\chi^2_\nu = 91$ and for the fit with no degrader measurements is $\chi^2_\nu = 4.6$. The intercept parameter (b) is fixed at zero because zero counts are measured when the accelerator is not delivering any particles and counts from sensor activation are negligible at a few hundred counts per second. The degrader measurements are the first three data points from the left and are not included due to large systematic uncertainties as described in the text and seen in figure 5. The fit algorithm was the non-linear least-squares fit (NLLS) Marquardt-Levenberg.

within the designated irradiation area. In typical clinical settings, the designated irradiation area is determined by the treatment plan.

Secondly, systematic errors are introduced because each proton triggers more than a single pixel as a result of the current detector front-end configuration and sensor thickness. It is known that the detector will count too many protons from previous work [6] with 62.4 MeV protons, this was also observed at very low count rates during initial testing at MedAustron. A first order correction for this effect could be given by calculating mean cluster size resulting from 62.4 MeV protons in 500 μm silicon of either simulated or measured data. The measurements would require extremely low flux in order to have average occupancy of a frame low enough to observe individual clusters. This measurement is not trivial because producing this low flux (approximately 1 kHz) of 62.4 MeV protons is completely out of the design parameters of medical accelerators. Methods




Figure 4. An example dose distribution for a TPS (treatment planning system) using a PBS (pencil-beam scan) with a given spot FWHM (full width half maximum). (A) (Left): large inter-spot spacing and (B) (right): small inter-spot spacing and many spots. The outer solid black square represents the bounding area for irradiation, the inner dashed black squares represent the detector and the grey colour shows the relative dose in arbitrary units.

used to reduce the beam intensity broaden and shift the peak energy down, therefore the energy is no longer known.

There is also an uncertainty in the summed counts from the activation of the silicon sensor, the Medipix3 chips and other surrounding material, including the readout system, aluminium cooling block and foam topped table. It was not possible to distinguish between the sources of the activation using the Medipix3 itself; alpha, beta and gamma decays were observed based on the shape of tracks. The components in the direct beam (silicon sensor, Medipix3 chips and the aluminium cooling block) are expected to be the most activated components by far. The magnitude of activation varies over time due to a combination of random radioactive decay rates, exponentially decreasing activity and the particle beam would cause increasing activation. This is expected to be proportional to dose in the sensor, chip and surrounding materials since the half-life time was in the order of days. Due to the highly mixed radiation field from various decay chains, an estimation of the dose is not made. The detector is activated to an average of 378 counts per second over the whole detector measured over 1000 seconds between two 800 MeV proton measurements half way through the day.

These effects introduce systematic uncertainties, as supported by the quality of the linear fit ($\chi^2_\nu = 4.6$) for the dataset ignoring the degrader measurements. This assumes an uncertainty in the expected number of particles of 3% due to shot-to-shot variation in the number of particles coming from the beam nozzle.

Figure 5 shows the count rate linearity against the degrader nominal percentage with 800 MeV protons, with and without measurements with degrader 10. Table 2 summarises the data from figure 5, showing the shot-to-shot intensity variation using 800 MeV protons with all 4 degrader plates (10, 20, 50, 100%). Given more measurement time, we would investigate more clinically relevant proton beam energies and also carbon ion beams. 800 MeV protons were prioritised as this was a new modality for MedAustron which was undergoing testing at the time. It is expected

that this would also be linear within the uncertainties for lower energy proton beams and for carbon ion beams also.




Figure 5. The integrated counts over all frames against the degrader percentages 10, 20, 50 and 100% using 800 MeV protons. The linear fit only uses degrader 100 measurements and fixes the y-intercept to 0. No error bars are plotted because the degrader percentage is a discrete quantity and the summed counts over all frames uses the Poisson error, the square root of the counts as the error which results in errors in the order of 0.001% which is not visible on this scale. There are 33 data points in this plot.

The percentage uncertainties in superscript and subscript in table 2 show the expected 2 to 3% uncertainty in beam intensity as measured by other devices at MedAustron during commissioning when no degraders were used. Degradation 100% means that no degrader was used. For the other degraders, larger uncertainties were measured. Degradation 50 consistently reduces the particle count to ~75%. Similarly, degradation 20 reduces the particle count to ~50%. Measurements with degradation 10 show greater relative variation, anywhere from 2–25% of the expected particle count is detected. This data indicates that the various methods used to obtain lower beam currents are typically producing significantly more protons than expected based on the degradation percentage alone. It is possible that the degradation 100 measurements are suffering from saturation effects or other such losses, however this is unlikely given consistently the higher than expected counts for the degradation 10, 20 & 50 measurements (the first three from the left) in figure 3. A mismatch between the targeted and achieved particle count has been documented by L. Adler (tables 6.9 & 6.10), although these measurements may not be comparable since they are measured in different locations along the beam-line. The large systematic uncertainties could be addressed with repeat measurements by measuring certain accelerator parameters. The total number of extracted particles can be calculated non-invasively via the differential of the main ring current transformer plus there is an active measurement of the DDS (dose delivery system) giving the exact number of particles deployed. This is calibrated in the medical energy ranges for protons: 62–252 MeV and for carbon

ions: 120–400 MeV/amu. This was not measured for this work and retrieving such information from the log files is no longer possible.

If unaccounted for, this mismatch would adversely affect patient treatment because when degraders are used, one would actually be delivering more dose than intended. However, the accuracy of the degrader ratings is actually not important due to the Dose Delivery System (DDS) which is an essential component of every medical particle accelerator. The DDS is designed to measure the number of particles delivered and move to the next spot as soon as the number of particles delivered on that spot matches the requested number. Therefore, the degraders should not affect the dose delivery in actual patient treatment.

3.2 Temporal beam intensity variation with frequency decomposition

The mean count of every raw frame over time is calculated after which a Fourier transform is applied, resulting in frequency components as shown in figure 6. The 49.98 Hz component is exactly the same frequency as Austrian AC mains electricity, this appears to be the most likely explanation for this frequency. The 150 Hz component is likely to be the 3rd harmonic of the mains AC frequency. The 30.55 and 252.51 Hz components are related to the spill ripples which are caused by power converter ripples and not the DDS (dose delivery system). The DDS is used for these measurements to scan over the whole detector surface.




Figure 6. A histogram of the frequency components of the intensity variation of the beams using the 1000 FPS data only. The top three in amplitude (as marked with the arrows) are 49.98 Hz with standard deviation, $\sigma = 0.29$, secondly 30.55 Hz $\sigma = 0.55$ and thirdly 252.51 Hz $\sigma = 0.83$. The DC component (first bin of the FFT data), has been removed for visualisation purposes.

3.3 Dead/unresponsive and noisy pixels over time

The total number of pixels is 262144 (512×512 pixels). Therefore, the percentage of dead pixels over time varies between only 0.019–0.158% and the percentage of noisy pixels varies between 0.000–0.011% in figure 7. A slow and precise equalisation (a procedure to flatten pixel noise baseline levels over the chip [10]) was started at 09:30 and finished an hour after. This procedure was run to correct for perceived damage as well as an increased number of bad pixels, i.e. noisy, dead, unresponsive pixels. There was a significant reduction in bad pixels compared to pre-irradiation levels, and these numbers were relatively stable for subsequent runs. There is a strong

correlation between bad pixels for the high flux region between 09:15 and 09:30 whereas the other points show low correlation.




Figure 7. The absolute number of zero counting (dead or temporarily unresponsive) pixels and noisy pixels over time (measurement start times). The percentage of dead/unresponsive pixels varies between 0.019 and 0.158% of the total number of pixels ($512 \times 512 = 262144$). The percentage of noisy pixels varies between 0 and 0.011% of the total number of pixels.

It is possible that during this high flux region in time, the Si-SiO₂ interface between the pixel implants and the sensor bulk accumulated charge to the point where the affected pixel's noise baseline was shifted out of range. If a pixel's noise baseline is shifted out of range, it will either never respond or would be noisy at that threshold. This effect is observed by measuring increased leakage current during relatively high intensity x-ray irradiation resulting in higher noise baselines over time. The leakage current decreases back to the pre-irradiation levels consistently.

Given that an equalisation and some time with no radiation fixed this issue and similar behaviour is observed at a much slower rate with x-rays, this hypothesis appears to be consistent. More studies would be relevant to probe the exact underlying mechanism.

3.4 Carbon ion energy dependency on total counts

Carbon ion spills lasting 22 to 25 seconds at 120, 260 and 400 MeV/amu were measured in order to verify the expectation that carbon ion energy should be inversely proportional to the total counts recorded. As the carbon ion energy increases, the probability of interaction per unit length in the silicon decreases. As table 3 shows, the carbon ion energy is inversely proportional to the total number of counts. This dependency was measured to be linear within this region and the number of carbon ions requested was kept constant.

Table 3. The carbon ion energy dependency on the total number of counts per spill. A linear fit on this data results in a slope of -1.26×10^7 , an intercept of 1.27×10^{10} , with $R^2 = 0.978$.

| Energy (MeV/amu) | Total number of counts |
|------------------|------------------------|
| 120 | 1.10×10^{10} |
| 260 | 9.74×10^9 |
| 400 | 7.53×10^9 |

3.5 Radiation damage

3.5.1 During proton and carbon ion measurements

Figure 8 shows an approximate 10% response decrease in the counts over the selected region of interest. Ideally, the detector would have been uniformly irradiated via requests to the DDS, which scans over the surface using a number of spots at a particular target spot weight. Variations in the extracted number of particles per spot and between spots introduce uncertainty to the homogeneity of the delivered particle distribution, these effects were not calculated.




Figure 8. Left: an integrated image of all 4 chips over all frames with 62.4 MeV protons scanned over the surface using the treatment planning system, a uniform exposure was intended. The darker oval shaped region in the centre shows the area suspected of radiation damage. The view is optimised for non-cross pixels which is why the cross pixels are all white. The x -axis shows detector columns and the y -axis shows detector rows. Right: a histogram of the integrated image, showing the regions of interests: dead pixels, most pixels and the cross pixels.

3.5.2 X-ray imaging 37 days after the proton and carbon ion measurements

The aim of this measurement was to find evidence of radiation damage and if it was consistent with supposedly uniform proton irradiation. This is achieved by using a cone beam x-ray tube to produce a relatively homogeneous radiation field. X-rays are the lowest energy, individually detectable particles with the Medipix3, can be produced at high rates ($> 10^{11}/s$) with common x-ray tubes and do not damage the detector at this flux. X-rays are therefore appropriate for investigating the homogeneity of the detector response over the surface after irradiation with particles causing nuclear interactions in the silicon sensor such as protons and carbon ions in the MeV range and above.

Relevant parameters: x-ray tube peak voltage 50 kVp, tube current 0.92 mA, 5 minute exposure and the detector was 15 cm from the tube exit window. These parameters were chosen in order to produce a homogeneous field with a very high number of x-rays resulting in a very low statistical uncertainty of $<0.001\%$. The tube peak voltage and the currents are the maximum possible for this Jupiter 5000 series x-ray tube [11].

One can observe from figure 9 that there is not a decrease in response in the centre of the image consistent with the uniform proton irradiation image 8. This implies that either the proton response is different from the x-ray response, the supposed radiation damage observed during irradiation, as seen in figure 8, has annealed or a combination of both. Damage is expected there due to the proton beam being fixed on the central area in the first part of the day.




Figure 9. Left: a raw x-ray image of all 4 chips over 5 minutes where the x-axis shows detector columns and the y-axis shows detector rows. The view is optimised for the majority of the pixels, not the cross pixels. Right: a histogram of the integrated image, showing the regions of interests: dead pixels, most pixels and the cross pixels.

As for detector uniformity, 9.9% variation in counts across the detector surface is typical in this configuration and does not indicate radiation damage. 95% of counts are within 19% of the mean, averaged over the 4 chips, excluding the cross. No response variation is observed in the centre.

Regarding the two patterns visible in figure 9, the wave-like pattern across the 4 chips is known to be due to doping concentration variation in the p-on-n silicon sensor during the crystal growth [12]. P type doping with boron is used for the implant and phosphorus is used for N type doping for the n bulk. A single silicon sensor cut from a single wafer is bump bonded to 4 Medipix3 chips in a 2×2 grid in this configuration. This explains why the waves are continuous across all the chips. The vertical gradient is due to the detector being close enough to the x-ray tube that the cone beam has significant intensity variation.

The number of dead/unresponsive pixels was 112 (0.043%) at the start of the proton and carbon ion measurements and is 123 (0.047%) in this x-ray test. There are only 0.0004% more dead/unresponsive pixels than before any proton and carbon ion irradiation. There is variation in this number as shown in section 3.3. Given that the variation in that number is much larger than the difference here, it is not expected that this difference is significant.

An average increase of 2.5 DAC (digital-to-analogue) units is observed; the mean of the noise of the chips is 1% more than before the measurement. Given that the temperature was not monitored, this is within the expected variation and is therefore not a conclusive indication of increased chip noise. Simulations of the Medipix3 chip response to temperature were done during the design process [13, 5.4.5 & appendix III] where Ballabriga simulates that the shaper output signal should vary by $0.16\%/^{\circ}\text{C}$ with nominal settings. The shaper output signal magnitude is directly proportional to the aforementioned DAC values. This would imply a temperature difference of 6.25°C which is consistent with expectations.

In conclusion, no significant increase in dead/unresponsive pixels is observed. The cause of the wavy pattern is well known. The total variation in response over the detector is in the normal range. No significant increase in chip noise is observed. No reasonable estimation of detector lifetime can be calculated from this. Studies have investigated the radiation hardness of the Medipix3 chip with x-rays and neutrons, for x-rays the chip was still operational after 460 MRad at a high dose rate of 3.5 kGy/s [14] and for neutrons, the chips tolerated 1 MeV neutron equivalent dose of $5 \times 10^{14}\text{ cm}^{-2}$ [15]. This is far in excess of typical requirements for space grade radiation hardness of 300 kRad and more in line with LHC inner tracker requirements of 300 MRad. Simulations are required to calculate the dose in the detector, in order to do this, one would need to know the ratios of elements in order to estimate the relation between particles of a particular energy and the absorbed dose in the chip, this was not in the scope of this analysis.

Additional tests at clinically relevant dose rates ($>1\text{ Gy/s}$ in water) with proton and carbon ion beams are needed to determine the lifetime of the detector in this radiation environment. Due to the limited beam time and detector supply, this is not in the scope of this study.

Based on this study, it is expected that the Medipix3 chips would be used as beam profile monitors for relatively infrequent quality assurance (QA) measurements. Additionally, these chips are relatively effective at verifying the accuracy and precision of the DDS. This was the first measurement to simultaneously show the spatial and temporal distribution of delivered protons and carbon ions at MedAustron; the Medipix3 fills that niche.

Many different detector geometries could be suitable for this application depending on the exact intended use, for example, one could exclusively measure the edges of the beam so the system would not disrupt the beam and could run indefinitely. The other extreme would be to use a retractable large area detector for occasional quality assurance measurements.

Scaling of Medipix3 based detectors is possible with TSV (Through Silicon Via) technology which enables $N \times N$ scaling, subject to sensor wafer size primarily [16–18] with a 0.8 mm non-active area, the periphery. The Medipix4 is being designed to further improve on this by eliminating the non-active area, enabling 100% active detection area.

4 Conclusions

The Medipix3 chip with a $500\text{ }\mu\text{m}$ silicon sensor has been used for a series of measurements using high energy protons and carbon ions with a wide range of particle flux and energies. Protons with energies of 62.4, 148, 252 and 800 MeV were used at flux rates between 10^4 and 10^8 . 120, 260, 400 MeV/amu carbon ions were used at flux rates varying between 10^7 and 10^8 carbon ions per second impinging on the detector surface.

The temporal beam intensity variations were decomposed into frequency components showing several peaks including Austrian mains frequency and two others which are related to the spill ripples in the synchrotron. None of these degrade the patient treatment due to the design of the Dose Delivery System (DDS).

During the period of highest flux, the number of zero counting and noisy pixels increased rapidly and were correlated. After running a software procedure to equalise the pixel response over the matrix, the number of zero counting and noisy pixels returned to approximately pre-irradiation levels. Further studies would be relevant to investigate this effect.

There is evidence that the Medipix3 can be used as a beam instrumentation device. It shows good count rate linearity with 62.4 MeV protons over the full flux range available, reliable performance at 1000 FPS and is sensitive to single particles. Proton and carbon ion beams have been measured at the full energy range, respectively 62.4 to 800 MeV and 120 to 400 MeV/amu. No conclusive evidence of radiation damage was observed, further measurements are necessary to determine detector lifetime. The Medipix3 front-end settings (DACs) could be optimised and tested from the default low energy x-ray (<30 keV) configuration with more beam-time.

Acknowledgments

This work is part of the research programme of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organisation for Scientific Research (NWO). It was carried out at the National Institute for Subatomic Physics (Nikhef) in Amsterdam, The Netherlands.

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant Agreement No. 675265 — OMA (Optimization of Medical Accelerators).

A Measurement overview

An overview of all measurements is displayed in tables 4 and 5.

Table 4. Measurements overview, 1 of 2. Spot weight has the units of numbers of particles specified.

| Run | Sub-run | Particle | Energy (MeV) | Degrader (%) | Spot weight |
|-----|---------|----------|--------------|--------------|-------------|
| 1 | | Proton | 800 | 20 | N/A |
| 2 | | Proton | 800 | 20 | N/A |
| 3 | | Proton | 800 | 20 | N/A |
| 4 | Test | Proton | 800 | 10 | N/A |
| | 1 | Proton | 800 | 10 | N/A |
| | 2 | Proton | 800 | 10 | N/A |
| 5 | 1 | Proton | 800 | 10 | N/A |
| | 2 | Proton | 800 | 10 | N/A |
| | 5 | Proton | 800 | 10 | N/A |
| | 7 | Proton | 800 | 10 | N/A |
| | 8 | Proton | 800 | 10 | N/A |
| | 9 | Proton | 800 | 10 | N/A |
| | 10 | Proton | 800 | 10 | N/A |
| 6 | 1 | Proton | 800 | 20 | N/A |
| | 2 | Proton | 800 | 20 | N/A |
| | 3 | Proton | 800 | 20 | N/A |
| | 4 | Proton | 800 | 20 | N/A |
| | 5 | Proton | 800 | 20 | N/A |
| | 6 | Proton | 800 | 20 | N/A |
| | 7 | Proton | 800 | 20 | N/A |
| | 8 | Proton | 800 | 20 | N/A |
| | 9 | Proton | 800 | 20 | N/A |
| | 10 | Proton | 800 | 20 | N/A |

Table 5. Measurements overview, 2 of 2. Spot weight has the units of numbers of particles specified.

| Run | Sub-run | Particle | Energy (MeV) | Degrader (%) | Spot weight |
|-----|---------|------------|--------------|--------------|-----------------|
| 7 | 1 | Proton | 800 | 50 | N/A |
| | 2 | Proton | 800 | 50 | N/A |
| | 3 | Proton | 800 | 50 | N/A |
| | 4 | Proton | 800 | 50 | N/A |
| | 5 | Proton | 800 | 50 | N/A |
| | 6 | Proton | 800 | 50 | N/A |
| | 7 | Proton | 800 | 50 | N/A |
| | 8 | Proton | 800 | 50 | N/A |
| | 9 | Proton | 800 | 50 | N/A |
| | 10 | Proton | 800 | 50 | N/A |
| | 11 | Background | 800 | N/A | N/A |
| 8 | 0 | Proton | 800 | 100 | N/A |
| | 2 | Proton | 800 | 100 | N/A |
| | 3 | Proton | 800 | 100 | N/A |
| | 4 | Proton | 800 | 100 | N/A |
| 9 | | Proton | 62 | 20 | 5×10^6 |
| 10 | | Proton | 148 | 20 | 1×10^7 |
| 11 | | Proton | 252 | 20 | 1×10^7 |
| 12 | | Proton | 62 | 10 | 1×10^6 |
| 13 | | Proton | 62 | 100 | 1×10^8 |
| 14 | | Proton | 62 | 100 | 5×10^7 |
| 15 | | Proton | 62 | 50 | 1×10^7 |
| 16 | | Proton | 62 | 100 | 1×10^9 |
| 17 | | Proton | 62 | 100 | 5×10^8 |
| 18 | | Carbon | 120 | 20 | N/A |
| 19 | | Carbon | 400 | 20 | N/A |
| 20 | | Carbon | 260 | 20 | N/A |
| 21 | | Carbon | 120 | 100 | N/A |

References

- [1] R. Ballabriga Suñé et al., *The Medipix3RX: a high resolution, zero dead-time pixel detector readout chip allowing spectroscopic imaging*, [2013 JINST 8 C02016](#).
- [2] G.W. Paterson, R.J. Lamb, R. Ballabriga Suñé, D. Maneuski, V. O'Shea and D. McGrouther, *Sub-100 nanosecond temporally resolved imaging with the Medipix3 direct electron detector*, [Ultramicroscopy 210 \(2020\) 112917](#).
- [3] J.P. Ronaldson, *Quantitative soft-tissue imaging by spectral CT with Medipix3*, Doctoral Dissertation (2012).
- [4] R. Ballabriga Suñé and X. Llopart, *Medipix3RX manual*, CERN, Medipix Group (2012).
- [5] E. Fröjdén et al., *Count rate linearity and spectral response of the Medipix3RX chip coupled to a 300 μm silicon sensor under high flux conditions*, [2014 JINST 9 C04028](#).

- [6] J. Yap et al., *Beam characterisation using MEDIPIX3 and EBT3 film at the Clatterbridge proton therapy beamline*, in proceedings, *10th International Particle Accelerator Conference (IPAC2019)*, Melbourne, Australia, 19–24 May 2019, THPMP033.
- [7] J. Susini et al., *New challenges in beamline instrumentation for the ESRF upgrade programme phase II*, *J. Synchrotron Radiat.* **21** (2014) 986.
- [8] W. Newhauser, A. Giebeler, K. Langen, D. Mirkovic and R. Mohan, *Can megavoltage computed tomography reduce proton range uncertainties in treatment plans for patients with large metal implants?*, *Phys. Med. Biol.* **53** (2008) 2327.
- [9] C.G. Ainsley and C.M. Yeager, *Practical considerations in the calibration of CT scanners for proton therapy*, *J. Appl. Clin. Med. Phys.* **15** (2014) 202.
- [10] J. Rinkel, D. Magalhães, F. Wagner, E. Fröjdh and R. Ballabriga Suñé, *Equalization method for Medipix3RX*, *Nucl. Instrum. Meth. A* **801** (2015) 1.
- [11] Oxford Instruments, *Jupiter 5000 series radiation shielded x-ray tube*, xray.oxinst.com/x-ray-tube-products/x-ray-tube-assembly-radiation-shielded/jupiter-5000-series (accessed 12 March 2020).
- [12] S.Y. Lim et al., *Applications of photoluminescence imaging to dopant and carrier concentration measurements of silicon wafers*, *IEEE J. Photovoltaics* **3** (2013) 649.
- [13] R. Ballabriga Suñé, *The design and implementation in 0.13 μm CMOS of an algorithm permitting spectroscopic imaging with high spatial resolution for hybrid pixel detectors*, Ph.D. Thesis, CERN (2009).
- [14] R. Plackett, X. Llopart, R. Ballabriga, M. Campbell, L. Tlustos and W. Wong, *Measurement of radiation damage to 130 nm hybrid pixel detector readout chips*, in *Topical Workshop on Electronics for Particle Physics* (2009).
- [15] K. Akiba et al., *Characterisation of Medipix3 silicon detectors in a charged-particle beam*, [2016 JINST 11 P01011](#).
- [16] T. Tick and M. Campbell, *TSV processing of Medipix3 wafers by CEA-LETI: a progress report*, [2011 JINST 6 C11018](#).
- [17] C. Ponchut et al., *SMARTPIX, a photon-counting pixel detector for synchrotron applications based on Medipix3RX readout chip and active edge pixel sensors*, [2015 JINST 10 C01019](#).
- [18] M. Sarajlić, D. Pennicard, S. Smoljanin, T. Fritzsche, K. Zoschke and H. Graafsma, *Progress on TSV technology for Medipix3RX chip*, [2017 JINST 12 C12042](#).
- [19] D. Bodensteiner, *RayStation: external beam treatment planning system*, *Med. Dosim.* **43** (2018) 168.
- [20] M.T.F. Pivi et al., *Status of the carbon commissioning and roadmap projects of the MedAustron ion therapy center accelerator*, in proceedings of *IPAC (2019)*, doi:[10.18429/JACoW-IPAC2019-THXXPLS1](https://doi.org/10.18429/JACoW-IPAC2019-THXXPLS1).

4

X-ray Computed Tomography

This chapter serves as an introduction of x-ray CT (computed tomography), types of artefacts that can be observed in CT reconstructions, the experimental setups which I made for XRF (x-ray fluorescence), 2D imaging and CT measurements. This experimental setup was developed and used for all measurements using x-rays in this work.

4.1 INTRODUCTION

With Johann Radon finding a mathematical solution to the tomographic problem in 1917 and subsequently Sir Godfrey Hounsfield reconstructing the first human head in the late 1960s; it is safe to say that computed tomography is a very well established field. This introduction is therefore exclusively intended to give the reader an overview. Further details can be found in the literature in great detail.

Intensity attenuation due to photo electric absorption can be described by the ‘Beer-Lambert Law’. I is the intensity after traversing a material of thickness, Z , with a linear attenuation coefficient dependent on energy, $\mu(E)$, and an initial intensity, I_0 :

$$I = I_0 e^{-\mu(E)Z} \quad (4.1)$$

rearranges to:

$$\mu(E)Z = -\ln\left(\frac{I}{I_0}\right) \quad (4.2)$$

This assumes a zero-width, monochromatic photon beam entering a homogeneous material which does not refract or diffract when coming into contact with any surfaces. In reality, the beam will have a non-zero width and common x-ray sources are significantly polychromatic. Higher order

effects are also visible from refraction and diffraction even with common x-ray sources. Another complication arises when one considers the case of imaging an inhomogeneous object, essentially any real object of interest. In this situation, the equation becomes:

$$I = \int_0^{E_{max}} I_0(E) \cdot e^{-\int_0^Z \mu(E,z) dz} dE \quad (4.3)$$

4

Not knowing the precise energy distribution or the atomic composition of every path within an object limits the usefulness of equation 4.3. If one assumes a single energy, which can be approximated with proper combinations of x-ray sources and filters, the equation becomes:

$$I = I_0 \cdot e^{-\int_0^Z \mu(z) dz} \quad (4.4)$$

Equation 4.4 is rearranged to the ‘Radon Transform’ and describes the line integral of the summed attenuation of all materials in the sample.

$$\ln \left(\frac{I_0}{I} \right) = - \int_0^Z \mu(z) dz \quad (4.5)$$

We now go through foundational computed tomography geometry, this explanation very closely follows a web-page on MathWorks (53). Figures have been remade based on the figures in this web-page.

A projection of a two-dimensional function $f(x, y)$ is a set of line integrals. An image on our detector is generated with a cone shaped beam where the object being imaged is rotated and the detector and x-ray source are static. Figure 4.1 shows a single projection at one of many rotation angles. Given that this is a 2D image, the cone beam is visualised as a fan. The line integral of $f(x, y)$ in the vertical direction is the projection of $f(x, y)$ onto the x-axis; the line integral in the horizontal direction is the projection of $f(x, y)$ onto the y-axis.




Figure 4.1: A fan-beam projection at a rotation angle, θ .

An example of horizontal and vertical projections for a simple 2D function is shown in figure 4.2.




Figure 4.2: Horizontal and vertical projections of a simple two dimensional function.

Projections can be measured along any angle theta, θ . In general, the Radon transform of $f(x, y)$ is the line integral of f parallel to the y' -axis.

We define this line integral $R_\theta(x')$ as:

$$R_\theta(x') = \int_{-\infty}^{\infty} f(x' \cos \theta - y' \sin \theta, x' \sin \theta + y' \cos \theta) dy' \quad (4.6)$$

where:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (4.7)$$

Figure 4.3 illustrates the geometry of the Radon transform.




Figure 4.3: Geometry of the Radon transform where $R_\theta(x')$ is the Radon transform of the axis x' .

The primary method to invert the Radon transform is the filtered back-projection, which is an analytical (not iterative) algorithm. The filtered back-projection can be used to reconstruct the density per voxel (3D pixel), however the exact solution requires an infinite number of rotations and detector pixels. In reality, tomographic x-ray measurements yield a finite number of rotations with a finite number of pixels. Solutions to this inversion problem include direct Fourier methods, filtered back-projection and iterative methods (algebraic and statistical). A thorough introduction to computed tomography can be found in print and freely online in ‘Principles of Computerized Tomographic Imaging’ (54) and also Enrico Schioppa’s PhD thesis (55).

Spatial resolution is determined by the pixel size divided by the magnification factor. When using x-ray tubes, spatial resolution is ultimately limited by x-ray spot size. The magnification factor is always more than one for cone beam x-ray sources, such as the type used in this work. Figure 4.4 visually summarises this.

In the same way as beams of visible light, cone beam (diverging) x-ray beams magnify objects the closer they are to the source. The magnification of the object (M) is the ‘source to detector distance’

(SDD) divided by the ‘source to object distance’ (SOD).

$$M = \frac{SDD}{SOD} \quad (4.8)$$

The projected pixel size (v) can be calculated by dividing the detector pixel pitch (d) by the magnification (M).

$$v = \frac{d}{M} \quad (4.9)$$

Since the objects measured always have some thickness, the source to object distance varies. Magnification therefore is different within the object depending on how far away it is from the source. The points of minimum and maximum magnification are the object point closest to the source and the point closest to the detector as the object rotates. These are marked as points A and B in figure 4.4.

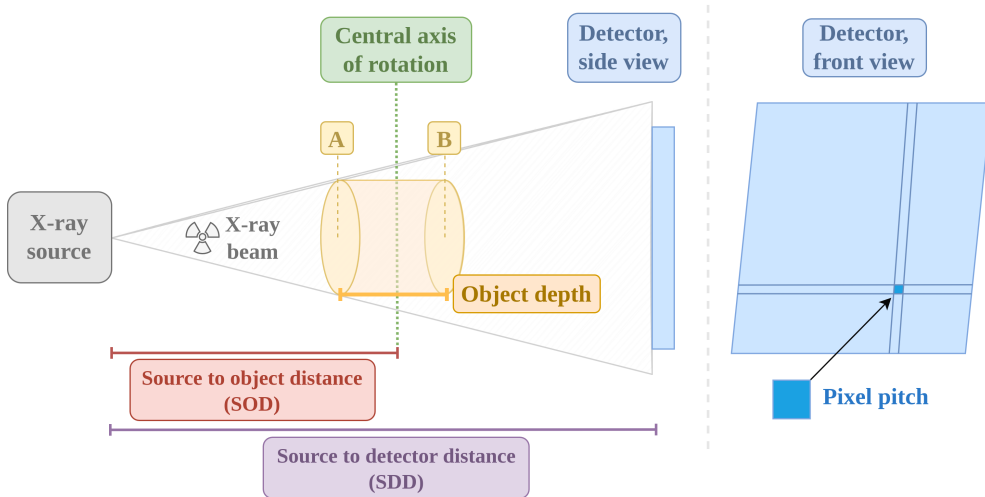


Figure 4.4: An overview of a cone beam computed tomography (CBCT) scan focused on magnification of the object being imaged.

The following calculations of the central and extreme magnifications and voxel size use realistic numbers based on the geometry of a real measurement in the Nikhef X-ray cabinet:

- Detector pixel pitch = 110 μm ;
- SOD = 100 mm;
- SDD = 160 mm;
- Object depth = 20 mm.

$$M_{\text{centre}} = \frac{160\text{mm}}{100\text{mm}} = 1.60 \quad (4.10)$$

$$M_A = \frac{160\text{mm}}{100\text{mm} - (20\text{mm}/2)} = 1.78 \quad (4.11)$$

$$M_B = \frac{160\text{mm}}{100\text{mm} + (20\text{mm}/2)} = 1.45 \quad (4.12)$$

$$v = \frac{110\mu\text{m}}{1.60} = 68.8 \mu\text{m} \quad (4.13)$$

The magnification using these example values is 1.60 at the centre, the minimum is 1.45 and the maximum is 1.78.

It should be noted that all CT reconstructions in this work were enabled by the ‘ASTRA’ toolbox using 3D cone beam reconstruction methods and GPU acceleration via the Python CUDA bindings (56; 57). This cone beam reconstruction is there to take the ‘non-uniform’ object magnification into account.

To match the small size of the detector a small custom phantom was produced based on my design at Nikhef. Figure 4.5 shows the engineering schematic of this phantom. A custom design was necessary because existing commercial CT phantoms are overwhelmingly meant for larger systems. It is a cylinder with a lid, both made of PMMA-c (polymethyl methacrylate, also known as acrylic, Plexiglas, Perspex etc.). It has seven 4.00 mm diameter holes along the length of the cylinder which contain other materials. These materials were other plastics and were chosen due to their relatively similar x-ray attenuation coefficients and densities. Water and air were also included.

This phantom is used in this chapter to explain sinograms and one type of CT reconstruction artefact.




Figure 4.5: A schematic of the custom designed microCT phantom. Made using FreeCAD.

4. X-RAY COMPUTED TOMOGRAPHY

4.2 EXPERIMENTAL SETUP

The x-ray cabinet was a fire safe that, by means of interlocks and lead cladding, was made fit for safely sealing the X-ray imaging set-up.

Two images of the cabinet can be found in figure 4.6. The upper half shows the cabinet when closed and actively imaging with x-rays. The lower half is when the cabinet is open with the detector in the 2D or CT imaging layout. The x-ray tube is the bright red box on the right, the detector is in the middle and is mounted to the fine y-z stage. The cabinet is lined with a few mm of lead.



Figure 4.6: Two images of the Nikhef x-ray cabinet, closed and open.

Now we know what the cabinet looks like, we switch to the experimental setup overview diagram in figure 4.7. It shows the system components and their two layouts for the same detector. The first

was for XRF measurements and the other was for 2D and CT imaging. The detector was mounted to the fine-y-z stage when doing 2D or CT imaging.

The internal dimensions of the cabinet are $140 \times 52 \times 37 \text{ cm}^3$. The sensor bias voltage cable, data and power cables are only shown for the 2D and CT imaging position; they are all present for all layouts.

The layout labelled 'XRF' shows how the detector is positioned for XRF measurements. The aim in this configuration is to measure the x-ray absorption edge energies (commonly known as 'k-edges') of a specific target material, not the primary beam itself. Several of these well known reference energies are used for detector energy calibrations. This is discussed further in the next chapter.

The detector is positioned as far outside the primary x-ray beam as possible in order to maximise the XRF signal relative to background events. Most background events in this configuration are expected to come from scattered photons from the primary x-ray beam, they could be scattered from the target material or anything else in the cabinet, including the lead lined walls.

An example of what this looks like when measuring copper is in figure 4.8. The x-ray spectrometer used was a SDD (silicon drift detector), a low count rate, high energy resolution detector. The make and model were: Amptek XR-100SDD.

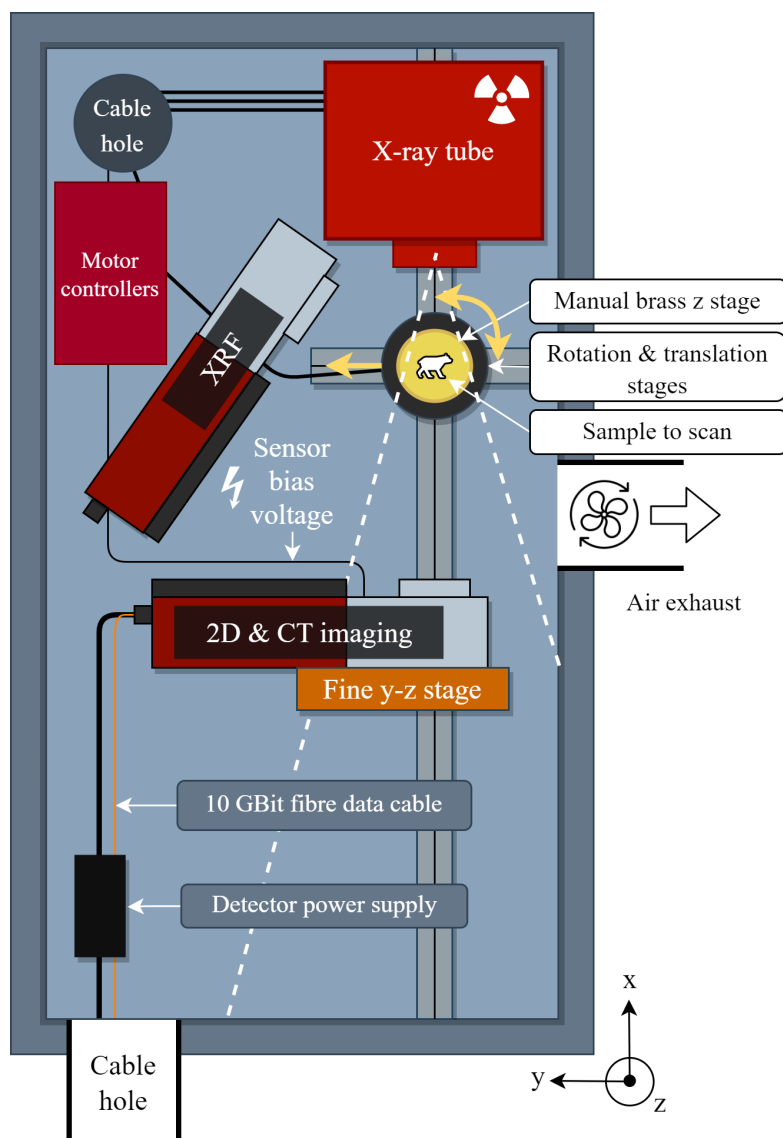


Figure 4.7: The experimental setup overview with two layouts, one for XRF measurements (outside the main beam) and the other for 2D and CT imaging. The white dashed lines diverging from the x-ray tube represent the approximate shape of the x-ray radiation emission. The manually controlled brass z-stage is on the rotation stage which rotates in the x-y plane. The coordinate system is in the lower right.




Figure 4.8: An image of the experimental setup for a XRF measurement of copper. The x-ray tube is the bright red box on the right, the copper target material is on the left and the Medipix3 and a x-ray spectrometer are above the x-ray tube.

The most frequently used detector layout was the one for 2D and CT imaging. In this configuration, we are interested in the x-ray attenuation differences both spatially and as a function of energy - effectively this is ‘colour imaging’. An example of the 2D and CT imaging layout is shown in figure 4.9.




Figure 4.9: An image of the experimental setup for 2D and CT imaging of a blueberry, which is on the manually controlled brass stage. This is on the black rotation stage with white letters indicating the current angle. There is a filter on the detector.

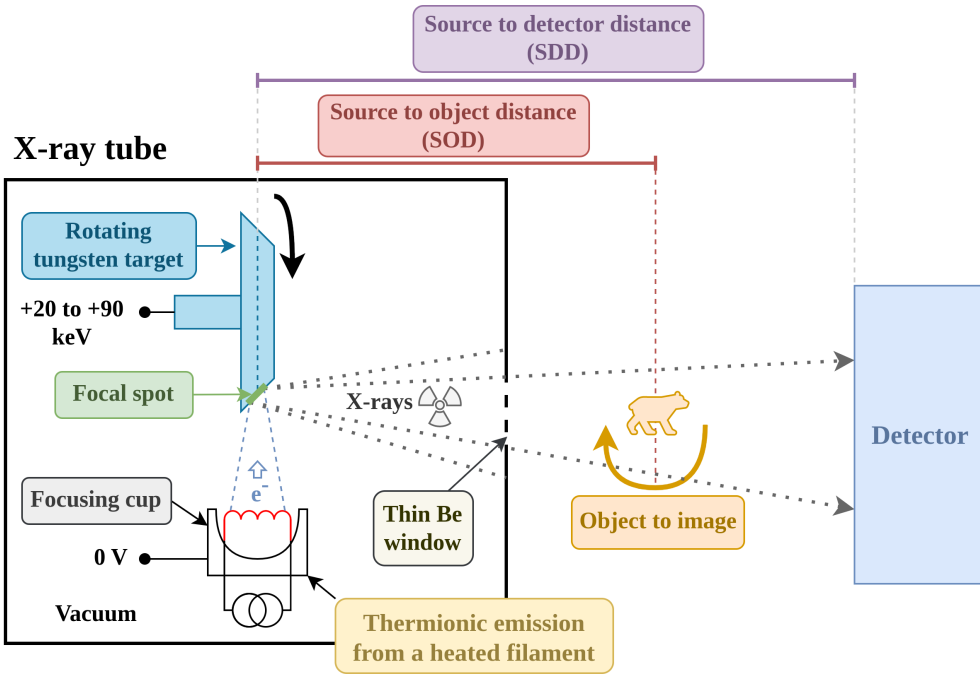


Figure 4.10: An overview diagram of a basic cone beam CT (CBCT) setup with a focus on the interior of a rotating target x-ray tube.

Figure 4.10 is a deeper dive into the x-ray tube and it shows the most basic form of a CBCT (cone beam computed tomography) measurement setup. A brief description of the relevant physics within a CBCT setup follows.

Electrons are emitted via thermionic emission from a heated filament. The electrons are accelerated towards the tungsten target and focused by the focusing cup into a small area, typical values are in the range of the μm to mm scale. The electrons decelerate as they interact with the tungsten target and emit x-rays via bremsstrahlung. The x-rays are emitted in a cone profile through a thin beryllium window and outside of the x-ray tube. The x-rays are then attenuated by an object and some pass through to finally reach the detector. The source to object (centre of rotation) and the object to detector distances are critical parameters for CT (computed tomography) reconstructions.

4. X-RAY COMPUTED TOMOGRAPHY

The fine 'y-z' stage has a range of 0–25 mm displacement in both the y and z axis respectively. This allows scanning of objects larger than otherwise possible without this control. The maximum effective scan area achievable in this setup is $53 \times 53 \text{ mm}^2 = 2809 \text{ mm}^2$ in a 2×2 grid layout, this is determined by both combination of the detector size ($28 \times 28 \text{ mm}^2$) and the translation stage (0–25 mm).

4

We can also magnify the images due to the cone shaped x-ray beam by changing the source to detector distance (SDD) and/or the source to object distance (SOD), see equation 4.8. This enables scans of smaller objects at higher magnification and therefore better spatial resolution, ultimately limited by the spot size of the x-ray tube.

A visual aid for this is figure 4.11. It shows the two-by-two grid layout and the overlapping areas which aid stitching. Ultimately, the x-ray tube would also limit the maximum scan size due to the diverging beam. The primary disadvantage of tiled scanning is that the measurement time now scales with the area. Secondly, drift effects would be introduced which complicates reconstruction.




Figure 4.11: A diagram showing the two by two grid layout of the quad chip Medipix₃ detector. This layout maximises the scan area specifically for the Nikhef cabinet to $53 \times 53 \text{ mm}^2$ using a detector with an active area of $28 \times 28 \text{ mm}^2$. The coloured squares represent the detector areas. The squares at the centre of the detector areas mark the central point of the detector. There is 3 mm of overlap in this configuration.

4.3 SINOGRAMS

In short, sinograms are a representation of CT data-sets where one axis is the rotation angle and the other is one column or row of the original images. They appear as multiple overlaid sinusoidally varying bands. The in-depth explanation of sinograms and the image acquisition and basic processing chain follows.

First, we start with the images, this is the ' I ' term in equation 4.1:

$$I = I_0 e^{-\mu(E) \cdot Z}$$

Keeping in mind the coordinate system introduced in figure 4.7, the projection images are the object as seen by the detector (in the y-z plane) as the object is rotated on a stage, around the z axis. We can then perform a flat-field correction by dividing the images (I) by the flat field images (I_0), taking the natural logarithm and multiplying by -1. Flat-field images are images taken under the same conditions as the projection images but without the object in view. At this point, we are at equation 4.2:

$$\mu(E) \cdot Z = -\ln\left(\frac{I}{I_0}\right)$$

We now have a three dimensional dataset, where the axes are y, z and object rotation angle - we measure from 0 to 360 degrees. There are typically many hundreds of angles measured. The x axis of the sinograms is the axis perpendicular to the rotational axis of the object, row. The y axis is the rotation angle of the object. This results in sinusoidally varying bands, they are symmetric along the object rotation angle at 180 degrees. Often this property is used to only take 0 to 180 degree scans and achieve a faster scan that, in theory, samples the volume space equally.

These sinograms are then fed into the CT reconstruction algorithms which output a three dimensional volume - the reconstructed slices in the x-y plane. The thickness of these slices is one voxel. The units here are x-ray linear attenuation coefficients (m^{-1}). These are subsequently visualised and analysed. There are many ways to use this data, it could be inspecting the slices directly, making volume renderings with thresholds for air, or analysing many CT reconstructions as a function of energy.

A diagram which summarises this process is figure 4.12. Note that there is a steel wire wrapped around the phantom in a helix shape which is not shown in the pictures. The dead pixels can be seen in the raw images - they are the small, randomly distributed black dots. They can be seen in the raw sinograms as the thin vertical black lines.

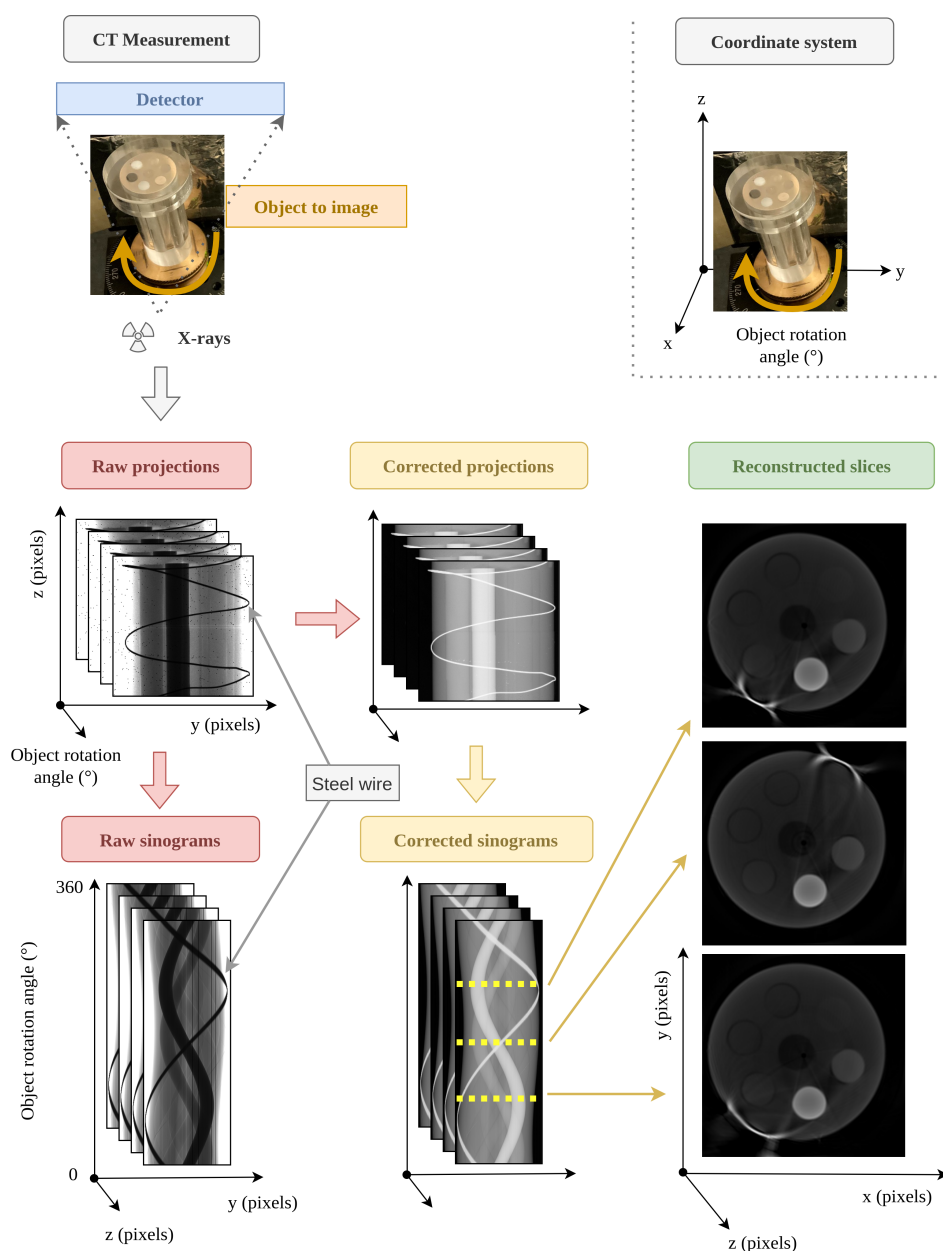


Figure 4.12: A diagram of the most simplified CT scan and reconstruction. The coordinate system used is in the upper right corner. X-rays are emitted from a source, they pass through an object and then are detected. These raw images are then corrected using a flat field image, the corrected sinograms are reconstructed using computed tomography algorithms. The output is a three dimensional volume of x-ray attenuation coefficients.

4.4 ARTEFACTS

We have many types of artefacts in our CT reconstructions, including cupping, beam hardening and ring artefacts. They are significant and therefore all need to be minimised before any work on spectral CT analysis makes any sense.

4.4.1 BEAM HARDENING

An x-ray beam from most x-ray sources is polychromatic, which means it has a wide energy distribution, anything up to the peak voltage used to accelerate the electrons. X-ray absorption in matter is both strongly atomic number (Z) and energy dependent (E). For a given homogeneous material and a polychromatic spectrum, one would therefore expect the x-ray beam to have a different energy distribution after passing through than before. Due to how the attenuation tends to reduce as a function of energy, this results in the energy distribution being dampened in the lower energy region. An alternative description is that the energy distribution has relatively more high energy than low energy x-rays after transiting a given material.

This phenomenon is referred to as ‘beam hardening’. Beam hardening in CT reconstructions introduces streaking and/or cupping artefacts in the reconstructed slices (58).

4.4.2 STREAKING ARTEFACTS

Streaking artefacts are found in areas with relatively high local attenuation variations, such as from bone to muscle or metal to bone and so on. They affect both the object itself and outside of the object, causing the artefact. They are the result of both beam hardening and Compton scattered x-rays. In the medical context, MAR (metal artefact reduction) algorithms are often used. Since streaking artefacts are caused by beam hardening, having energy sensitivity provides physical information to address the problem. An approach called ‘dual energy CT’ using existing detectors and two peak acceleration voltages is used to generate ‘virtual monochromatic images’. This partially fixes the issue.

Streaking artefacts are visible in figure 4.12 in the reconstructed slices section (lower right). For the sake of clarity, one of these reconstructions is displayed in figure 4.13. This figure indicates three regions of interest (ROI). The area between the two circles is the same material of the same density as the rest of the phantom bulk. However, the reconstructed attenuation coefficients in the green ROI is lower than the average value of the PMMA, one of several streaking artefacts here. The second artefact is between the brightest circle in the ROI and the steel wire in the yellow ROI. Yet another example can be seen around the wire in the red ROI, the wire is wrapped around the phantom and has a circular profile. The reconstruction shows excessively high and low intensity areas around the wire for all slices, curving both into the phantom and into the air.




Figure 4.13: An example reconstruction of the calibration phantom with three examples of streaking artefacts with regions of interests shown in red, yellow and green dotted rectangles.

4.4.3 CUPPING ARTEFACTS

As a polychromatic x-ray beam passes through a homogeneous cylinder, the centre of the cylinder absorbs more photons relative to the edges. Given the energy dependent absorption, the resulting x-ray beam through the thin parts at the outside cause a disproportionately large absorption of low energies. This results in a reconstructed profile resembling a cup, hence why these are referred to as cupping artefacts. The values reduce approaching the centre and increase approaching the edges. Software approaches exist to mitigate this effect, they are not perfect. There are many potential physical approaches to this including pre-hardening the beam with filters and detector based methods. An example of a cupping artefact is shown in figure 4.14b with the corresponding line profile in 4.14c and volume rendering in figure 4.14a.

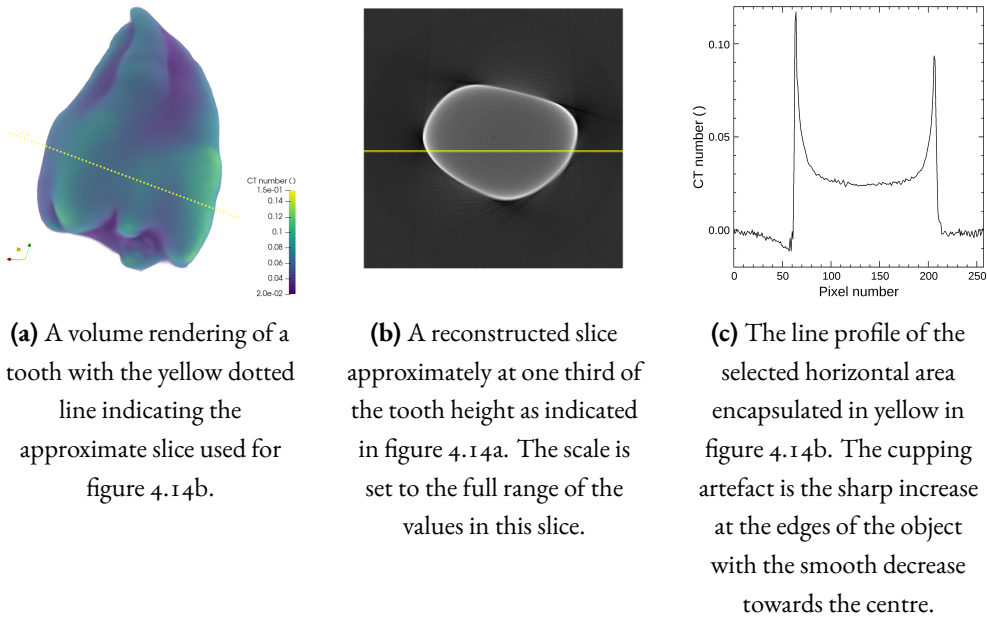


Figure 4.14: An example of cupping artefacts using a human wisdom tooth. All subplots show the same CT reconstruction data (CT number).

Beam hardening artefacts can be reduced by many methods. For example, by selecting a narrow energy window by using filters, detector thresholds or using a monochromatic source, the strong energy dependence is no longer a factor. Secondly, scanning objects with lower variations of atomic number will result in lower beam hardening artefacts; however, even perfectly pure and uniform materials will show cupping artefacts. And finally, using sensors that are fully absorbing at the energy range of interest will also help. For higher energies this is achieved with higher Z sensors such as GaAs (Gallium Arsenide), CdTe (Cadmium Telluride) & CZT (Cadmium Zinc Telluride).

4.4.4 RING ARTEFACTS

Ring artefacts look like tree growth rings in reconstructed slices and occur primarily due to detector element mismatches and beam hardening from sub-optimal flat field corrections. Variations in radiation source intensity and peak voltage can also introduce ring artefacts but this is typically negligible with current technology. They can be full or partial rings. As the object rotates, a given pixel will see different parts of the projected object. If that pixel has a different gain or noise baseline than its neighbours, it will appear different (brighter or darker) in the sinograms. Darker or brighter columns are then observed in the sinograms. These are propagated as rings in the CT reconstructions - these rings are systematic errors.

Detector elements have both systematic errors and random variations. Systematic errors can be

caused by different gains, noise offsets per detector element, or even variations in sensor thickness (59). This can be seen in figure 4.15, column four, row D. Random variations can be due to electronics or sensor effects which can depend on radiation, time, temperature and so on. For example, CdTe polarises with exposure to x-rays, which changes the response as a function of dose (60).

Beyond flat field corrections, the best methods to eliminate ring artefacts is to either manufacture identical detector elements or to be able to calibrate each detector element perfectly and for the calibration to be completely accurate over time. There are practical issues with all three of these points; this is discussed in the next chapter. The physical approaches, therefore, do not offer a practical and complete solution. Software methods can complement these in sinogram space or on the reconstructed slices.

A custom column-based filtering approach was developed in sinogram space and has been applied to all reconstructions shown in this work. The idea behind the algorithm is that the average value per column should be smoothly varying. It is intended to reduce the impact of pixel-to-pixel gain and noise baseline variations. The Python code for this with comments is included in Appendix E. A short summary of the algorithm follows and a diagram illustrating the filtering can be found in figure 4.16. There are two inputs to the function, one is a three-dimensional array of the rotation data (raw or corrected) and the other is a threshold for above the air but below the object (`above_air_value_but_below_interesting_data`). The data axes are rows, object rotation angle and columns. For example it could be shaped as (256, 600, 256) where there are 600 rotation angles and the detector has 256 rows and columns. For each row, the mean along the rotation angle is calculated (`column_array`). There is a nested loop where the window size is iterated over up to a fixed maximum. Within this, there is another iteration for the specific column. If the `column_array` is greater than the `above_air_value_but_below_interesting_data`, this column is set to the mean of the left and right windows. The size of the left and right windows is defined by the `window_size`. The input array is not modified, copies of arrays are made to ensure the intended behaviour. The sinogram is corrected by multiplying the ratio of the new modified array divided by the original unmodified array. In terms of the code length, this is a very simple algorithm at only 15 lines of Python code, excluding whitespace and comments. Execution time scales with the square of the `window_size` - in Big O notation, this is $O(n^2)$. It could also be easily parallelised to run on many CPU cores. For the data used in this work, a `window_size` of five was found to be approximately the point of diminishing returns.

The overall approach used in this work includes the column-based filtering, energy calibrations per pixel, corrections based on those calibrations (Energy Calibrated Images) all on top of standard flat field corrections. This is visually summarised in figures 4.5 and 4.15.




Figure 4.15: Ring artefact reductions with raw uncorrected data, flat field corrections only and energy calibrated images reconstructed at 6 keV with flat field corrections with sinograms and reconstructed slices. Software-based methods to correct for ring artefacts are common between columns two to four, this plot is exclusively showing the impact of ECI. Column four shows the difference between columns two and three, the effect of the energy calibration per pixel and subsequent reconstruction at 6 keV. Row C is row B cropped to a given area and the visualisation range is set to fit the data within the yellow boxes for each image. It shows how there are successively lower ring artefacts (vertical straight lines in sinograms) going from columns one to three. The object scanned here is the CT HU calibration phantom with a thin steel wire wrapped around in a helix. Column four shows how column three is better than column two. Column one is an x-ray image and its sinogram, column two, and column three are corrected images with their respective sinograms and reconstructions. The horizontal dashed yellow line indicates corresponding sinogram slice for the projections in row A.

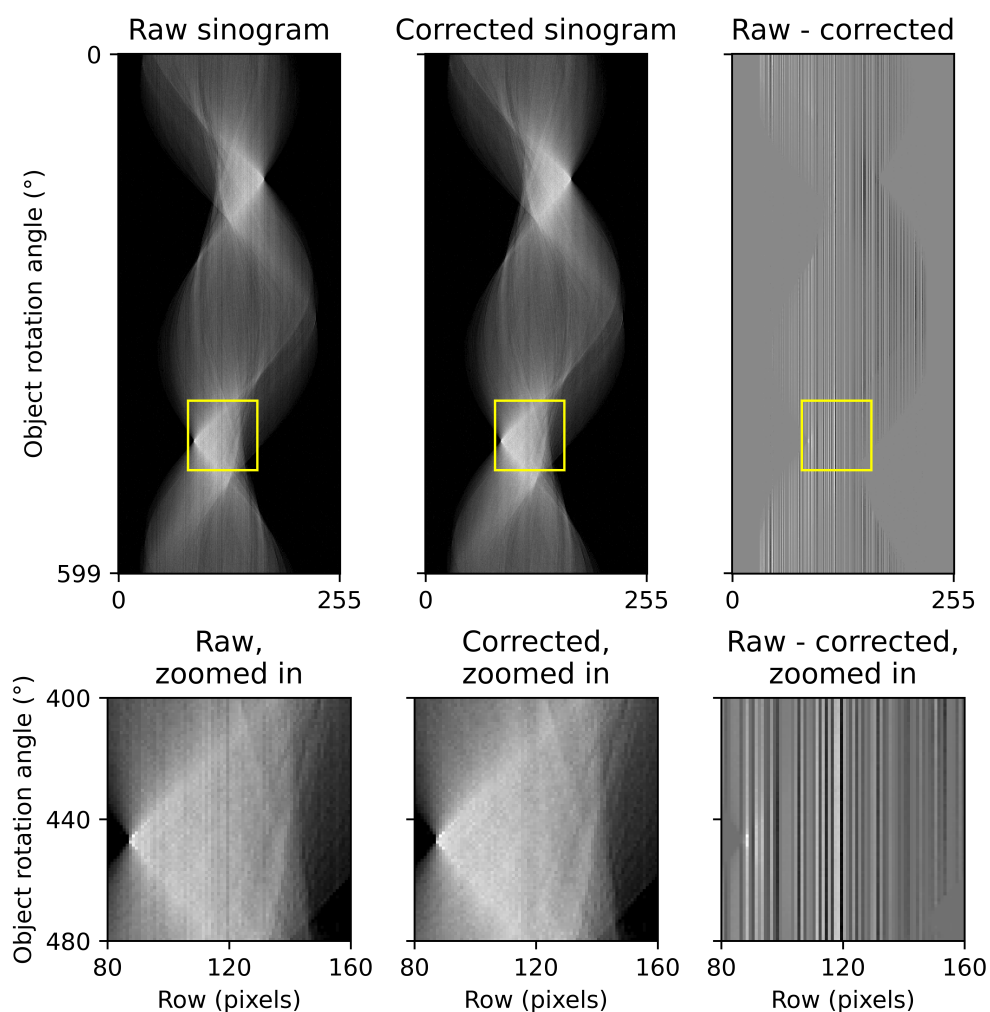


Figure 4.16: A diagram showing an example sinogram before and after using the column based filtering method. The uncorrected raw sinogram is on the left, followed by the corrected sinogram, and finally the difference between the raw and corrected sinograms is shown. The second row shows the zoomed-in region of interest as indicated by the yellow squares in the first row. It can be seen that the relatively prominent vertical lines in the raw sinograms are reduced in the second column. Therefore, this reduces ring artefacts.

4.4.5 OTHER ARTEFACTS

As discussed in section 2.6.4, random telegraphic noise is a process which results in random variations in pixel response over a wide frequency range. RTN affects both single pixels and chips, and so it propagates to CT scans in multiple ways. Randomly varying pixels cause partial or full ring artefacts. Randomly varying chips cause artefacts (horizontal or vertical lines) in sinograms over the rotation angle axis; this also propagates to the CT reconstructions.

Instabilities in the radiation source can also cause artefacts, either the energy distribution or the intensity can vary. Variations in the energy distribution can result in differences in beam hardening; this has not been observed in this work.

Variations in the x-ray tube current would result in corresponding detector-wide variations in counts. The result of this would be similar to RTN at chip level except that it would be relatively easy to correct for by multiplication and/or addition to smooth out the variations. Variations of the x-ray tube intensity could be measured with an additional detector as a cross check to the Medipix3 since the Medipix3 chips are expected to be less stable than the x-ray tube.

Work expands so as to fill the time available for its completion.

Cyril Northcote Parkinson, *The Economist*, 1955

5

5

Spectral X-ray Imaging

Hybrid pixel detectors can measure the x-ray spectrum and also enable very high speed measurements by high frame rates in the order of kilohertz. The high frame rate enables x-ray videos of moving organs like the heart with truly negligible motion artefacts.

Conventionally, the energy information of the x-ray spectrum is entangled by integrating over all energies. This is a detector limitation. When designing a detector, there must be a compromise between spatial resolution and energy resolution. You end up with imaging detectors (fast and no/low energy resolution), spectroscopic detectors (good energy resolution but slow) or something between. This gap is filled with the advent of hybrid pixel detector chips; these chips have comparable spatial resolution to imaging detectors while having energy resolving power. This is interesting for certain applications, one such application is discussed in the subsequent chapter.

X-ray attenuation for all materials is strongly dependent on energy, especially in the 1–100 keV range. This is the core of the reason why having any energy sensitivity is beneficial. It can be used to reduce artefacts in CT scans and allow the selection of specific contrast agents with transition energies in relevant energy ranges. These include iodine, gadolinium and gold nanoparticles. Since a new dimension, energy, is measured, it could even enable entirely new imaging techniques.

5.1 LINEAR ATTENUATION COEFFICIENT OVER ENERGY

The mass attenuation coefficient, $\frac{\mu(E)}{\rho}$, is heavily dependent on energy and is discussed in section 5.2.

Within the mass attenuation coefficient, the linear attenuation coefficient ($\mu(E)$) is the term that is dependent on energy. It is possible to reconstruct the energy dependency of the linear attenuation coefficient. Medipix3 specifically allows this because it is able to set 1–8 energy thresholds

depending on the configuration. Flat panel detectors do not share this feature and so the peak x-ray voltage is varied to accomplish a similar goal; this is used in dual energy CT. As a bonus, the Medipix3 is able to record at multiple energy thresholds simultaneously. This is important because the quantum noise is precisely the same for different thresholds; thus eliminating that error source when subtracting energy thresholds from each other. This is further discussed in subsection 5.5.4.

In order to uniquely determine elements and compounds, we can use the linear x-ray attenuation coefficient as functions of energy. We can build up the relative linear attenuation coefficient as a function of energy by measuring at many energies. Breaking that down further, we need to derive the ratio of the linear attenuation coefficients at two energies, $\frac{\mu_{E_2}}{\mu_{E_1}}$. To do that, we start from equation 4.1.

$$I = I_0 \cdot e^{-\mu(E) \cdot Z} \quad (5.1)$$

$$Z = \frac{-1}{\mu(E)} \cdot \ln \left(\frac{I}{I_0} \right) \quad (5.2)$$

$$Z_{E_1} \equiv Z_{E_2} \quad (5.3)$$

$$\therefore \frac{-1}{\mu_{E_1}} \cdot \ln \left(\frac{I_{E_1}}{I_{E_{0,1}}} \right) = \frac{-1}{\mu_{E_2}} \cdot \ln \left(\frac{I_{E_2}}{I_{E_{0,2}}} \right) \quad (5.4)$$

$$\frac{\mu_{E_2}}{\mu_{E_1}} = \frac{\ln \left(\frac{I_{E_2}}{I_{E_{0,2}}} \right)}{\ln \left(\frac{I_{E_1}}{I_{E_{0,1}}} \right)} \quad (5.5)$$

To precisely calculate the ratios of linear attenuation coefficients at different energies, ideally the energy resolution and threshold dispersion would be zero. In reality, the energy resolution FWHM and the threshold dispersion are in the order of keV. For this application with soft tissues and a silicon sensor, this is a significant issue. The linear attenuation coefficient varies the most at low energies (< 10 keV) and silicon becomes almost transparent above 25 keV. On top of this, the minimum energy threshold that can be reliably used with this configuration is 5 keV.

The threshold dispersion is minimised at the noise baseline, so at the minimum usable level (5 keV), the threshold dispersion is minimal. As the energy increases the threshold dispersion increases proportionally so pixels respond at increasingly different energies. This shows itself as progressively noisier images as a function of energy. Energy calibration per pixel offers a first order correction for this effect; this is discussed in section 5.3.

5.2 SOFT TISSUE X-RAY CONTRAST

The x-ray mass attenuation coefficients as a function of energy for various tissue types have been measured and tabulated in NIST's x-ray mass attenuation coefficients table three

<http://physics.nist.gov/PhysRefData/XrayMassCoef/tab3.html>. Mass attenuation coefficients are strongly atomic number (Z) and energy dependent since the photoelectric cross section σ_{ph} is pro-

portional to $\frac{Z^4}{E^3}$. However, almost all elements found in the human body by mass are low Z such as oxygen, carbon, hydrogen, nitrogen and calcium. Therefore, apart from bones, the differences in attenuation are mostly due to density variations. These density variations result from biological aspects, for example blood in one part of the body will have a slightly different density to elsewhere because the number of red blood cells is different.

In most biological materials, the variations in density and chemical composition are larger than the variation in x-ray mass attenuation coefficients over the energy range we are capable of measuring over. It is therefore not sufficient to use these x-ray mass attenuation tables for determining soft tissue contrast as a function of energy.

This subsection explores how x-ray attenuation differences affect imaging contrast. The differences are demonstrated between soft tissue and hard tissue such as bone and within soft tissues. We start with a fundamental overview of photon interactions with matter in the x-ray energy region. We then cover the absolute attenuation of various materials of interest over the relevant energy range and compare these to each other.

We start with the absolute x-ray attenuations for various biological materials in figure 5.1. Note that the mass attenuation coefficient is used as materials do vary in terms of density and so a density-independent metric is useful for material comparisons.




Figure 5.1: The mass attenuation coefficients, μ/ρ , for six materials of interest: cortical bone, liquid water, dry air near sea level, skeletal muscle tissue, adipose tissue and polyethylene. The materials in the legend are ordered in terms of descending average mass attenuation coefficient. The black dashed rectangle indicates the region of interest used in the subsequent plot. The data for these coefficients is compiled, updated and distributed by NIST (National Institute of Standard and Technology), USA (61) under the X-Ray Mass Attenuation Coefficients section: <https://physics.nist.gov/PhysRefData/XrayMassCoef/tab4.html>.

As one might expect, bone has the highest average mass attenuation coefficient over most of this energy range. Polyethylene and adipose (fat) tissue have the lowest coefficients for most of the energy range plotted. The ‘ICRU-44’ postfix as seen in the legend of figure 5.1 is in reference to the International Commission on Radiation Physics Report 44: ‘Tissue Substitutes in Radiation Dosimetry and Measurement’ published in January 1989; it standardised the definition of certain composite materials as specific ratios of elements. The rest of the materials (air, muscle & water) are clustered together in this log-log scale plot, showing that there are minimal differences between them over the energy range of interest. In addition to this, the K-shell photoelectric effect transition energies of carbon, nitrogen and oxygen are 277, 392.4 and 524.9 eV respectively, which is far below the noise floor of the Medipix3 under ideal conditions (~ 4 keV with a Si sensor). There are therefore no transition energies of the main components of soft tissue in the achievable energy range. Even ignoring the problem of the noise floor, x-rays of such low energies strongly interact

with air and so measurements in vacuum would become necessary and only very thin objects could be imaged.

Zooming to a more relevant region in figure 5.2, we see differences starting to emerge. Note that the plot is now zoomed into a smaller range in both axes and is now a log-linear plot rather than log-log. The lines are not smooth due to the finite sampling of these coefficients from NIST. At this point, one may think that it is trivial to separate these materials. This is not the case because there are significant density variations for these materials, especially muscle and adipose tissues.




Figure 5.2: The photon mass attenuation coefficients for various composite materials from zero to 100 keV. The data used in these plots are from NIST’s x-ray mass attenuation coefficient table 4. The vertical dashed grey line at 5.0 keV indicates the approximate minimum achievable energy threshold with Medipix3 with a silicon sensor.

One method to help with this is to normalise a given measurement with the equivalent air measurement (without the object of interest); this is known as a flat-field correction. We now multiply by density to arrive at the attenuation coefficient, $\frac{\mu}{\rho} \times \rho = \mu$. We now add in $\pm 3\sigma$ error bands using data compiled by The Foundation for Research on Information Technologies in Society (IT’IS), Switzerland (62).

Figure 5.3 uses information from these data; it shows how the selected tissues now overlap, given the variation within the population. Bone is the only tissue in this list that can be unambiguously separated from the others over this energy range. In the context of this work, we are interested in soft tissues such as breast, muscle and fat. These have the biggest separation at low energies and have

significant overlap by 40 keV. Therefore, even with a perfect measurement of these coefficients, it is not possible to unambiguously distinguish soft tissues from each other.




Figure 5.3: The photon attenuation coefficients for several biological materials from 5 to 90 keV. The data used in these plots are from NIST’s x-ray mass attenuation coefficient table four. The error bands used are the 3σ values from the tissue density database compiled by the IT’IS Foundation (The Foundation for Research on Information Technologies in Society) (62). The vertical grey dashed line at 5.0 keV indicates the approximate minimum achievable energy threshold with Medipix3 with a silicon sensor.

5.3 ENERGY CALIBRATION

OVERVIEW

This subsection exclusively discusses the energy calibration approach. The application of this for imaging are discussed later in section 5.5.

Energy calibration of the detector can be achieved with XRF (x-ray fluoroscopy) where one uses multiple pure and thin metal foils with electronic transitions to obtain single energy references. Typically the transitions used are called K_α and K_β in the Siegbahn notation and respectively also called $K - L_3$ and $K - L_2$ in the modern IUPAC (International Union of Pure and Applied Chemistry) notation.

Another option for single energy references is electronic test pulses which are built into every

pixel. Test pulses need to be calibrated themselves as well and would only test the front end pixel electronics. With XRF you calibrate the entire system, pixel and sensor, at once. Therefore test pulses are not used.

Threshold scans are defined as follows. With the radiation source active, one acquires an image, changes a threshold(s) by some step size and then repeats until finished. They are frequently used for chip and radiation characterisation and some types of measurements. Subsequent analysis is application specific. Since this scan type is common, it was implemented in software where one enters parameters like step size and frames per threshold and where the files will be saved.

Both XRF and test pulse calibration methods scan over the threshold to probe the relationship between threshold and counts. Threshold DACs are linearly proportional to energy within the energy range of interest. The threshold scan produces distributions of counts over the threshold (in DAC units) which are ideally error functions ('erf'), often referred to as 'S-curves' in the Medipix group due to their shape. The differential represents the number of counts per energy bin (DAC step). Ideally you would get a normal distribution. However, due to charge sharing and weaker emission lines at slightly different energies the distribution will not be perfectly normally distributed, figures 5.5 and 5.6 show this. The normal distribution corresponding to the XRF line of interest can be disentangled from the background. The mean of the normal distribution maps to the DAC value at which the energy reference is. The standard deviation of the normal distribution gives the energy resolution, which can be expressed in both DAC units and in energy. With multiple known energies, this process can be repeated for sufficient data points; three would be the absolute minimum. This analysis is used similarly for chips and individual pixels and is graphically summarised in figure 2.7.

Many factors complicate this story. Firstly, charge sharing between pixels results in a skewed distribution towards lower energies. Secondly, the experimental setup is not unlikely to result in x-rays at other energies than the reference level; this can happen for many reasons. Finally among other reasons, when using sensor materials with K edges in the energy range one is attempting to calibrate, such as GaAs, CdTe and CZT, additional error sources are introduced.

Even with monochromatic synchrotron radiation, such measurements do not result in a perfect error function, there is always a low energy tail using both SPM and CSM as shown in figure 5.4.

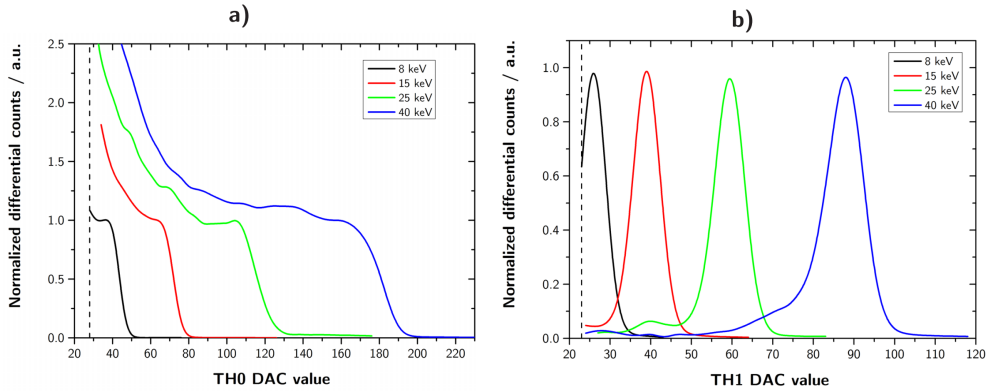


Figure 9.13: Normalized energy spectra in SPM (a) and in CSM (b) for monochromatic synchrotron radiation of 8 keV, 15 keV, 25 keV and 40 keV. The electronic noise level is indicated by the vertical dashed black line.

Figure 5.4: Figure 9.13 with the original caption (27).

Supposing that the aforementioned challenges are overcome, one might assume to be in possession of a perfect energy calibration for every chip and for every pixel within that chip. While that should be true at that time, any energy calibration will need to be refreshed after some experiment specific time. As detailed in chapter 2, the system response does vary over time beyond photon counting statistics. The most extreme practical option would be to calibrate before and after individual measurements.

If thresholds are set too close to each other in energy, the effect of the energy resolution and inter-pixel variation will produce increasingly non-physical results when subtracting one from the other. There is therefore a minimum inter-threshold energy spacing which should be maintained. This spacing should be more than the energy resolution of the chip.

A practical example follows. When using a polychromatic x-ray source (like an x-ray tube), if th_0 is set to 10 keV and th_1 is set to 11 keV, one might expect every pixel in th_0 to count more than th_1 . In reality, this does not occur, sometimes the same pixel in th_0 will count less than th_1 even though th_0 is set at a lower threshold.

Each physical pixel contains a set of transistors and each set of transistors corresponds to a threshold. Due to engineering tolerances in production, there are noise and gain variations in these sets of transistors. Therefore, different thresholds have different energy calibrations. This results in needing between one and eight thresholds per chip. For a multi-chip detector, the number of energy calibrations is then multiplied by the number of chips, so for the four chip systems used in this work, between four and 32 energy calibrations are required per detector just for the thresholds themselves. The analysis of calibrating each pixel is presented in the subsequent subsection.

ANALYSIS

The following subsection outlines the full chain analysis from the observed counts over the threshold to various high level statistical metrics. The examples shown in this analysis utilise a 500 μm thick silicon sensor biased to +150 V with a pixel pitch of 110 μm and with SPM. The analysis is independent of all of these properties. The main difference with CSM is that figure 5.5 would be closer to a horizontally flipped error function, which propagates in the analysis accordingly.

5

Starting with a threshold scan of Sn in the XRF experimental setup with SPM, the data for figure 5.5 is directly measured. This figure shows one pixel with x-y coordinates of (100,100) arbitrarily chosen. Sn has three fluorescence energies close to each other, $K - \alpha_1$ at 25.2714 keV, $K - \alpha_2$ at 25.0440 keV and $K - \beta_1$ at 28.4860 keV. The $K - \alpha$ energies are not separable with a Medipix3 because they are only 227.4 eV apart, which is less than the energy resolution per pixel. Therefore the two features (changes in slopes) at ~ 118 and ~ 132 DAC units correspond to the mixed $K - \alpha$ energies and the $K - \beta_1$ energy. The increase in the number of counts to the left of these features is expected and due to charge sharing primarily.




Figure 5.5: Counts as a function of threshold for a single pixel ($x = 100, y = 100$) for an example measurement using Sn x-ray fluorescence. There are mixed $K - \alpha$ energies and the $K - \beta_1$ energy respectively at ~ 118 and ~ 132 DAC units.




Figure 5.6: The numerically differentiated counts as a function of threshold for a single pixel ($x = 100, y = 100$) for an example measurement using Sn x-ray fluorescence. After the gradient is removed from the data, a Gaussian distribution is fitted. Only the main peak is used for the fit due to both the noise and the presence of a second peak at a slightly higher threshold (energy).

Figure 5.6 moves on to show the result of numerically differentiating the data from figure 5.5, see the blue circles labelled ‘(1) Differentiated raw data’. Fitting a Gaussian distribution to this data leads to poor quality fits due to both the charge sharing at lower thresholds than the main peak and the higher relative noise at lower thresholds. By subtracting the gradient in the differentiated raw data using a straight line fit ‘(2) Straight line fit of (1)’, the fit quality and reliability increases dramatically because the low energy tail is removed and only a subset of the data around the peak is used for the fit. This method was chosen after developing a more physically accurate model which produced very similar results and took a few orders of magnitude longer to compute. Specifically, three points immediately before the peak and ten points after are used. This approach could be optimised many ways, for example a composite function could be fitted to the data instead. This could take the form of an error function with a Gaussian distribution. The fit to the resulting distribution ‘(3) (1) minus (2)’ is labelled ‘(4) Gaussian fit of (3)...’ which has corresponding mean (μ) and sigma (σ) values. The maximum point of the distribution is labelled as ‘(5) Maximum point = ...’. In this analysis, only the strongest peak is considered. This is potentially another point for improvement.

Figure 5.7 has two subplots which look at the same information as the previous two plots but at chip level instead of for a single pixel. The left subplot shows the undifferentiated counts for all four chips, similarly to figure 5.5. The right subplot shows the differentiated counts for all four chips with the subset of data points shown as the circles and the resulting Gaussian fit shown as the

solid lines. The fit parameters are shown in the legend per chip.




Figure 5.7: The counts and numerically differentiated counts as a function of threshold in DAC units for an example measurement using Sn fluorescence. Left: Counts per chip as a function of the threshold in DAC units. The main drop in counts at 100—115 DAC units is accompanied by a smaller drop at 115—130 DAC units. These correspond to the two K edges of Sn. Right: the numerically differentiated plots per chip with the data as circles and the fitted normal distribution as solid lines. The fit parameters are in the legend of the right plot. The raw data used for the fit is a subset of the input data to aid fitting.

Figures 5.8 and 5.9 represent information in similar ways. They show their respective values per pixel first on the left and then the histogram of that data on the right with a log scale and the various statistical distribution values in the legends.

Figure 5.8 shows how the four chips have different mean μ values. There is a darker than surrounding area in the top left chip, the cause of it is not known with certainty and was not observed on different metal foil measurements with this detector in the same conditions. The black pixels at the top of the top right chip (chip one) and bottom left chip (chip three) correspond to the peak at roughly 85 DAC units as seen in the histogram. As a reminder for the chip layout, see figure 2.3.

Figure 5.9 shows how the four chips have different mean σ values without any obvious spatially dependent features like gradients or large local variations. There are a number of pixels with a reported standard deviation of zero, these correspond to the unresponsive pixels, also known as ‘dead pixels’.




Figure 5.8: The mean ($\mu(\text{Th})$) fitted DAC values per pixel for an example measurement using Sn fluorescence. The mean of the mean distribution is $\mu(\mu(\text{Th})) = 111.77$ DAC units and the standard deviation of the mean distribution is $\sigma(\mu(\text{Th})) = 4.34$ DAC units.




Figure 5.9: The standard deviation ($\sigma(\text{Th})$) fitted DAC values per pixel for an example measurement using Sn fluorescence. The mean of the standard deviation distribution is $\mu(\sigma(\text{Th})) = 1.99$ DAC units and the standard deviation of the standard deviation distribution is $\sigma(\sigma(\text{Th})) = 0.18$ DAC units.

The correlation of the fitted mean and standard deviation distributions is shown in figure 5.10. One would expect a 2D Gaussian distribution per chip. With four chips, this expectation becomes four overlaid 2D Gaussian distributions when plotted together. This expectation appears to be consistent with this plot.




Figure 5.10: A 2D histogram of the mean distribution plotted against standard deviation distribution for fitted DAC values per pixel for an example measurement using Sn fluorescence. Note the log scale for the colour bar. The plot limits were set to the mean $\pm 4\sigma$ with the Y lower limit fixed at 0. The horizontal red line is the lower limit for the standard deviation of $\frac{1}{2\sqrt{2\ln 2}} = 0.425$ DAC units. The $2\sqrt{2\ln 2}$ term comes from the relation $FWHM = 2\sqrt{2\ln 2}\sigma$.

Repeating the above steps for multiple pure metal targets results in figure 5.11. The x coordinate is the mean of the mean distribution ($\mu(\mu(\text{Th}))$) and the y coordinate is the x-ray energy of the highest intensity emission line of the metal target used. In this example series of experiments, Cu, Mo, Cd & Sn were used with respective energies of 8.047, 17.479, 23.174 & 25.271 keV. The goodness-of-fit number, R^2 , is near perfect at 0.99997 which shows the excellent linearity in terms of energy for Medipix3 in this energy range and configuration.




Figure 5.11: The global per-pixel detector energy calibration using four metals (Cu, Mo, Cd & Sn) with a near perfect linear fit. ‘a’ is the slope and ‘b’ is the intercept value in units of keV/DAC and keV respectively.

Finally, the analysis is concluded with figure 5.12. It shows the 2D distributions and histograms of the three parameters related to the straight line fits for per-pixel energy calibration. The slopes can be considered to be the same as gain and have units of keV/DAC. See figure 2.3 for a reminder for the chip layout. The top left chip (chip zero) has a lower gain than the top right (chip one) and bottom right (chip two) chips while the bottom left chip (chip three) has a higher average slope. The intercepts are equivalent to noise baselines per pixel in keV. The noise baselines are relatively uniform across the four chips except for the near the edges for the top right (chip one) and bottom left (chip three) chips. These have lower mean values in figure 5.8 and lower slope values in the 2D slope distribution in figure 5.12. Almost all the pixels have excellent fits, as seen in the R^2 histogram - almost all the entries are in the last histogram bin. Fits with insufficient quality can be identified with this subplot by using an arbitrarily chosen R^2 value based on the desired fit quality and the number of pixels that can be ‘sacrificed’. Pixels with insufficient quality fits can then be masked or otherwise processed.

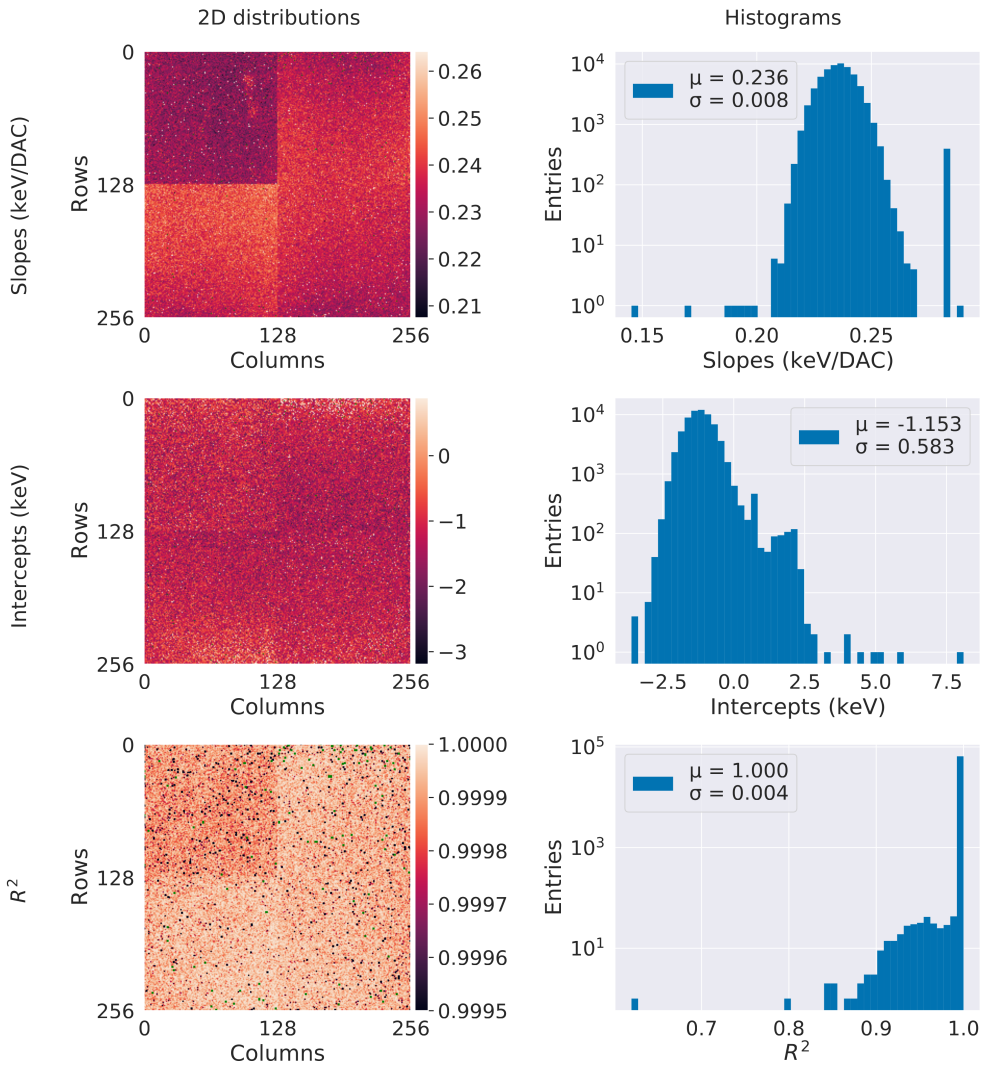


Figure 5.12: The 2D distributions and histograms of the slopes (gains) in keV/DAC, the intercepts in keV and the R^2 goodness of fit values for th1 using Cu, Mo, Cd & Sn in a series of example XRF experiments.

Practically speaking, it would be ideal to set the gain of each pixel. In Medipix3, this feature was not implemented and so we are left with only being able to set the noise baselines per pixel (equalisation) and chip wide thresholds (threshold DACs). A description and algorithm overview for the equalisation procedure can be found in appendix D - this shows one of several possible variations depending on the chip and sensor configuration. Setting chip wide thresholds effectively sets the mean of the distribution since all pixels behave slightly differently. It would also be ideal to

be able to apply an additive or multiplicative set of per-pixel corrections. In addition to the pixel noise baseline, it would be ideal to adjust the gain per pixel to further homogenise the detector response. To first order, this is addressed later on.

5.4 XRF EXPERIMENTS

5

All of the following experiments ideally require a mono-energetic x-ray source. This can be achieved by the following approaches. First, an x-ray tube in an x-ray fluoroscopy (XRF) configuration. An overview of this arrangement is in figure 4.10. Secondly, with a radioactive source, eg. Am-241 emits 59.54 keV photons. Finally, one ideally has access to a synchrotron.

The following three setups are now presented:

1. Silicon drift detector
2. Medipix3 with a Si sensor and pixel pitch of 110 μm
3. Medipix3 with a GaAs sensor and pixel pitch of 55 μm

5.4.1 SILICON DRIFT DETECTOR

We start with the high energy resolution spectrometer, the silicon drift detector. Note that the SDD has one detector element, so this is equivalent to having pixels in the order of $1 \times 1 \text{ mm}^2$ as opposed to $55 \times 55 \mu\text{m}^2$ pixel pitch. Figure 5.13 shows a XRF measurement using Cd as the target with the equivalent Medipix3 energy resolution, using a representative example FWHM value of 2.1 keV. The data is normalised to the peaks of each curve respectively.

It shows the x-ray emission lines of cadmium (Cd) at 23.1736 ($K - \alpha_1$), 22.9841 ($K - \alpha_2$) and 26.0955 ($K - \beta_1$) keV. There are five more L shell emission lines from 3.1–3.7 keV. One can observe many background peaks/signals that are not associated with Cd. This background is a result of every other interaction from the 90 kVp x-ray incident beam not related to the Cd XRF signal. The causes of this include lead shielding, steel, copper, tungsten, tin, air and any other element within the cabinet. Significant changes would be required in this cabinet to achieve better signal to background ratios; the current ratio is sufficient.




Figure 5.13: Measured Cd XRF spectrum using an Amptek XR-100SDD and a PX5 digital pulse processor with an energy bin size of 11.6 eV and the equivalent energy resolution available with Medipix3 (2.1 keV). In order of decreasing intensity, Cd fluorescence is 23.1736 ($K - \alpha_1$), 22.9841 ($K - \alpha_2$) and 26.0955 ($K - \beta_1$) keV as marked with the vertical dashed grey lines. Five L shell x-ray emission lines are included from 3.1 to 3.7 keV, only four are visible due to two lines overlapping on this scale.

5.4.2 Si 110 μm PIXEL PITCH

With a 500 μm Si sensor with a 110 μm pixel pitch (colour mode), it is possible to measure S -curves with XRF experiments with all modes - SPM, SPMa and CSM. Figure 5.14 presents differentiated S -curves in DAC units. We are using DAC units and not keV here because different gain modes are used. This means that for a given energy, the peak will be at a different DAC unit. Also, these peaks correspond to the same energies as each other because the same Cd XRF energy spectrum is used between each measurements.

As expected, it shows that the ratio of the mean values for SPM th1 compared to CSM for the same threshold is close to two. This also implies that the minimum noise baseline is also approximately double for CSM compared to SPM. The FWHM of the SPMa fit is between the SPM and CSM values, yet again, as expected. When comparing distributions, keep in mind that tho corresponds to physically different pixel electronics compared to th1. The same DAC values for tho and th1 do not correspond to the same analogue outputs.




Figure 5.14: Using a Cd XRF experiment and a Medipix3 with a 500 μm Si sensor and pixel pitch of 110 μm biased to +150 V using SPM, SPMa and CSM, all in HGM (High Gain Mode). The fitted distributions are the sum of a Gaussian distribution and a straight line. All pixels in all chips are used in these plots. The mean and FWHM of the fitted distributions are in the legend. The distributions were all normalised to their respective peaks.

In the context of this work with this configuration, using CSM is not useful because the noise baseline is relatively high and the transparency of silicon increases significantly at around 25 keV. This limits the useful energy range in this configuration to roughly 10–25 keV with an energy resolution (FWHM) of 5 keV. The ratio of the energy range to the energy resolution is therefore only approximately 3. For SPM, this ratio is approximately $\frac{25 \text{ keV} - 5 \text{ keV}}{2.5 \text{ keV}} = 8$. In this case, the benefit of a more correct spectrum with CSM is outweighed by this unfavourable ratio compared to SPM combined with the lower count rate capability of CSM. The same conclusion would not be drawn for sufficiently thick high Z sensors such as GaAs, CdTe and CZT where the energy range is far larger. More on this topic follows in the next subsection.

5.4.3 GaAs 55 μm PIXEL PITCH

With a 500 μm GaAs sensor with a 55 μm pixel pitch (Fine Pitch Mode), it was not possible to observe an S-curve with an XRF experiment with even a single pixel using SPM. However with CSM, this is not the case. Instead the three functioning gain modes were explored. Super low (SLGM), low (LGM) and high (HGM) gain modes were used with CSM as shown in figure 5.15. The scale of both axes matches figure 5.14 for easier direct visual comparison.

This GaAs sensor is an n-type semiconductor material because it is compensated with chromium (Cr), we therefore collect electrons and use a bias voltage on the sensor of -130 V. This is the opposite polarity to silicon because the silicon sensor material used in this work is p-type and we collect holes.




Figure 5.15: Using a Cd XRF experiment and a Medipix3 with a 500 μm GaAs sensor and pixel pitch of 55 μm biased to -130 V in super low (SLGM), low (LGM) and high (HGM) gain modes, all use CSM. The fitted distributions are Gaussian distributions. All pixels in all chips are used in these plots. The mean and FWHM of the fitted distributions are in the legend. The distributions were all normalised to their respective peaks.

As one would expect, the gain modes are proportional to the mean value of these distributions for a given energy. The FWHM also scales with the gain modes because while they are the same in terms of energy, they are not in terms of DAC units because of the increased resolution with higher gain modes. The relevant gain mode will be partly determined by the trade-off between energy resolution and energy range.

5.5 ENERGY CALIBRATED IMAGES (ECI)

5.5.1 INTRODUCTION

Hybrid pixel detectors suffer from differences in pixel gain, threshold and crystal defects in the sensor bulk, among others. One can compensate for these effects by using per-pixel energy calibrations and at least two images to reconstruct images effectively at a single energy. This method can be optimised in terms of acquisition speed and reconstruction quality. Ideally only non systematic error sources would remain, including electronic noise on the preamplifier input, unstable pixels and beam hardening in the sample. In reality, the energy calibration per pixel is not perfect and even if it were, the pixels do change response over time so the calibration effectively is a first order correction.

5.5.2 METHODS

When a chip wide threshold is set which corresponds to a specific energy, the effective threshold per pixel is distributed over energy with the centre at the specific chip wide energy that was set, see figure 5.8. The width of this distribution is the threshold dispersion, typically it is expressed in DAC units, number of electrons or energy. The method is increasingly beneficial as the threshold dispersion increases. For Medipix3, the threshold dispersion is proportional to energy. As mentioned in the previous subsection, it is possible to minimise this dispersion.

If every pixel is calibrated in terms of energy as described in section 5.3, the slopes [keV/DAC] and intercepts [keV] are known per pixel. This method assumes that the pixels have sufficiently stable slopes and intercepts over time. Radiation damage, random telegraphic noise (RTN), significant temperature variations and varied leakage currents would all degrade that stability.

Scanning over the threshold (energy) or simultaneously measuring multiple energies, effectively measures the counts as a function of the threshold for every pixel. Chip wide thresholds can then be used to map the chip wide threshold in DAC units to the actual energy in keV per pixel. This mapping uses the linear energy calibration per pixel, as shown in equation 5.6.

$$Energy_{x,y,thr} = slope_{x,y,thr} \times threshold_DAC_{chip,thr} + intercept_{x,y,thr} \quad (5.6)$$

In order of the variables, the units are keV, keV/DAC, DAC and keV. ‘x’ and ‘y’ refer to the pixel coordinates and ‘thr’ indicates the threshold (zero to one or zero to seven). All terms depend on the threshold used (‘thr’) as it maps to different physical pixels which can share the same x-y coordinates. For a single chip, the number of slope and intercept values is the number of physical pixels (256×256) multiplied by the number of thresholds per physical pixel (two), therefore $256 \times 256 \times 2 = 131072$. For a four chip system, there are over half a million slopes and intercept values which compromise one complete per pixel energy calibration.

An example measurement follows using one simultaneous measurement with eight thresholds. The method can be used with a single threshold by scanning that threshold over energy.

Plotting the raw counts for a single arbitrarily chosen pixel ($x = 184$, $y = 111$) against the chip wide threshold using eight thresholds simultaneously results in figure 5.16a. From equation 5.6, we can plot the energy calibration per threshold using the chip dependent threshold DAC value ($thr0 - 7$, $threshold_DAC_{chip,thr}$), pixel specific slope (a , $slope_{x,y,thr}$) and intercept (b , $intercept_{x,y,thr}$). This results in figure 5.16b. The effective energy per threshold can then be plotted using these calibration values, see figure 5.16c. Finally, the threshold in DAC units can be mapped to energy for that single pixel. This is figure 5.16d which shows the counts for each energy calibrated threshold that this single pixel was set to. The uncertainty in counts is the square root of the number of counts, \sqrt{N} . This is shown in the inset plot in figure 5.16d.

This particular example shows an interesting edge case involving thresholds one and two. It shows how it is possible to have a higher threshold value in DAC units but a lower correspond-

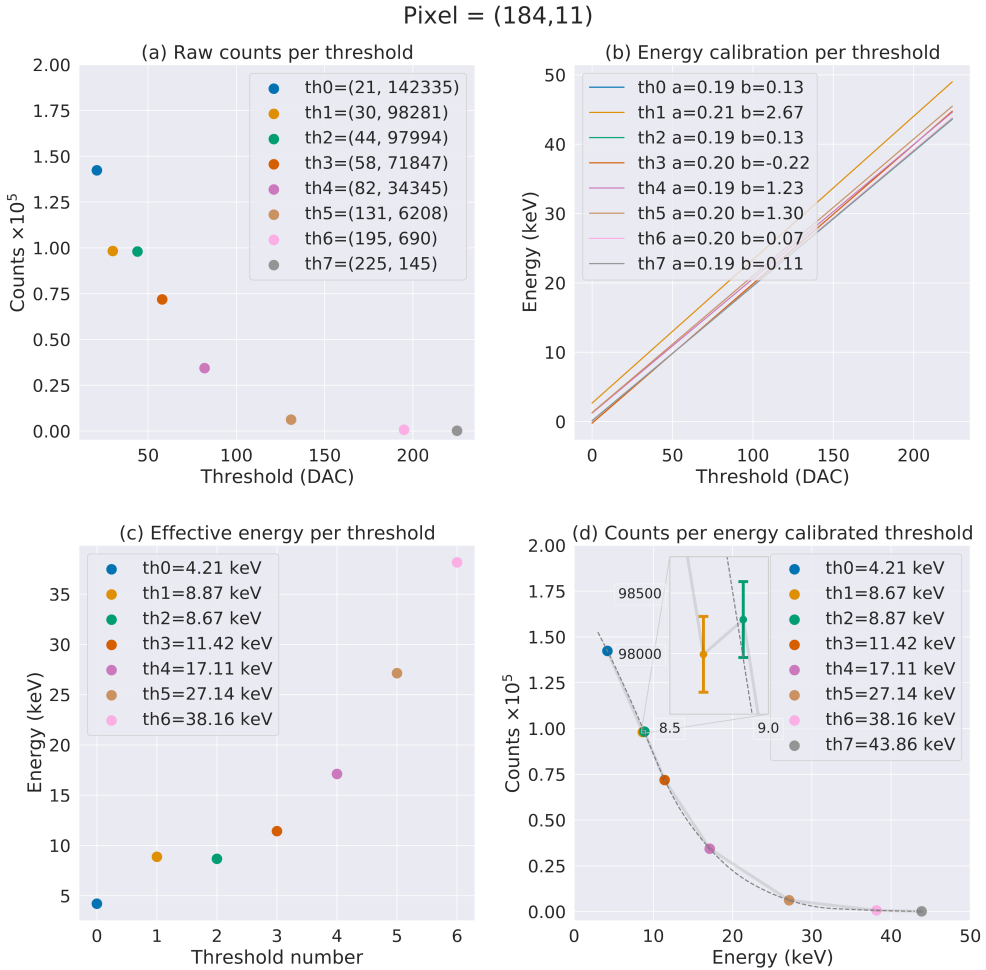


Figure 5.16: An overview plot of the energy calibration per pixel process. This used all eight thresholds and an arbitrary pixel was chosen with coordinates of $x = 184$ and $y = 11$. A Si based detector was used in SPM with a peak x-ray tube voltage of 50 kVp. The thresholds were set from 21 to 225 DAC units as seen in (a). The energy calibrations per threshold are in (b). This leads to the effective energy per threshold for this pixel in (c). Finally, the counts per energy calibrated threshold are shown in (d), the uncertainties plotted are the photon counting errors, \sqrt{N} .

ing energy. For the same threshold value in DAC units, ‘th1’ will always map to a higher energy compared to all other thresholds because ‘th1’ has a higher y intercept value ($b = 2.67$).

At this point, it is possible to estimate the actual number of photons that were detected at any energy within the measurement range. Linear interpolation of the counts over energy yields an estimation of the number of counts at a specific set of energies. This is the core of the method. It depends on spectrum being smooth between sampling points and sampled relatively densely - not having a transition between points. It does not need fitting to a function.

The counts of a single pixel have an uncertainty. Part of that is due to photon counting and is equal to the square root of the number of counts, \sqrt{N} , as discussed in section 2.6.4. Given this fundamental uncertainty, it is ideal to use this method with the minimal achievable photon counting uncertainty. If the spectrum is known to be free of transition energies and is therefore smooth, it should be possible to fit a function to the data and then do the interpolation. This is a topic for further research.

COMPUTING NOTE

This energy calibration per pixel algorithm is highly parallelisable. Almost all the time is spent in the parallelisable portion of the code. It has been implemented in Python using the numpy and joblib libraries in a partially parallelised approach.

Specifically, the work is split into as many chunks as CPU threads are available, handed to individual CPU (central processing unit) threads, they then work serially within those chunks and finally, the results are then summed along the fourth axis to form a complete data set. On a modern desktop (‘AMD Ryzen 7 5800X’ CPU with eight physical cores & 16 logical threads), it takes eight seconds per image stack to process, resulting in as many files as interpolated energies are requested. The number of interpolated energies requested only impacts the run-times in terms of the time taken to save the output to disk. Improvements to execution time could of course be made.

In terms of speeding up the run-time, there are two reasonable options broadly speaking. One would be to use a lower level compiled language such as C/C++ to process the parallel portion of the code and the second is using non CPU hardware.

GPUs (graphical processing units) could be used to fully parallelise this method using frameworks such as NVIDIA’s CUDA (Compute Unified Device Architecture), OpenCL (Open Computing Language) from the Khronos group, or even with FPGAs (Field Programmable Gate Array). This algorithm has no dependency on neighbouring pixels or previous states and so is an ‘embarrassingly parallel’ problem. GPUs are particularly interesting due to how GPU capabilities have been scaling significantly faster than CPUs in terms of throughput and how relatively inexpensive they are and are not excessively hard to program.

5.5.3 MEASUREMENT OVERVIEW

For the following measurements the geometry (relative position x-ray source, object, detector) was kept constant while changing other parameters. The varied parameters were the x-ray tube peak voltage and current, the detectors and the objects. For these measurements, 30 and 90 kVp were used for both detectors, see table 5.1. The source to detector distance was maintained at 285 ± 10 mm and the detector to object distance was kept at 6—6.5 mm. For each of these, an air measurement with the same acquisition settings was taken over the same measurement time.

Table 5.1: An overview of the main measurement parameters used in this section, there are four configurations in total. The tube voltages and currents are the first column. The first row contains the detector material, thickness and bias voltage. The other contents show which threshold was used and in which mode (CSM or SPM).

| | Si 110 μm +150 V | GaAs 55 μm -130 V |
|--------------------------|-----------------------------|------------------------------|
| 30 kVp 200 μA | th1, SPM | th1, CSM |
| 90 kVp 89 μA | th1, SPM | th1, CSM |

The Si based detector was in SPM and is the same device used in the rest of this work; the thickness was 500 μm and biased to +150 V. The GaAs detector was operated in CSM with a sensor thickness of 500 μm and biased to -130 V, a significant fraction of the pixels were not well connected to the sensor. Only threshold one was analysed for all data-sets. The eight thresholds available using Si in colour mode (110 μm pixel pitch) do not behave very differently to each other and so only analysing one threshold was deemed sufficient for these measurements. For GaAs, SPM or SPMa were not possible to calibrate in terms of energy either with single pixels or with entire chips and so was not used. This is not the case for the CSM threshold as shown in subsection 5.4.3.

5.5.4 RESULTS AND DISCUSSION

This analysis is an exploration of the various methods in which ECI improves both imaging and spectroscopic performance.

Energy Calibrated Image (ECI) combined with a Flat Field Correction (FFC) are referred to as ECI_FFC. The idea is that both the image and flat field should be corrected in terms of energy and then normal flat field correction is used, resulting in the best possible image.

The first analysis is SNR as a function of energy for raw images, flat-field corrected (FFC), energy calibrated images (ECI) and ECI_FFC (Energy Calibrated Image (ECI) combined with a Flat Field Correction (FFC)). This is calculated for both detectors at both x-ray tube settings. SNR is defined in subsection 2.6.1 as $\text{SNR} = \frac{\mu^2}{\sigma^2}$. The effects of ECI_FFC compared to FFC images have been hinted at in figure 4.15. More examples are shown here with different samples.

Comparisons are made between ‘ground truth’ high energy resolution scans and simultaneous multi-energy imaging in two ways in terms of generating ECI.

The ratio of linear attenuation coefficients in equation 5.5 is explored in combination with ECI. As part of this, transmission spectra at the same relatively high intensity are briefly explored with Medipix3 using molybdenum (Mo) both with and without a filter.

SNR AS A FUNCTION OF ENERGY

The SNR as a function of energy for the Si and GaAs detectors at 30 and 90 kVp is shown in figure 5.17. The same experimental conditions were kept between measurements; it is important to note that these conditions are not the optimum in terms of SNR with the CDMAM v3.4 phantom from the University Medical Centre in Nijmegen, the Netherlands. A picture of the phantom is included in figure 5.18. As of 2014, the CDMAM phantom v3.4 was ‘the standard test object used to assess the imaging performance of digital mammography systems in Europe’ (63; 64).




Figure 5.17: SNR (signal to noise) ratios of a CDMAM v3.4 (Contrast-Detail Phantom for Mammography) phantom from the University Medical Centre in Nijmegen, the Netherlands. The peak voltage of the x-ray tube was 30 and 90 kVp and Si and GaAs detectors were used.




Figure 5.18: A picture of the CDMAM phantom type 3.2. It has an array of gold discs of various thicknesses and diameters on a black aluminium plate, embedded within a PMMA case. There are three metal sheets taped to the front surface to help identify which part of the grid is being imaged.

The x axis is the energy of the threshold in keV, this means all photons at that energy and above are included for each datum. The SNR uncertainties plotted are the differences between the values from the two gold disks within a single square of the CDMAM phantom with thicknesses of $25\text{ }\mu\text{m}$ and diameters of 1.25 mm . The uncertainties in energy are the sum of the threshold dispersion and the energy resolution for the non-ECI data (raw and FFC). Given that the step size is approximately 0.2 keV for Si and 0.5 keV for GaAs, half of that step size respectively was used for the uncertainty in energy in the ECI data. The horizontal pink line in every subplot indicates the Rose criterion, an SNR equal to five as discussed in section 2.6.1.

One would expect that the highest SNR over energy would be ECI_FFC, then ECI or FFC and the lowest should be the raw images. In these experimental conditions, across both 30 and 90 kVp measurements and by comparing the average SNR over energy, the option with the best SNR as a function of energy is ECI_FFC, followed by FFC, ECI and finally raw images.

The CDMAM phantom is a combination of aluminium and PMMA, the total x-ray attenuation is equivalent to 10 mm of PMMA (65). The flat field images taken for these measurements were just using air, not 10 mm equivalent of PMMA. This mismatch is the reason why for the Si detector, the FFC does not improve SNR over all energies and in fact reduces it at certain energies. This demonstrates the importance of optimising the material used for the flat field correction, simply using air is insufficient. Previous measurements using the same Si detector and similar conditions achieved SNR of over 200 compared to a maximum of $30\text{--}40$ with these conditions. The fact that the ECI_FFC is not equal to or above the ECI at all energies indicates that the flat field is not a good correction for the raw image or the ECI. The GaAs data shows that while the ECI provides a boost to the SNR from the raw images, the FFC is by far the dominant factor in this case.

Another example follows with a measurement of a bee, viewed at both 7.0 and 27.0 keV in figure 5.19. At 27.0 keV , the relative normalised differences between the raw images compared to ECI and the FFC compared to the ECI_FFC are significant. The improvements are qualitatively evident in the image quality which matches the SNR as a function of energy findings. The fact that at 27.0 keV the image becomes uniform after energy correction proves that the method is consistent and robust, otherwise the pixel responses would not converge and the SNR would be degraded.

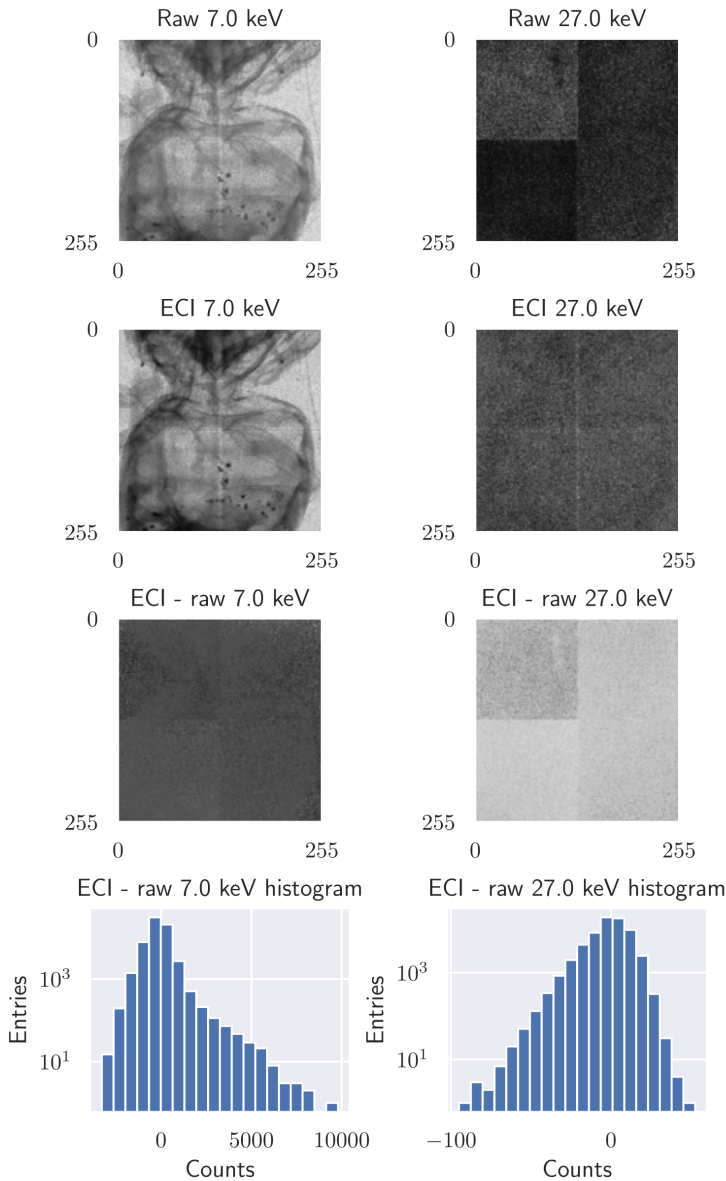


Figure 5.19: The differences between raw images and ECI at 7.0 and 27.0 keV of a bee thorax and abdomen with histograms of the differences. The x-ray tube voltage used in this measurement was 30 kVp with a tube current of 200 μ A. A high energy resolution threshold scan was used as input data for these reconstructions using one threshold. At 7.0 keV, there were between 1000—10000 counts within the image and for 27.0 keV there were 0—130 counts. The scan lasted 56 seconds and did not use a flat field correction. The bee is barely visible at 27.0 keV due to the x-ray absorption of the bee and air converging.

Qualitatively it can be observed that the order of increasing image quality is raw, ECI, FFC and then ECI_FFC. This is consistent with the SNR as functions of energy plots. When comparing in image space, the systematic errors fixed with ECI are increasingly more evident as energy increases. This is caused by the design of the chip where the gain of pixels cannot be fine-tuned; this is planned to be addressed in Medipix4. Ideally, the pixels would be designed and fabricated in such a way that they respond identically.

HIGH ENERGY RESOLUTION SCAN COMPARED TO SIMULTANEOUS MULTI-ENERGY IMAGING FOR ECI

A threshold scan of air using one threshold and the minimum available step size (1 DAC unit) which corresponds to 0.236 keV ($\sigma = 0.008$ keV/DAC) for this Si detector are now compared against a single multi-threshold measurement. For the threshold scan, threshold one was scanned from 400 down to 30 DAC units corresponding to ~ 90 to ~ 5.9 keV. For the single multi-threshold measurement, all eight thresholds available with the Si detector were used and were not scanned. The tube settings were maintained at 90 kVp and 89 μ A with the same distances.

The system would have been affected by various drift effects during the threshold scan as previously discussed. This is not the case for the single multi-threshold measurement; the incident photons are precisely the same between each threshold and correspond to precisely the same starting and ending frame times.

In terms of generating ECI, there are compromises to be made between energy range, energy resolution and measurement time. For example, a short measurement with a large energy range can be made at the expense of energy resolution. Inversely, a long measurement with a large energy range can be made with the best energy resolution. The final example, a short measurement can be made with the best energy resolution but with a small energy range. All of these options are valid and potentially relevant depending on the application. This is related to the lower right subplot (d) in figure 5.16.

A metric used to consider which approach is preferable is the relative residual (between reconstructed number of counts vs ground truth) as a function of energy. For this measurement, the ground truth is considered to be the threshold scan data. This metric is used in figures 5.20 and 5.21, it shows three scenarios. The first is the naive approach of spacing thresholds equally across the full energy range; this is ideal in cases of a totally unknown spectrum. The second is an optimisation of the first approach, spacing the thresholds closer together in the low energy region at the expense of higher relative residuals at high energies. Finally, if only the low energy region is interesting, as is often the case for biological and thin samples, the thresholds are equally spaced with a spacing equal to the energy resolution at the expense of a significantly smaller energy range. It can be observed that the best relative residuals are found with the third approach.

The upper half of figure 5.20 shows the same type of data as in figure 5.16d. Figure 5.16d shows one measurement using eight thresholds simultaneously and is displayed using a linear scale on the

y axis, whereas figure 5.20 scans one threshold with the minimum step size in energy ($\sim 0.2\text{keV}$) and uses a logarithmic scale on the y axis.




Figure 5.20: The black line is the counts as a function of energy for a high energy resolution scan for a single arbitrary pixel ($x=100, y=100$). Three approaches are presented for ECI in terms of threshold spacing and their respective relative residuals as shown in the legend. A Si detector was used with a x-ray tube voltage of 90 kVp and only air (flat-field) was measured.

The optimal approach therefore depends on the shape of the spectrum and the compromises as mentioned. With knowledge of the shape of the spectrum, it would be possible to improve on the relative residual of all approaches. This is due to how ECI is implemented in this work using linear interpolation only. While this is a very fast operation and is a generic implementation (completely spectrum shape independent), it is not the most accurate for x-ray spectra given that we have knowledge about how the linear attenuation coefficient of all materials scales with energy, $\propto \frac{1}{E^3}$. This statement, of course, does not take transition energies into account. This is only an issue if the object measured has at least one transition energy in the energy range scanned.

Figures 5.20 and 5.21 show the counts of the high resolution scan, the three approaches outlined above and the relative residuals of respectively both a single arbitrary pixel and the mean counts of all pixels in this detector.




Figure 5.2.1: The counts as a function of energy for a high energy resolution scan for all pixels in this Si detector with three approaches for ECI in terms of threshold spacing and their respective relative residuals. A Si detector was used with a x-ray tube voltage of 90 kVp and only air (flat-field) was measured.

Another interesting aspect to consider is the difference between a simultaneous eight threshold measurement and, selecting the same eight threshold values, a subset of a high energy resolution threshold scan. Ideally these would be equal, however there are time dependent drift effects and quantum noise so this is not the case.

Simultaneous measurements benefit from measuring the exact same photons at the exact same time as each other, thus reducing it to the quantum noise corresponding to number of photons in that bin, instead of the much larger number of photons above that threshold of two images. Since this uncertainty only dominates at very low counts, they therefore benefit very low statistics or signal measurements primarily. They can also be used to scan the energy range of interest with greater effective resolution than scanning with a single threshold or scan a given range faster, inversely proportional to the number of thresholds using a ‘comb’ type scan. A comb scan is one where inter-threshold spacing is maintained and the thresholds are incremented/decremented until they overlap in energy.

The subset approach is ultimately limited by the amount of drift present in the system. Therefore the optimal approach depends on the compromise between the statistics of a single simultaneous measurement and drift rate during a threshold scan.

RATIO OF LINEAR ATTENUATION COEFFICIENTS AS FUNCTIONS OF ENERGY

Now we will cover a series of measurements using both Si and GaAs detectors; some measurements have materials with transition energies in the measured energy range. Transmission spectra are shown for the materials with transition energies in the scanning range. These energies can be observed and matched to a limited range of specific elements.

For the Si detector, example measurements are shown using the same beam and data as used for figure 5.19 along with the CDMAM phantom with Mo (molybdenum) foils at 30 and 90 kVp. One measurement with transition energies is shown for the GaAs detector - the two Mo foils with a Pb foil overlapping each other.

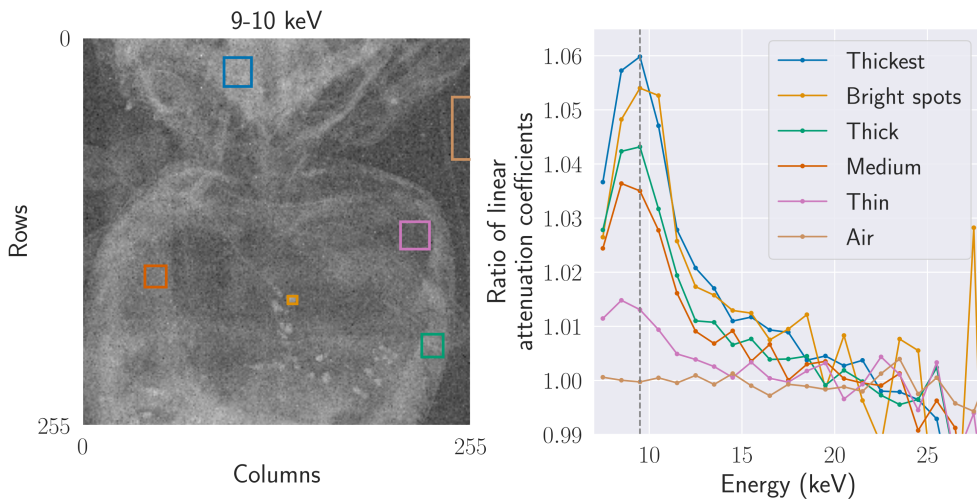


Figure 5.22: The relative linear attenuation coefficients over energy of a bee's thorax and abdomen with a corresponding labelled image at 9.0–10.0 keV as indicated by the vertical dashed line. The x-ray tube voltage used in this measurement was 30 kVp with a Si detector. The coloured squares show the manually selected region of interest where the median of that part of the image over energy was used to generate the line plots.

The dried out bee is the most simple of the samples in terms of composition. As seen by the detector, it is a combination of just two base materials, chitin and air. Chitin is a polymer produced by bees for their exoskeletons. Figure 5.22 plots the ratios of linear attenuation coefficients as a function of energy for several regions of interest. It can be observed that since the correction is using air, the air is flat across energy until the energy approaches the peak tube voltage at 30 kVp where the counts approach 0. Based on the data, and the fact that this was very dried out bee, it appears that the 'thin' and 'thick' regions of the bee are combinations of the bulk material (chitin) and air.

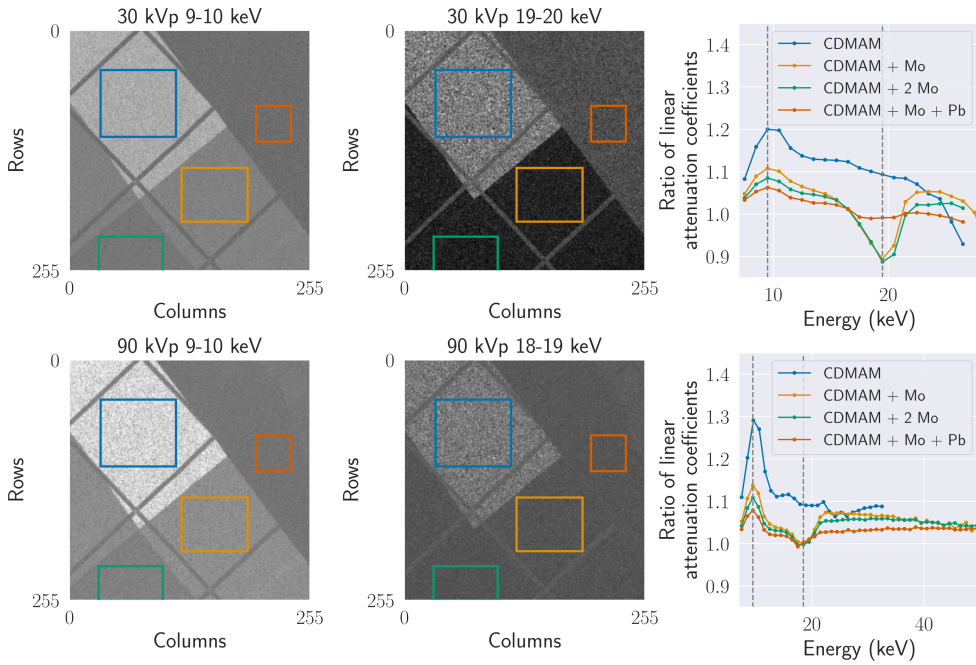


Figure 5.23: The relative linear attenuation coefficients over energy of a CDMAM v3.4 phantom with corresponding labelled images at 9.0—10.0, 18—19 and 19—20 keV as indicated by the vertical dashed lines. The x-ray tube voltage used in these measurements were 30 (first row) and 90 kVp (second row) with a Si detector. The coloured squares show the manually selected region of interest where the median of that part of the image over energy was used to generate the line plots. The viewing windows are matched across all images from 0.82 to 1.38.

The CDMAM phantom was used for figure 5.23, however there were Mo and Pb foils overlaid to easily identify which square was being imaged. This was necessary because the detector is a fraction of the area of the CDMAM phantom. These foils were useful for both masking a specific square and as an example transmission spectrum measurement of Mo foil with a filter effectively. Mo has three transition energies in the scanning range, they are at 17.479, 17.374 and 19.608 keV and have respective relative intensities of 1.00, 0.52 and 0.15. One can observe minimum points in both 30 and 90 kVp data, however they are slightly shifted relative to each other. At 30 kVp, the minima for the ‘CDMAM + Mo’ and ‘CDMAM + 2 Mo’ is at the 19—20 keV point whereas for 90 kVp, it is at 18—19 keV. Both are higher in energy than the true maximum intensity transition energy corresponding to the $K-\alpha_1$ transition at 17.479 keV. That material is the CDMAM combined with the air which corresponds to at least an equivalent attenuation of 10 mm of PMMA. There is no such minimum for the ‘CDMAM’ only data; this shows that the effect is actually originating from

the Mo foil and not the CDMAM or any other effect such as the detector response or spectrum shape. The minima is relatively smaller for the ‘CDMAM + Mo + Pb’ data at 30 kVp compared to 90 kVp because more photons within that energy range were absorbed at 30 kVp compared to 90 kVp.




Figure 5.24: The relative linear attenuation coefficients over energy of Mo foils and air, with the corresponding labelled images at 17.0—18.0 and 25—26 keV as indicated by the vertical dashed lines. The x-ray tube voltage used in this measurement was 90 kVp with a GaAs detector. The coloured squares show the manually selected region of interest where the median of that part of the image over energy was used to generate the line plots.

The measurement in figure 5.24 in shows the differences in relative linear attenuation coefficients between one Mo foil (‘Mo x1’) and two overlaying Mo foils (‘Mo x2’) of the same thickness as ‘Mo x1’ along with the air. The consistency between both single Mo foil measurements can be seen. The differences between the double foil and the single foil data are that the peak is slightly higher for the double foil data and it drops to low values after the transition energy and stays lower until the end of the scanning range. It can be seen that the air data has deviations from the expected value of one across the energy range. The key result from this plot is that the peaks of the Mo foil measurements, with only air in the beam path apart from the Mo foil, match the most intense transition energy of Mo at 17.479 keV. This supports the argument that the CDMAM acting as a filter in figure 5.23 caused a shift of 1—2 keV.

Ultimately, these measurements open up possibilities with the Medipix3 detector that were not possible before - massively parallel transmission spectra measurements with energy calibrations per pixel over a relatively large area with a minimum spatial resolution of 55 μm . While this will never provide better energy resolution than dedicated spectrometers, such as a silicon drift detector or high purity germanium detectors, this solution enables very high speed spectra measurements (in the order of $100 \text{ kHz} \times 512 \times 512 = 2.6 \times 10^{10}$ counts/second over a FPM two by two Medipix3 detector) with energy resolutions as good as the sensor and chip combinations enable. Both of these aspects are continually optimised over time due to the design of hybrid pixel detectors, the sensor

layer (e.g. Si) is separated from the readout ASIC (e.g. Medipix3).

5.5.5 SUMMARY

The impact of extending ECI to CT reconstructions with no impact to scan time has been shown in figure 4.15 using a microCT phantom object. Addressing the underlying primary physical reasons causing ring artefacts by calibrating each pixel in terms of energy is a powerful, generic method. It does not rely on any a priori knowledge about the spectrum or the object being imaged, only knowledge about the detector is required. It also does not impact measurement time at all since it is post-processing and is software-based. As shown in figure 5.17, ECI always has higher SNR than raw images and ECI_FFC has higher SNR than FFC. Figure 5.17 also shows that the flat field correction itself should be optimised for the object by introducing equivalent x-ray attenuation material of the average thickness of the object being imaged. It is not perfect however, the energy calibration per pixel will change over time, there is still random telegraphic noise with Medipix3 and there are various drift effects.

Future work could involve investigating the energy bins between the reconstructed images from ECI. The optimal energy bin size would be a point of interest to investigate; it would likely be limited by the energy resolution. This is related to figure 5.17. This could be used for beam hardening artefact reduction by calculating the differences between energy bins. In the medical context, this is known as metal artefact reduction from typically implanted prostheses like teeth fillings, screws and so on.

ECI_FFC CT could be investigated further with the acquisition of a threshold scan for each projection of a CT scan. Depending on the measurement time, drift effects may become significant.

Using the CDMAM v3.4 it should also be possible to conduct a detailed study into the differences between imaging performance of the Medipix3 with different sensors and configurations to existing clinical systems. This could involve dose calibration using clinically calibrated dose measurement systems.

5.5.6 EXISTING EFFORTS

There has been an attempt to solve the deficiencies of hybrid pixel detectors with a similar method published in 2017 using a 110 μm pixel pitch CdTe sensor and a Medipix3 detector operated in CSM (66). It used a variance and mean-based cut and a relatively small region of interest (a 10×10 pixel area) of seemingly representative pixels. It exclusively presents the differences in terms of energy resolution improvements and not how this propagates for imaging. The main limitation of this method is that the reconstructed energy spectrum depends on the neighbouring pixels.

The main difference between the reference and this work are detector related. We use two different detector configurations, both a 55 μm pixel pitch GaAs sensor operated in CSM and a 110 μm Si sensor using SPM. All pixels are used for the analysis. Both impacts of this analysis on energy resolution and imaging are explored. The energy range used in this work starts and ends at

lower values. The lower limit is lower in this work due to the reduced noise floor and the maximum energy is limited by the x-ray tube to 90 kVp. The overall conclusions are in agreement with and complement the previous work from 2017. This work presents a method which does not depend on neighbouring pixels. We know the energy response per pixel for all thresholds, and that is all the information required to estimate the number of counts at a specific threshold level at that energy.

A paper in the life sciences field shows the possibilities of x-ray fluorescence imaging with spectrometer level energy resolution (~ 150 eV) using synchrotron radiation(67). It covers a number of imaging modalities including x-ray fluoroscopy tomography. It includes examples of high spatial and high energy resolution imaging of biological samples such as human brain and larval zebrafish with iron, zinc, copper, mercury and phosphorus separated into overlaid RGB (red, green, blue) channels. This shows the wealth of information that can be extracted using the best radiation source and energy sensitive x-ray detectors. These methods are effectively restricted to those with access to synchrotrons, require raster scanning of a micro-focused beam and the maximum count rates are limited by the capabilities of the spectrometers used. As an approximate estimate, this paper shows a saturation count rate of approximately 800 kHz over the entire detector; this is 3–4 orders of magnitude lower than for a single Medipix3 chip if the radiation was distributed evenly over the chip. The best measurement method to use depends primarily on the available equipment, but ultimately, the acceptable compromise between precision and measurement time.

5.6 SUMMARY

Spectral CT reconstruction methods are an area of active research.

One such recent example is an ‘iterative, hybrid reconstruction technique which combines the spectral properties of PCD (photon counting detector) data with the resolution and signal-to-noise characteristics of EID (energy integrating detector) data’ (68). Pre-clinical and clinical systems are starting to enter usage (69). They have ‘demonstrated superior spatial resolution relative to current CT systems and improved noise properties and multi-energy temporal resolution relative to similarly configured, energy-integrating detector CT’ (70). Attempts to optimise k-edge imaging using spectral CT have been made (71). There have been efforts to compare spectral to conventional CT for breast and dental imaging and with dual contrast (72–74).

The benefits of hybrid pixel detectors can be realised with further development of the platform, mainly by increasing the area of the detector to at least $10 \times 10 \text{ cm}^2$, improving system stability, improving the count rate capabilities and decreasing the system cost through mass production.

The SDD has an energy resolution approximately one to two orders of magnitude better than the Medipix3. On the other hand, Medipix3 has a spatial resolution of one to two orders of magnitude better than the SDD, $55 \mu\text{m}$ vs. $\sim 1 \text{ mm}$. GaAs seems to really require CSM to do any spectral studies. GaAs being more absorbing at higher energies relative to Si makes it a complementary sensor material for Si. In order to boost the higher energy absorption of Si, thicker sensors would be required. Thicker Si sensors have their disadvantages, including increased charge sharing, increased

research and development costs and reduced spatial resolution.

ECI relies on per-pixel energy calibrations which result from the linear response of Medipix3 pixels over this energy range. These calibrations are then combined with energy dependent information, this could be obtained with a threshold scan or with multiple thresholds set to different energies. Taking the different pixel responses into account, the interpolated counts can then be extracted per pixel for any energy within the measured energy range. This results in at least one ‘energy calibrated image’.

Examples of how ECI impact 2D projection and 3D CT imaging are included in figures 4.15 (columns 2—4) and 5.19. They show the improvements relative to raw images and flat field corrected images at two energies, one near the noise floor and one near the peak tube energy. This shows how ECI improves over equalisation alone and that the method scales.

The optimal approach for ECI is ultimately application dependent. Three options are explored in subsection 5.5.4. Energy resolution can be traded for energy range.

Spectral X-ray CT applied: Re-excision after positive margin

6.1 CURRENT CLINICAL PRACTICE

According to the World Health Organisation, ‘In 2020, there were 2.3 million women diagnosed with breast cancer and 685,000 deaths globally. As of the end of 2020, there were 7.8 million women alive who were diagnosed with breast cancer in the past five years, making it the world’s most prevalent cancer’ (75).

Currently, after the tissue is located with a CT or MRI (Magnetic Resonance Imaging) scan, it is excised, sent to the pathologist who fixes the tissue in a formalin solution, embeds it in paraffin, slices, stains and inspects it. This method results in FFPE (formalin-fixed paraffin-embedded) tissue blocks which is a commonly used protocol to archive tissue for many years with minimal degradation. By visual inspection, the surgical margin is assessed.

The surgical margin is the shortest distance between the tumour and the boundary of the excised tissue. If that distance is too small, there is an excess risk that tumour tissue has been left behind in the patients. This results in reduced local control (stopping the local cancer growth) and possibly also reduced patient survival rates. Definitions of which minimum distance is required are dependent on tumour type, stage, and country. In case of a positive margin (when cancer cells are still present in the tissue) additional treatment is needed, most often in the form of a re-surgery.

The current gold standard method does have shortcomings however. The whole process of using FFPE blocks and histology typically takes several days to several weeks, requiring the patient to be closed up and increasing the probability of reoperation. The amount of tissue assessed is typically < 1% of the excised tissue. By inspecting 2D histology slices, only the in-plane margin can be assessed

6. SPECTRAL X-RAY CT APPLIED: RE-EXCISION AFTER POSITIVE MARGIN

for a given slice whereas this is actually a 3D problem; there is relevant variation between slices. There is significant noise in the observation by the pathologists; if you take several pathologists, give them the same image, their answers will not be the same as each others'. The act of very thinly slicing these brittle FFPE tissue blocks is not risk free, it frequently damages the slices and depends on the equipment used and operator skill. The fixation in formalin and subsequent embedding in paraffin is a chemical process which alters the structure due to the dehydration and clearing.

Clinically, the issues arising from this include 'substantial physical, psychologic, and financial burdens for patients, with higher risk of complications, worse cosmesis, and additional costs, on average, of >\$10,000 per patient' (76).

6

Depending on many factors including tumour properties and country, up to 30% of patients of such surgeries require a re-surgery. For example, there was a study of patients in English NHS trusts entitled 'Reoperation rates after breast conserving surgery for breast cancer among women in England: retrospective study of hospital episode statistics'. The study was summarised as 'One in five women who had breast conserving surgery in England had a reoperation. Reoperation was nearly twice as likely when the tumour had a carcinoma in situ component coded. Women should be informed of this reoperation risk when deciding on the type of surgical treatment of their breast cancer.' The study makes an important point - 'Few studies have examined reoperation rates after breast conserving surgery, and uncertainty exists about how likely women are to need reoperation. Estimated rates from international hospital based studies have ranged from 17% to 68%' (77–81).

Evidently, this is a significant societal issue due to the number of cancer related deaths per year and therefore many approaches have been attempted. Visible light-based methods offer excellent SNR, spatial resolution and measurement time at the cost of being depth limited to ~ 2.5 mm. They produce results comparable to the current gold standard method, H&E staining, and include ultraviolet photoacoustic remote sensing, micro-elastography and optical coherence tomography (OCT). It is important to note that these methods are experimental and not yet widespread.

Further details can be found in the following works:

- Ultraviolet photoacoustic remote sensing microscopy, $0.7\ \mu\text{m}$ resolution (82);
- Reflective objective-based ultraviolet photoacoustic remote sensing virtual histopathology, $0.39\ \mu\text{m}$ resolution (83);
- Optical coherence tomography and micro-elastography, $2\text{--}10\ \mu\text{m}$ resolution, up to $1\text{--}2$ mm depth (76);
- Photoacoustic microscopy (84);
- Finger mounted OCT probe (85);
- Raman spectroscopy and imaging, 'sub-cellular' resolution, $\approx 20\text{--}100\ \mu\text{m}$ depth (86–88).

M4i (Maastricht MultiModal Molecular Imaging Institute) introduced a mass spectrometry based method to address this problem and was awarded the NWO physics valorisation prize in 2019. This offers the possibility of tumour specific treatments.

The most commonly used method in clinical practice today is 2D x-ray projection imaging (radiography). There was a related systematic review and meta-analysis of 798 patients ‘using radiography, digital breast tomosynthesis (DBT), micro-CT, and ultrasound for intraoperative margin assessment with the histological assessment as the reference method’ (89). The authors found that while this is very fast and does not require a microCT machine, it has insufficient sensitivity (true positives) and also has noise in the observation of where the tumour boundary is. The net result is that in practice, it performs the worst out of the four methods reviewed: radiography, digital breast tomosynthesis, micro-CT and ultrasound. Based on this paper, for radiography the pooled sensitivity (true positive rate) was only 52% and specificity (true negative rate) was 77%. This means that the most commonly used method has a 48% false negative rate, it misses nearly half of imaged tumours. Clearly, there is room for improvement.

The targeted scan time was less than 10 minutes and a statistical uncertainty contribution was chosen to be $< 1\%$ which requires at least 10,000 counts per pixel within the region of interest. This choice was a trade off between measurement time and minimising the statistical uncertainty per projection. This value depends on this particular x-ray tube, peak tube voltage, tube current, pixel size, magnification factor among other aspects. This is worth studying in further detail for optimisation of this process.

6.2 HUMAN BREAST TUMOUR - 2018

In 2018, there was an exploratory joint measurement with the NKI (Dutch National Cancer Institute) involving a spectral microCT system for online evaluation of surgical margins in three dimensions. The results are shown in Figure 6.1. For the spectral image (right), the data is clustered into 18 categories using the energy information, each colour is categorised as a different tissue combination. The homogeneous region in the top left with a high density of calcifications (bright yellow) is the tumour. Material types in this sample are air (corners), microcalcifications, cancer tissue, functional breast tissue and fat.




Figure 6.1: Breast tissue specimen (approximately $20 \times 20\text{mm}^2$) imaged with the Medipix3 detector. Left image photograph of the specimen; Centre: a classic x-ray projection image (all energies are integrated); Right: the spectral image after an energy based clustering reconstruction.

This measurement was considered to be promising enough to continue subsequent studies. This led to the mouse breast tumour scan three years later in 2021.

6.3 MOUSE BREAST TUMOUR - 2021

In 2021 it became possible to resume tumour scanning due to new access to samples via Dr Jasper Nijkamp, at this time based at the Danish Centre for Particle Therapy (DCPT) at Aarhus University Hospital. During these three years, detector characterisation and spectral microCT methods were developed and refined with this type of measurement in mind. This includes addressing all artefact types, such as ring, beam hardening and RTN effects.

The mouse breast tumour sample (labelled ‘c7006-25’) was a FFPE tissue block fitting with a volume of $1 \times 1 \times 1$ cm. FFPE samples were used over flash frozen fresh samples primarily so that the samples did not need to be cooled during transport or measurements. Secondly, there are no motion artefacts due to the sample thawing during measurements. Finally, measurements can be repeated essentially without limit which is ideal for exploratory measurements. However, it is not known if contrast within FFPE samples is comparable to contrast within fresh samples. It is expected that because FFPE has very little to zero water content compared to fresh tissue, the x-ray contrast should be higher for the flash frozen and fresh tissue. This expectation should be confirmed with further tumour measurements of FFPE blocks, flash frozen samples and fresh tissue.

c7006-25 was an artificially induced very rapidly growing tumour as part of radiotherapy research at the DCPT. This was from a non-irradiated control mouse. As determined by the calibrated values from the H&E stained slice, the maximum width, height and depth of the sample was $5.73 \times 7.28 \times 1.93$ mm.

Like the measurement in 2018, this was an exploration of whether spectral x-ray microCT could lead to a prototype imaging system which could enable high speed margin assessment and tumour

segmentation during surgery.

This method would be no more invasive than current protocols. It would not require surgeons or operating room personnel to learn new methods or equipment. No staining or slicing is needed, this method is non-destructive; therefore the sample could be processed in exactly the same way as it is currently. This would require the presence of a microCT scanner near or in the operating room with one person to interact with the scanner.


These characteristics indicate minimal disruption and do not present any ‘show stoppers’. No other system is known to exist with such characteristics. If the aims were achieved, this system has the potential to revolutionise surgical margin assessment by providing online feedback to the surgeons. This will allow for highly directed removal of extra tissue when needed, preventing re-surgeries and adjuvant treatments. Adjuvant treatments are treatments given after the primary treatment to reduce the risk of a recurrent cancer.

6.3.1 MEASUREMENTS AND RESULTS


We will look at overviews of the various methods used to achieve the best CT reconstruction quality; ECI_FFC, registration and the ‘optimal’ measurement time. This is followed by H&E (hematoxylin and eosin) stained slices of the tumour compared to the reconstructed CT slice from Medipix3. Finally, there will be a volume rendering with colour corresponding to energy bins as a presentation of the spectral information within this tumour.

Figure 6.2 presents a H&E stained slice both unlabelled and partially labelled by Dr Jasper Nijkamp. The tissue types included are tumour, tumour outline, necrotic tumour, blood vessels, an inflammation area and the hair / follicles. There are three circles caused by air bubbles in the lower right which are artefacts.

Figure 6.3 shows the registered CT slice corresponding to the same volume as the H&E slice along with two pictures of the sample, the tumour is encased in paraffin.




(a) Unlabelled slice.




(b) Labelled slice.

Figure 6.2: The unlabelled (a) and labelled (b) H&E stained slice scan, the current ‘gold standard’. This is one 2D slice where the slice thickness is in the order of $5-10\ \mu\text{m}$. All dark purple areas are associated with active dividing tumour cells, some of which are indicated with the green outline. The ‘tumour outline’ indicates the boundary between tumour and non-tumour, healthy tissue.



(a) Reconstructed CT slice.



(b) Two pictures of the sample.

Figure 6.3: (a) is the corresponding Medipix3 CT slice from the ECI_FFC reconstructed at 5 keV. It should be compared directly to figure 6.2. This is one of more than 50 reconstructed slices with a thickness of $\simeq 33 \mu\text{m}$. The grey scale is $\mu = 0.055$ to 0.018 mm^{-1} where white is the maximum and black is the minimum. (b) contains two pictures of the mouse breast tumour embedded in paraffin, on the left is the front face and the right is the rear side; these are raw outputs with matched levels in all colour channels with a camera model (Sony SLT- $\alpha 58$) and ISO specific denoising filter.

6. SPECTRAL X-RAY CT APPLIED: RE-EXCISION AFTER POSITIVE MARGIN

Measurement details relating to figure 6.3a follow:

- Peak x-ray tube voltage = 30 kV;
- X-ray tube current = 200 μ A;
- Source to detector distance = 275 ± 5 mm;
- Source to object distance = 110 ± 5 mm;
- Effective magnification $\simeq 3.3$;
- Reconstructed voxel size $\simeq 33$ μ m;
- Number of projections = 600;
- Flat field correction exposure time = 200 seconds;
- Mean photon counts in the flat field at the lowest threshold $\simeq 5 \times 10^6$ counts;
- Projection exposure time = 10 seconds;
- Mean photon counts in the object at the lowest threshold $\simeq 6 \times 10^4$ — 1.5×10^5 counts;
- Total exposure time = 200s + (600 \times 10s) = 6200s = 1.7 hours.

The CT slice intensity scale is proportional to density per voxel in figure 6.3a. The higher the density of the contents of that voxel are, the more white they are. A voxel is a value on a regular grid in three-dimensional space, a three dimensional pixel. The values for the CT reconstruction are consistent with the NIST x-ray attenuation coefficient tables for the closest tissue types.

In terms of aligning the images to a common coordinate system (registering), this was performed by Dr Jasper Nijkamp by using the python interface of ‘SimpleElastix’.

The main deviations between scans were due to the skin layer which became partially detached from the tumour during the slicing procedure. A significant fraction of slices were damaged by rips and tears in the centre of the sample and were not useful for registration with the CT data. Another issue with registration was that the CT slices cover the full three dimensional volume of the tumour whereas the H&E scans only cover thin slices (in the order of 5—10 μ m) with 19 slices over the entire tumour. Of those 19 slices, not all were even considered usable for registration purposes.

NDPITools C software and their ImageJ plugins were used to extract the H&E stained slices into TIFF (Tag Image File Format) files and work with them (90). NDPI is a pathology slide specimen specific file format which is similar to the TIFF file but supports significantly larger dimensions.

The tumour is outlined along with the other tissue types. In this particular case, almost the entire sample is tumour tissue with a thin layer of skin with hair follicles within the dermis. As expected, the non-necrotic tumour tissue is centred around the blood vessels which can be observed by the red colour of the remaining blood cells in the H&E scans. The tumour outline is the entire shape of what we would normally call the tumour. However, this specific tumour is very aggressive. There are therefore regions with active dividing tumour cells and necrotic areas that have grown beyond their blood supply. Therefore, the green tumour labels indicate active dividing tumour areas.

This scan does not give information about the differences between non-tumour breast tissue and tumour breast tissue. This is due to the lack of non-tumour breast tissue in this sample. It does however, show the differences between necrotic and non-necrotic tumour tissue. It can be observed that the non-necrotic tissue (most purple in the H&E scan, labelled as green) has a higher reconstructed linear attenuation coefficient (more white) in the CT slice compared to the necrotic tissue. Per unit volume, necrotic tissue should, by definition, have fewer cells due to a restricted blood and oxygen supply. They should also be the same material as the non-necrotic tumour tissue; leading to a lower linear attenuation coefficient (more black) than non-necrotic tissue. This is therefore consistent with expectations.

The hair follicles are the dots along the upper edge of the skin and are visible in both scans; they are dark purple in the H&E scan and medium grey in the CT slice. The particularly high intensity white regions on the extreme right and bottom of the CT slice contain a high number of non-tumour cells and blood vessels; these are areas of inflammation and have appeared as lesions in other similar tissue samples. These skin injuries are hypothesised to be a result of the tumour pushing the skin which is subsequently damaged when the mouse walked around.

In terms of achieving the best image quality within a 10 minute period, several methods were attempted. Starting from raw images, the FFCs improve on that. FFCs are the most commonly used method for image correction for most x-ray imaging applications. The energy calibration per pixel with a flat field correction (ECI_FFC) builds on that by correcting for inter-pixel variations using the spectral information and per pixel energy calibrations. Software-based post-processing in 'sinogram space' provides a significant boost to reconstruction quality. In sinogram space, certain patterns and spatial frequency components cannot exist. Filtering and correction can act based on this knowledge to suppress any remaining artefacts, systematic or random, from any source.


Another improvement can be achieved by oversampling - the imaged object is translated in steps of several pixels perpendicular to the beam, registered and averaged over the steps. This increases the signal to noise ratio at the cost of longer measurement times; this method is beneficial for particularly low contrast objects and for adaptive imaging modalities. There is a small penalty to the overall field of view due to the translation. The following is a curiosity driven example of adaptive imaging in this context: if one imaged an entire object, found something interesting that required small voxel sizes, zooms in, acquires another scan and overlays the new scan on the original. These scans can be combined to enable higher spatial resolution in regions of interest. Another advantage of this imaging approach would be to decrease the time for an initial view of an object and then periodically improve the reconstruction by re-scanning. This could be interesting in time critical situations where a lower quality measurement sooner adds value. If there is sufficient contrast to make a decision, why wait longer?

A qualitative example of this is shown in figure 6.4. It shows the result of a scan with two minutes of exposure time along with the 10 minute scan. The 10 minute scan is the mean value of the CT reconstructions of five sequential measurements with the same configuration except for a translation of 100 μm per measurement.


6. SPECTRAL X-RAY CT APPLIED: RE-EXCISION AFTER POSITIVE MARGIN

The measurement details are the same as before except for the following items:

- Flat field correction exposure time = 4 seconds;
- Mean photon counts in the flat field at the lowest threshold $\simeq 10^5$ counts;
- Projection exposure time = 0.2 seconds;
- Mean photon counts in the object at the lowest threshold $\simeq 10^3 - 3 \times 10^3$ counts;
- Total exposure time = $4 \text{ s} + (600 \text{ projections} \times 0.2 \text{ s}) = 124 \text{ s} = 2.1 \text{ minutes}$.



(a) A slice from a single two minute scan. This is relatively noisy compared to (b). The faint horizontal lines across the centre of the tumour bulk are ring artefacts.




(b) A slice with the mean values of the CT reconstructions of five sequential measurements with the same configuration. The total exposure time was 10 minutes for this.


Figure 6.4: Two slices of a CT reconstruction from the ECI_FFC reconstructed at 6 keV made using the Medipix3. The grey scale is $\mu = 0.055$ to 0.018 mm^{-1} . Note the 90° rotation relative to previous figures.

As one would expect, a measurement with one to two orders of magnitude lower photon counts in the raw images has a lower signal-to-noise ratio. However the quality of the two minute scan is qualitatively similar to the measurement lasting 1.7 hours. In terms of distinguishability (not a detailed statistics analysis), it is sufficient to discriminate between non-necrotic and necrotic tumour tissue. This is the metric by which success is judged for this problem - either these differences are visible or they are not. The primary cause of insufficient distinguishability is very likely due to insufficient photon statistics.

There are ring artefacts in the two minute scan that could not be completely suppressed which appear as faint horizontal lines in the centre of the tumour bulk along the entire height of the reconstruction. The 10 minute scan, however, appears to be a reasonable compromise between measurement time and image quality. Fortunately, this happens to be the approximate targeted measurement time mentioned before any measurements started.



(a) Classical x-ray CT, no energy information.



(b) Spectral x-ray CT, using energy bins for colour information.

Figure 6.5: Volume renderings of a CT reconstruction of the c7006-25 tumour. This contains the full 3D information with more than 50 reconstructed CT slices and can be viewed at any angle or sliced in any axis. The grey colour scale is $\mu = 0.055$ to 0.018 mm^{-1} . The colour (red, green, blue) channels are derived from the normalised energy bins from a single energy sensitive CT scan. Red corresponds to an energy bin of 6–10 keV, green is 10–21 keV and blue is 21+ keV. The total irradiation time of the sample was 43 minutes plus one minute for the flat field measurement. The x-ray tube peak voltage was set to 30 kVp and the current was 200 μA .

Figure 6.5 shows a volume rendering using the energy sensitivity for the colour channels. Each colour channel (red, green and blue) corresponds to an energy bin from one energy value to a higher energy value. Each colour channel is normalised to the maximum within that channel. The bulk of the tumour tissue does not show a variation in terms of the energy above the apparent noise levels. The areas of inflammation on the surface of the skin are the white dots on the left side. These are surrounded by a pink-red area which means that the linear attenuation coefficient is relatively higher in the first energy bin than the rest of the skin which is green-blue. It is hypothesised that this pink-red area is showing an inflammatory response over a wider area.

6.4 SUMMARY

Based on the distinguishability between the necrotic and non-necrotic tumour tissue, the 1.7 hour and 43 minute scan both appear to have sufficiently low statistical uncertainties and the two minute scan appears to not have enough photons. If one is aiming to distinguish necrotic from non-necrotic tumour tissue, the optimum measurement settings for this sample with this system configuration appears to be somewhere between the two extremes. Optimising this would require a scan of different measurement times for both the flat field and object exposure times. With higher intensity x-ray tubes such as those using liquid metal jets, it would be possible to achieve similar statistical uncertainties as the 1.7 hour scan in significantly shorter times.

6.5 OUTLOOK

Another avenue of research would be to measure the spectral properties (x-ray mass attenuation coefficients) of various tissues in absolute terms. The following tissues would be relevant: muscle, fat, connective tissues and tumour (necrotic and non-necrotic). This could be achieved by high resolution threshold scans of tissue samples and calculating the relative linear attenuation coefficients as functions of energy as discussed in section 5.5.4. Cross-checking these resulting distributions with high energy resolution x-ray spectrometers with a high quality radiation source would provide a ‘ground-truth’ to compare against.

Other avenues of further development would be the visualisation and interpretation of the extra dimension of energy in the CT reconstructions. Also, what is the added value to physicians? This question can be answered in a medical context.

This work was used as preliminary evidence as part of a successful funding application made by Associate Professor Jasper Nijkamp to the Danish Cancer Society for a project titled ‘Spectral micro-CT imaging: online evaluation of tumor surgery’ (1,915,000 DKK, R279-A16481, 2019) and to the Aarhus University Research Fund for a project titled ‘Purify the gold standard of tumor segmentation by standard tumor imaging with a spectral microCT scanner at pathology’ (2,500,000 DKK, AUFF-E-2021-9-15, 2021) in Denmark. Further studies in this direction are being continued by myself, Dr Jasper Nijkamp and the rest of the team at the Department of Clinical Medicine at Aarhus University in Denmark, also using a Medipix3 based detector.

6. SPECTRAL X-RAY CT APPLIED: RE-EXCISION AFTER POSITIVE MARGIN

We're very good at answering questions that aren't asked.

Martin Fransen, Aug. 2019

7

Summary and outlook

7

7.1 EXECUTIVE SUMMARY

This work can be summarised as characterising and optimising a detector and showing its added value for several applications.

An overview list of the achievements within this work follows:

- Noise and contrast performance;
- Energy calibration;
- Use and performance in extreme conditions - clinical particle beams;
- Performance in spectral x-ray imaging;
- Setting up and optimising a data acquisition procedure and post processing methods for optimisation of spectral data;
- Spectral CT scans, focusing at identifying tumour tissue to reduce surgical margin.

7.2 SUMMARY

The work that went into producing this thesis is split into three distinct parts.

I

The first part was the development and commercialisation of the detector control and readout software, Dexter, in combination with the characterisation of the Medipix3. At the beginning of my project, Dexter was in an alpha state, many critical features were missing or incomplete but the core of the software was present and working. For this research, several new modes of operations were developed to commercial standard including threshold scanning, CT scanning and energy calibration integrations. These added new capabilities, made existing features faster and more robust in many aspects. At the time of writing, it is stable and fully featured for my purposes. With sufficient hardware, it can handle the Medipix3 detector at the maximum supported data rate, 10 Gbit/s. Significant time was spent learning about the Medipix3 and the correct methods to control it alongside the development of the firmware on the FPGA in the SPIDR3 system. Conversations with the chip designers at CERN were crucial for understanding the chip as it was intended to behave. Studies into signal to noise ratios, system stability and energy resolution directly support the x-ray imaging and CT scanning work.

2

The second part was the work to investigate the extremes of the Medipix3 chip in hadron therapy environments. Due to the nature of the work, significant collaboration within the OMA (Optimization of Medical Accelerators) network was critical for any of these experiments. This pushed the verified operational limits of the chip with protons, carbon ions and heavier ions. A limit was identified where the chip encounters the ‘volcano effect’ at a particular energy deposition. The chip’s count rate linearity, dynamic range and minimum sensitivity was found to be favourable for both protons and carbon ions compared to current detection methods, EBT3 film and gas chambers. Radiation damage of the chips was observed. However, it self-recovered over time, apparently back to the pre-irradiated state. Moving onto the accelerators, there were frequency components in the mean counts over time that were detected in both sets of measurements, one confirmed to be due to spill ripples of the synchrotron. These efforts resulted in three published manuscripts.

3

The third and final part is the spectral x-ray imaging and computed tomography. This work spans the entire range from the detector characterisation all the way through to the spectral x-ray CT imaging of the mouse breast tumour. It depends on all of the previous non-hadron therapy work combined. My efforts were also part of the FleX-ray project, resulting in a new CT scanning facility at the Computational Imaging group within CWI. Per pixel energy calibrations were the answer to improving image quality by addressing a specific error source while maintaining the same units and energy dependence. This is in contrast to beam hardening corrections which result in equivalent thicknesses of the correction material and in principle, eliminate energy dependent information as it is all mapped to thicknesses of a material. It is important to note that for the best x-ray imaging, beam hardening corrections are often one of the best options available. It was found that charge summing mode (CSM) indeed benefits higher Z sensor materials the most, where the charge cloud ends up in the order of the pixel pitch size by the time it reaches the chip. Due to the combination of the thickness of the silicon used and the degraded energy resolution, it was not deemed that CSM was useful for spectral x-ray imaging of biological samples. Comparisons were made between a dedicated x-ray spectrometer (silicon drift detector) and the Medipix3 along with comparisons of different gain modes and between SPM and CSM with a GaAs sensor. These were found to be consistent with expectations. Onto the challenge of imaging biological tissue, using experimentally derived values, the x-ray attenuation coefficients as functions of energy are shown to overlap for common soft tissue types. This demonstrates how, even with a perfect detector, it is not possible to unambiguously separate these soft tissues from each other using the linear attenuation coefficients alone, more so at higher energies. The signal to noise ratio as a function of energy was investigated for four combinations of detector configuration, with Si and GaAs, 110 and 55 μm pixel pitches, between SPM and CSM and with low (30 keV) and high (90 keV) peak x-ray voltages. This showed how the energy calibration benefits the SNR over energy both in image space and in energy space. It also showed how the flat field correction must be tuned for each imaged sample for optimal SNR. Examples of the relative linear attenuation coefficients over energy were shown for objects with and without transition energies and for biological tissue, a calibration phantom (CDMAM) and pure metal foils.

7.3 OUTLOOK

Future studies could use Medipix4 in similar ways to Medipix3, ideally there would be no RTN with the new design. Timepix2 has been shown to be exceptionally stable over time with no evidence of RTN so far. Timepix4 could be useful for hyper spectral x-ray imaging and CT with cluster analysis at relatively low count rates. Liquid metal jet-based x-ray sources are a significant improvement over conventional x-ray sources in terms of increasing x-ray flux while maintaining x-ray spot size. Given that the Medipix4 should be capable of an even higher count rate than Medipix3, this combination may prove fruitful. Direct comparisons between the Medipix3 and other x-ray detectors such as flat panel detectors would be of interest also.

We now move onto possible improvements for the existing experimental setup at Nikhef. The rotation stage used in this work has insufficient torque for the load; it skips steps at a very low acceleration and rotational velocity as reported by software. A different motor would be best. When trying to achieve a fast scan, waiting on motors to accelerate, move and decelerate start to dominate scan times. Latency communicating with system components is a much smaller effect in this setup ($\sim 5 \text{ ms} \times 600 = \sim 3 \text{ seconds}$ per typical scan) compared to mechanical motor aspects. Therefore, continuous rotation scans are necessary to cut down on scan times with a penalty of small motion artefacts, depending on the relative motion of the object relative to the frame rate. This is standard for medical systems which are able to do very fast full body scans by utilising helical scans.

Metal artefact reduction using energy bins and tumour detection are both areas of research where photon counting detectors are finding their place. The metal artefact reduction is achieved by multi-energy imaging and then calculating the difference between those bins. These artefacts are suppressed at higher energies because there is lower relative difference between the higher density and lower density parts of the object scanned (91).

The highest quality x-ray CT scans all use synchrotron radiation and phase-contrast imaging. They benefit from effective spot sizes in the order of nanometres, coherent radiation and the phase-contrast effect which boosts the contrast of low attenuation objects. This is showcased in a paper where whole human organs were non-destructively scanned down to cellular resolution ($1.3\text{--}2.5 \mu\text{m}$ per voxel) for COVID-19 patients (92). This method is among the best available in terms of ultimate scan quality, however, it requires significant expertise and access to the most expensive radiation source type (synchrotrons) in existence. Advances in accessibility to high quality radiation sources and detector technology are crucial for significant improvements in this field. As detailed in the 'Plasma Wakefield Accelerator Research 2019 - 2040: A community-driven UK roadmap compiled by the Plasma Wakefield Accelerator Steering Committee (PWASC)' (93), this is likely to come to fruition in the form of a compact synchrotron radiation source driven by a laser-plasma wakefield accelerator.

This project starts the long path towards making surgical removal of tumours significantly more efficient. Given the high and still growing number of people needing this type of surgery, many lives could be saved due to this work. I am now working on that future as a post-doctoral researcher at

Aarhus University Hospital in Denmark.

The future is small, bright and colourful!

Bibliography

- [1] Hannah Ritchie and Max Roser. Causes of death. *Our World in Data*, 2018. <https://ourworldindata.org/causes-of-death>.
- [2] IQVIA Inc. Global oncology trends 2021, 2021. <https://www.iqvia.com/insights/the-iqvia-institute/reports/global-oncology-trends-2021>.
- [3] W. C. Röntgen. On a new kind of rays, 1896.
- [4] Nicholas Rivera and Ido Kaminer. Light–matter interactions with photonic quasiparticles. *Nature Reviews Physics*, 2(10):538–561, Oct 2020.
- [5] Tom Schoonjans, Antonio Brunetti, Bruno Golosio, Manuel Sanchez Del Rio, Vicente Armando Solé, Claudio Ferrero, and Laszlo Vincze. The xraylib library for X-ray-matter interactions. Recent developments. *Spectrochimica Acta - Part B Atomic Spectroscopy*, 66(11-12):776–784, nov 2011.
- [6] D. Coster and R. De L. Kronig. New type of auger effect and its influence on the x-ray spectrum. *Physica*, 2(1-12):13–24, jan 1935.
- [7] Se Hee Lee, Won Seok Chang, Sung Mi Han, Duck Hyun Kim, and Jong Ki Kim. Synchrotron x-ray nanotomography and three-dimensional nanoscale imaging analysis of pore structure-function in nanoporous polymeric membranes. *Journal of Membrane Science*, 535:28–34, 8 2017.
- [8] Florian T. Gassert, Theresa Urban, Manuela Frank, Konstantin Willer, Wolfgang Noichl, Philipp Buchberger, Rafael Schick, Thomas Koehler, Jens von Berg, Alexander A. Fingerle, Andreas P. Sauter, Marcus R. Makowski, Daniela Pfeiffer, and Franz Pfeiffer. X-ray dark-field chest imaging: Qualitative and quantitative results in healthy humans. *Radiology*, 301(2):389–395, 2021. PMID: 34427464.
- [9] Konstantin Willer, Alexander A Fingerle, Wolfgang Noichl, Fabio De Marco, Manuela Frank, Theresa Urban, Rafael Schick, Alex Gustschin, Bernhard Gleich, Julia Herzen, Thomas Koehler, Andre Yaroshenko, Thomas Pralow, Gregor S Zimmermann, Bernhard Renger, Andreas P Sauter, Daniela Pfeiffer, Marcus R Makowski, Ernst J Rummeny, Philippe A Grenier, and Franz Pfeiffer. X-ray dark-field chest imaging for detection and quantification of emphysema in patients with chronic obstructive pulmonary disease: a diagnostic accuracy study. *The Lancet Digital Health*, 3:e733–e744, 11 2021.
- [10] P. Zhu, P. Duvauchelle, G. Peix, and D. Babot. X-ray compton backscattering techniques for process tomography: imaging and characterization of materials. *Measurement Science and Technology*, 7:281, 3 1996.
- [11] Peter Rez, Robert L. Metzger, and Kenneth L. Mossman. The dose from compton backscatter screening. *Radiation Protection Dosimetry*, 145:75–81, 4 2011.
- [12] M. Margret, M. Menaka, B. Venkatraman, and S. Chandrasekaran. Compton back scatter imaging for mild steel rebar detection and depth characterization embedded in concrete. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 343:77–82, 1 2015.

- [13] H Bethe. Bremsformel für elektronen relativistischer geschwindigkeit. *Zeitschrift für Physik*, 76:293–299, 1932.
- [14] B. Mustapha, B. Aydogan, J. Nolen, A. Nassiri, J. Noonan, M. Pankuch, J. Welsh, R. Schulte, and J. Robb. Prospects for an advanced heavy ion therapy center in the Chicago area. *AIP Conference Proceedings*, 2160:050009, 10 2019.
- [15] W. Adam, W. de Boer, E. Borch, M. Bruzzi, C. Colledani, P. D’Angelo, V. Dabrowski, W. Dulinski, B. van Eijk, V. Eremin, F. Fizzotti, H. Fraai-Köbl, C. Furetta, K. K. Gan, A. Gorisek, E. Griesmayer, E. Grigoriev, F. Hartjes, J. Hrubec, F. Huegging, H. Kagan, J. Kaplon, R. Kass, K. T. Knöpfle, M. Krammer, W. Lange, A. Logiudice, C. Manfredotti, M. Mathes, D. Menichelli, M. Mishina, L. Moroni, J. Noomen, A. Oh, H. Pernegger, M. Pernicka, R. Potenza, J. L. Riest, A. Rudge, S. Sala, S. Schnetzer, S. Sciortino, R. Stone, C. Suter, W. Trischuk, J. J. Velthuis, B. Vincenzo, P. Weilhammer, J. Weingarten, N. Wermes, and W. Zeuner. Radiation hard diamond sensors for future tracking applications. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2006.
- [16] R Plackett, X Llopart, R Ballabriga, M Campbell, L Tlustos, W Wong, and Richard Plackett@cern Ch. Measurement of radiation damage to 130nm hybrid pixel detector readout chips. *Proceedings of the Topical Workshop on Electronics for Particle Physics*, pages 157–160, 2009.
- [17] Martin Kroupa, Amir Bahadori, Thomas Campbell-Ricketts, Anton Empl, Son Minh Hoang, John Idarraga-Munoz, Ryan Rios, Edward Semones, Nicholas Stoffle, Lukas Tlustos, Daniel Turecek, and Lawrence Pinsky. A semiconductor radiation imaging pixel detector for space radiation dosimetry. *Life Sciences in Space Research*, 6:69–78, 2015.
- [18] Jieh-Wen Tsung. *Diamond and Silicon Pixel Detectors in High Radiation Environments*. PhD thesis, University of Bonn, 2012. <https://cds.cern.ch/record/1540669/files/CERN-THESIS-2012-302.pdf>.
- [19] Lawrence Pinsky, Son Hoang, John Idarraga-Muñoz, Martin Kroupa, Nicholas Stoffle, Amir Bahadori, Edward Semones, Hisashi Kitamura, Satoshi Kodaira, Jan Jakubek, Zdenek Vykydal, Daniel Turecek, and Stanislav Pospisil. Summary of the first year of Medipix-based space radiation monitors on the ISS. In *IEEE Aerospace Conference Proceedings*, pages 1–8, 03 2014.
- [20] Mark Storey. Proton radiation oncology, pll - protons 101 - where dose goes. integral dose, 2 2022. <https://www.protonradiationoncology.com/protons101/2019/3/8/intro-to-protons-101-where-dose-goes>.
- [21] Wayne D Newhauser and Rui Zhang. The physics of proton therapy. *Physics in Medicine & Biology*, 60:R155, 3 2015.
- [22] KT CERN. Success stories - medaustro, 2022. <https://kt.cern/success-stories/medaustro>.
- [23] Rafael Ballabriga-Suñé. *The Design and Implementation in 0.13 micrometer CMOS of an Algorithm Permitting Spectroscopic Imaging with High Spatial Resolution for Hybrid Pixel Detectors*. PhD thesis, Universitat Ramon Llull, Barcelona, Spain, 2009.

-
- [24] Rafael Ballabriga-Suñé and Xavier Llopart. *Medipix3 RX manual*. CERN, Medipix Group, 2012.
 - [25] Eva N Gimenez, Rafael Ballabriga-Suñé, Michael Campbell, Ian Horswell, Xavier Llopart, Julien Marchal, Kawal J S Sawhney, Nicola Tartoni, and Daniel Turecek. Characterization of Medipix3 with synchrotron radiation. *IEEE Transactions on Nuclear Science*, 58:323, 2011.
 - [26] Patrick Takoukam Talla. *Investigation of photon counting pixel detectors for X-ray spectroscopy and imaging*. PhD thesis, Der Naturwissenschaftlichen Fakultät der Friedrich-Alexander-Universität Erlangen-Nürnberg, 2011.
 - [27] Elias Hamann. *Characterization of High Resistivity GaAs as Sensor Material for Photon Counting Semiconductor Pixel Detectors*. PhD thesis, Albert-Ludwigs-Universität Freiburg im Breisgau, 11 2013.
 - [28] Helmuth Spieler. *Semiconductor Detector Systems*, volume 9780198527848. Oxford University Press, 9 2007.
 - [29] U. Fano. Ionization yield of radiations. ii. the fluctuations of the number of ions. *Physical Review*, 72:26–29, 7 1947.
 - [30] M. N. Mazziotta. Electron-hole pair creation energy and fano factor temperature dependence in silicon. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 584:436–439, 1 2008.
 - [31] M Barbero, V Bellini, V Belyaev, E Berdermann, P Bergonzo, H Bol, M Bruzzi, V Cindro, W De Boer, I Dolenc, P Dong, W Dulinski, V Ermin, R Eusebi, F Fizzotti, H Frais-Kölbl, A Furgeri, K K Gan, A Golubev, A Gorisek, E Griesmayer, E Grigoriev, F Hartjes, H Kagan, R Kass, G Kramberger, S Kuleshov, S Lagomkarsino, A Lo Giudice, I Mandic, C Manfredotti, A Martemyanov, M Mathes, D Menichelli, S Miglio, M Mikuz, M Mishina, S Mueller, H Pernegger, M Pomorski, R Potenza, S Roe, C Schmidt, S Schnetzer, T Schreiner, C Schrupp, S Sciortino, S Smith, R Stone, C Sutura, M Traeger, W Trischuk, C Tuve, J Velthuis, E Vittone, R Wallny, P Weilhammer, and N Wermes. Development of diamond tracking detectors for high luminosity experiments at the lhc the RD42 collaboration, 1 2007. <http://cds.cern.ch/record/1009654?ln=en>.
 - [32] G. Claps, F. Murtas, L. Foggetta, C. Di Giulio, J. Alozy, and G. Cavoto. Diamondpix: A CVD diamond detector with Timepix3 chip interface. *IEEE Transactions on Nuclear Science*, 2018.
 - [33] Lukas Tlustos, Georgy Shelkov, and Oleg P. Tolbanov. Characterisation of a GaAs(Cr) Medipix2 hybrid pixel detector. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 633:S103–S107, 5 2011.
 - [34] M. C. Veale, S. J. Bell, D. D. Duarte, M. J. French, A. Schneider, P. Seller, M. D. Wilson, A. D. Lozinskaya, V. A. Novikov, O. P. Tolbanov, A. Tyazhev, and A. N. Zarubin. Chromium compensated gallium arsenide detectors for x-ray and γ -ray spectroscopic imaging. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 752:6–14, 7 2014.

- [35] Arthur E. Burgess. The rose model, revisited. *Journal of the Optical Society of America A*, 16:633, 3 1999.
- [36] E. Fröjd, R. Ballabriga-Suñé, M. Campbell, M. Fiederle, E. Hamann, T. Koenig, X. Llopart, D. De Paiva Magalhaes, and M. Zuber. Count rate linearity and spectral response of the Medipix3RX chip coupled to a 300 μm silicon sensor under high flux conditions. *Journal of Instrumentation*, 9:C04028, 4 2014.
- [37] E. Hamann, T. Koenig, M. Zuber, A. Cecilia, A. Tyazhev, O. Tolbanov, S. Procz, A. Fauler, M. Fiederle, and T. Baumbach. Investigation of GaAs:Cr Timepix assemblies under high flux irradiation. *Journal of Instrumentation*, 10:C01047, 1 2015.
- [38] Raj Kumar Panta, Michael F. Walsh, Stephen T. Bell, Nigel G. Anderson, Anthony P. Butler, and Philip H. Butler. Energy calibration of the pixels of spectral X-ray detectors. *IEEE Transactions on Medical Imaging*, 34:697–706, 3 2015.
- [39] K. Heijhoff, K. Akiba, M. van Beuzekom, P. Bosch, J. Buytaert, M. Campbell, A. P. Colijn, P. Collins, E. Dall’Oco, T. Evans, R. Geertsema, M. L. E. Heidotting, D. Hynds, X. Llopart Cudie, H. Schindler, and H. Snoek. Timing performance of the LHCb VELO Timepix3 Telescope. *Journal of Instrumentation*, 15, 8 2020.
- [40] Enrico Schioppa, Jan Visser, and Els Koffeman. Prospects for spectral ct with Medipix detectors. *Proceedings of Science*, 0:246, 7 2014.
- [41] Rafael Ballabriga-Suñé, Michael Campbell, and Xavier Llopart. ASIC developments for radiation imaging applications: The Medipix and Timepix family. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 878:10–23, 1 2018.
- [42] Jacinta Yap. *Characterisation Studies of Proton Beamlines for Medical Applications and Beam Diagnostics Integration*. PhD thesis, University of Liverpool, 3 2021.
- [43] Anna Baratto-Roldán, María del Carmen Jiménez-Ramos, Maria Cristina Battaglia, Javier García-López, María Isabel Gallardo, Miguel A. Cortés-Giraldo, and José M. Espino. Feasibility study of a proton irradiation facility for radiobiological measurements at an 18 MeV cyclotron. *Instruments*, 2(4), 2018.
- [44] Giuseppe Battistoni, Till Boehlen, Francesco Cerutti, Pik Wai Chin, Luigi Salvatore Esposito, Alberto Fassò, Alfredo Ferrari, Anton Lechner, Anton Empl, Andrea Mairani, Alessio Mereghetti, Pablo Garcia Ortega, Johannes Ranft, Stefan Roesler, Paola R. Sala, Vasilis Vlachoudis, and George Smirnov. Overview of the FLUKA code. *Annals of Nuclear Energy*, 82:10–18, 8 2015.
- [45] T. T. Böhlen, F. Cerutti, M. P.W. Chin, A. Fassò, A. Ferrari, P. G. Ortega, A. Mairani, P. R. Sala, G. Smirnov, and V. Vlachoudis. The FLUKA code: Developments and challenges for high energy and medical applications. *Nuclear Data Sheets*, 120:211–214, 6 2014.
- [46] Akifumi Itano. Heavy ion medical accelerator in Chiba (HIMAC). *Hoshasen Seibutsu Kenkyu*, 30(1):1–17, 1995.
- [47] S Hoang, R Vilalta, L Pinsky, M Kroupa, N Stoffle, and J Idarraga-Muñoz. Data analysis of tracks of heavy ion particles in timepix detector. *Journal of Physics: Conference Series*, 523:012026, jun 2014.

- [48] S P George, M Kroupa, S Wheeler, S Kodaira, H Kitamura, L Tlustos, T Campbell-Ricketts, N N Stoffle, E Semones, and L Pinsky. Very high energy calibration of silicon Timepix detectors. *Journal of Instrumentation*, 13(11):P11014–P11014, nov 2018.
- [49] X. Llopart, R. Ballabriga-Suñé, M. Campbell, L. Tlustos, and W. Wong. Erratum to “Timepix, a 65k programmable pixel readout chip for arrival time, energy and/or photon counting measurements” [nucl. instr. and meth. a. 581 (2007) 485–494]. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 585(1):106 – 108, 2008.
- [50] S P George, M Kroupa, S Wheeler, S Kodaira, H Kitamura, L Tlustos, T Campbell-Ricketts, N N Stoffle, E Semones, and L Pinsky. Very high energy calibration of silicon Timepix detectors. *Journal of Instrumentation*, 13, 9 2018.
- [51] T E J Campbell-Ricketts, M Kroupa, and L S Pinsky. Spectroscopy of high-energy ions with Timepix3. *Journal of Instrumentation*, 11, 2016.
- [52] L. Pinsky, J Idarrága-Muñoz, Martin Kroupa, H Son, Nicholas Stoffle, E Semones, Amir Bahadori, Daniel Turecek, S Il, Jan Jakubek, Zdenek Vykydal, H Kitamura, and Y Uchihori. Medipix in space on-board the ISS. *Journal of radiation research*, 55:i62–i63, 03 2014.
- [53] Radon Transform - MATLAB & Simulink - MathWorks, 2022. <https://uk.mathworks.com/help/images/radon-transform.html>.
- [54] Avinash C. Kak and Malcolm Slaney. Principles of computerized tomographic imaging. *Principles of Computerized Tomographic Imaging*, 1 2001.
- [55] Enrico Junior Schioppa. *The color of X-rays. Spectral X-ray computed tomography using energy sensitive pixel detectors*. PhD thesis, Universiteit van Amsterdam, 2014.
- [56] Wim van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K. Joost Batenburg, and Jan Sijbers. The ASTRA toolbox: A platform for advanced algorithm development in electron tomography. *Ultramicroscopy*, 157:35–47, 2015.
- [57] Wim van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabrovolski, Jan De Beenhouwer, K. Joost Batenburg, and Jan Sijbers. Fast and flexible x-ray tomography using the ASTRA toolbox. *Opt. Express*, 24(22):25129–25147, Oct 2016.
- [58] Hyoung Suk Park, Yong Eun Chung, and Jin Keun Seo. Computed tomographic beam-hardening artefacts: mathematical characterization and analysis. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 373(2043):20140388, Jun 2015. 25939628[pmid].
- [59] William Vågberg, Jakob C. Larsson, and Hans M. Hertz. Removal of ring artifacts in microtomography by characterization of scintillator variations. *Optics Express*, 25:23191, 9 2017.
- [60] Vytautas Astromskas, Eva N. Gimenez, Annika Lohstroh, and Nicola Tartoni. Evaluation of polarization effects of e- collection schottky cdte medipix3rx hybrid pixel detector. *IEEE Transactions on Nuclear Science*, 63:252–258, 2 2016.
- [61] J.H. Hubbell and S.M. Seltzer. NIST: X-Ray Mass Attenuation Coefficients, 1996. <http://www.nist.gov/pml/data/xraycoef/>.
- [62] ITIS tissue density database, 2021. <https://itis.swiss/virtual-population/tissue-properties/database/density/>.

- [63] Celia J. Strudley and Kenneth C. Young. Evaluation of a new design of contrast-detail phantom for mammography: CDMAM model 4.0. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8539 LNCS:217–224, 2014.
- [64] Directorate-General for Health and Food Safety (European Commission). European guidelines for quality assurance in breast cancer screening and diagnosis - publications office of the eu, 10 2013. <https://op.europa.eu/en/publication-detail/-/publication/4e74ee9b-df80-4c91-a5fb-85efb0fdd42b/language-en>.
- [65] Mirion Technologies (Capintec), Inc. Cd phantom for mammography (cdmam 3.4) | mirion technologies (capintec), inc., 2021. <https://capintec.com/product/cd-phantom-for-mammography-cdmam-3-4/>.
- [66] A. Atharifarid, J. L. Healy, B. P. Goulter, M. Ramyar, L. Vanden Broeke, M. F. Walsh, C. C. Onyema, R. K. Panta, R. Aamir, D. J. Smithies, R. Doesburg, M. Anjomrouz, M. Shamshad, S. Bheesette, K. Rajendran, N. J.A. De Ruiter, D. Knight, A. Chernoglazov, H. Mandalika, S. T. Bell, C. J. Bateman, A. P.H. Butler, and P. H. Butler. Per-pixel energy calibration of photon counting detectors. *Journal of Instrumentation*, 12, 3 2017.
- [67] M. Jake Pushie, Ingrid J. Pickering, Malgorzata Korbas, Mark J. Hackett, and Graham N. George. Elemental and Chemically Specific X-ray Fluorescence Imaging of Biological Systems. *Chemical Reviews*, 114:8499–8541, 9 2014.
- [68] Darin P. Clark and Cristian T. Badea. Hybrid spectral CT reconstruction. *PLoS ONE*, 12, 2017.
- [69] James R. Bennett, Alex M.T. Opie, Qiong Xu, Hengyong Yu, Michael Walsh, Anthony Butler, Phillip Butler, Guohua Cao, Aaron Mohs, and Ge Wang. Hybrid spectral micro-CT: System design, implementation, and preliminary results. *IEEE Transactions on Biomedical Engineering*, 61:246–253, 2014.
- [70] Kishore Rajendran, Martin Petersilka, André Henning, Elisabeth R. Shanblatt, Bernhard Schmidt, Thomas G. Flohr, Andrea Ferrero, Francis Baffour, Felix E. Diehn, Lifeng Yu, Prabhakar Rajiah, Joel G. Fletcher, Shuai Leng, and Cynthia H. McCollough. First clinical photon-counting detector CT system: technical evaluation. *Radiology*, 12 2021.
- [71] Peng He, Biao Wei, Wenxiang Cong, and Ge Wang. Optimization of K-edge imaging with spectral CT. *Medical Physics*, 39:6572–6579, 2012.
- [72] Polad M. Shikhaliev and Shannon G. Fritz. Photon counting spectral CT versus conventional CT: Comparative evaluation for breast imaging application. *Physics in Medicine and Biology*, 56:1905–1930, 4 2011.
- [73] L. Vanden Broeke, M. Grillon, A. W.K. Yeung, W. Wu, R. Tanaka, and V. Vardhanabhuti. Feasibility of photon-counting spectral CT in dental applications—a comparative qualitative analysis. *BDJ Open*, 7:1–8, 12 2021.
- [74] David P. Cormode, Salim Si-Mohamed, Daniel Bar-Ness, Monica Sigovan, Pratap C. Naha, Joelle Balegamire, Franck Lavenne, Philippe Coulon, Ewald Roessl, Matthias Bartels, Michal Rokni, Ira Blevis, Loic Boussel, and Philippe Douek. Multicolor spectral photon-counting computed tomography: In vivo dual contrast imaging with a high count rate scanner. *Scientific Reports*, 7, 2017.

- [75] World Health Organization. Breast cancer fact sheet, 2021. <https://www.who.int/news-room/fact-sheets/detail/breast-cancer>.
- [76] Kelsey M. Kennedy, Renate Zilkens, Wes M. Allen, Ken Y. Foo, Qi Fang, Lixin Chin, Rowan W. Sanderson, James Anstie, Philip Wijesinghe, Andrea Curatolo, Hsien Ern I. Tan, Narelle Morin, Bindu Kunjuraman, Chris Yeomans, Synn Lynn Chin, Helen DeJong, Katharine Giles, Benjamin F. Dessauvagie, Bruce Latham, Christobel M. Saunders, and Brendan F. Kennedy. Diagnostic accuracy of quantitative micro-elastography for margin assessment in breast-conserving surgery. *Cancer Research*, 80(8):1773–1783, 2020.
- [77] Dalal Aziz, Ellen Rawlinson, Steven A Narod, Ping Sun, H Lavina A Lickley, David R McCready, and Claire MB Holloway. The role of reexcision for positive margins in optimizing local disease control after breast-conserving surgery for cancer. *The Breast Journal*, 12(4):331–337, 2006.
- [78] Dalal Aziz, Ellen Rawlinson, Steven A. Narod, Ping Sun, H. Lavina A. Lickley, David R. McCready, and Claire M. B. Holloway. The role of reexcision for positive margins in optimizing local disease control after breast-conserving surgery for cancer. *The Breast Journal*, 12(4):331–337, 2006.
- [79] Gary Freedman, Barbara Fowble, Alexandra Hanlon, Nicos Nicolaou, Douglas Fein, John Hoffman, Elin Sigurdson, Marcia Boraas, and Lori Goldstein. Patients with early stage invasive cancer with close or positive margins treated with conservative surgery and radiation have an increased risk of breast recurrence that is delayed by adjuvant systemic therapy. *International Journal of Radiation Oncology Biology Physics*, 44(5):1005–1015, 1999.
- [80] C. Cellini, S. T. Hollenbeck, P. Christos, D. Martins, J. Carson, S. Kemper, E. La Vigne, E. Chan, and R. Simmons. Factors associated with residual breast cancer after re-excision for close or positive margins. *Annals of Surgical Oncology* 2004 11:10, 11:915–920, 2004.
- [81] A.K. Talsma, A.M.J. Reedijk, R.A.M. Damhuis, P.J. Westenend, and W.J. Vles. Re-resection rates after breast-conserving surgery as a performance indicator: Introduction of a case-mix model to allow comparison between dutch hospitals. *European Journal of Surgical Oncology (EJSO)*, 37(4):357–363, 2011.
- [82] Nathaniel J. M. Haven, Kevan L. Bell, Pradyumna Kedarisetti, John D. Lewis, and Roger J. Zemp. Ultraviolet photoacoustic remote sensing microscopy. *Opt. Lett.*, 44(14):3586–3589, Jul 2019.
- [83] Nathaniel J. M. Haven, Pradyumna Kedarisetti, Brendon S. Restall, and Roger J. Zemp. Reflective objective-based ultraviolet photoacoustic remote sensing virtual histopathology. *Opt. Lett.*, 45(2):535–538, Jan 2020.
- [84] Terence T. W. Wong, Ruiying Zhang, Pengfei Hai, Chi Zhang, Miguel A. Pleitez, Rebecca L. Aft, Deborah V. Novack, and Lihong V. Wang. Fast label-free multilayered histology-like imaging of human breast cancer by photoacoustic microscopy. *Science Advances*, 3(5):e1602168, 2017.
- [85] Rowan W. Sanderson, Andrea Curatolo, Philip Wijesinghe, Lixin Chin, and Brendan F. Kennedy. Finger-mounted quantitative micro-elastography. *Biomed. Opt. Express*, 10(4):1760–1773, Apr 2019.

- [86] Sishan Cui, Shuo Zhang, and Shuhua Yue. Raman spectroscopy and imaging for cancer diagnosis. *Journal of Healthcare Engineering*, 2018, 2018.
- [87] Adukkadan N. Ramya, Jayadev S. Arya, Murali Madhukrishnan, Shanmughan Shamjith, Murukan S. Vidyalekshmi, and Kaustabh K. Maiti. Raman imaging: An impending approach towards cancer diagnosis. *Chemistry – An Asian Journal*, 16(5):409–422, 2021.
- [88] Andrew H. Hill, Bryce Manifold, and Dan Fu. Tissue imaging depth limit of stimulated raman scattering microscopy. *Biomedical Optics Express*, 11:762, 2 2020.
- [89] Irina Palimaru Manhoobi, Anne Bodilsen, Jasper Nijkamp, Anuj Pareek, Trine Tramm, Søren Redsted, and Peer Christiansen. Diagnostic accuracy of radiography, digital breast tomosynthesis, micro-CT and ultrasound for margin assessment during breast surgery: A systematic review and meta-analysis. *Academic Radiology*, 2022.
- [90] Christophe Deroulers, David Ameisen, Mathilde Badoual, Chloé Gerin, Alexandre Granier, and Marc Lartaud. Analyzing huge pathology images with open source software. *Diagnostic Pathology*, 8:1–8, 6 2013.
- [91] Chang-Lae ID Lee, Junyoung Park, Sangnam Nam, Jiyoung Choi, Yuna Choi, Sangmin Lee, Kyoung-Yong Lee, and Minkook Cho. Metal artifact reduction and tumor detection using photon-counting multi-energy computed tomography. *PLoS ONE*, 2020.
- [92] C. L. Walsh, P. Tafforeau, W. L. Wagner, D. J. Jafree, A. Bellier, C. Werlein, M. P. Kühnel, E. Boller, S. Walker-Samuel, J. L. Robertus, D. A. Long, J. Jacob, S. Marussi, E. Brown, N. Holroyd, D. D. Jonigk, M. Ackermann, and P. D. Lee. Imaging intact human organs with local resolution of cellular structures using hierarchical phase-contrast tomography. *Nature Methods* 2021, pages 1–10, 11 2021.
- [93] Bernhard Hidding, Simon Hooker, Steven Jamison, Bruno Muratori, Christopher Murphy, Zulfikar Najmudin, Rajeev Pattathil, Gianluca Sarri, Matthew Streeter, Carsten Welsch, Matthew Wing, and Guoxing Xia. Plasma wakefield accelerator research 2019–2040. Technical report, Plasma Wakefield Accelerator Steering Committee, 2019.
- [94] Eva N. Gimenez, Rafael Ballabriga-Suñé, Gabriel Blaj, Michael Campbell, Igor Dolbnya, Erik Fröjdj, Ian Horswell, Xavier Llopart, Julien Marchal, John McGrath, David Omar, Richard Plackett, Kawal Sawhney, and Nicola Tartoni. Medipix3RX: Characterizing the Medipix3 redesign with synchrotron radiation. *IEEE Transactions on Nuclear Science*, 62:1413–1421, 6 2015.
- [95] Jean Rinkel, Debora Magalhães, Franz Wagner, Erik Fröjdj, and Rafael Ballabriga-Suñé. Equalization method for Medipix3RX. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 801:1–6, 2015.
- [96] Navrit J. S. Bal. Studies into different detector materials, 2019. <https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5c8e1cf59&appId=PPGMS>.



Dexter - Readout and control software

The readout and control software for the Medipix3RX chip (CERN technology), named Dexter, is a C++ program using the Qt library. This software can now readout and save images at the maximum speed (2000 Hz, 12 bit frames) and control all chip features. It has had several contributors (mainly John Idarrága-Muñoz) and is more than 50k lines of C++ and 10k lines of UI code – it is a complex and relatively mature program which enables full chip characterisation.

Dexter is the readout and control software for the Medipix3RX chip. As such, it sits on top of the stack - it interfaces and communicates with the layers below through various APIs (Application Programming Interface). These APIs include SpidrMpx3Lib which is a C++ library used to communicate with the SPIDR (Speedy PIXel Detector Readout) system. One level lower, there is FPGA firmware and LEON software with which the C++ library communicates. These control the Xilinx Artix 7 FPGA (Field Programmable Gate Array) which in turn controls the SPIDR. The SPIDR then sends and receives signals over a 'flex' (flexible) cable to the assembly and finally the Medipix3RX chips. See Figure A.1 for a graphical overview.

Within Dexter, there is a very simple structure which makes it easy to extend and debug the program; there are no complex hierarchical structures. All main views can communicate to each other via the 'mpx3gui' object; they do this by acquiring a pointer to the specific objects for each view. This is a fast, extensible and safe method to solve this problem. Inside these important GUI (Graphical User Interface) classes, there are sometimes sub-dialogues. These are logically related and may feedback input or variables back up the chain. There are important non GUI classes which are integral to the functioning of the program such as the data taking and data consuming threads. They use the SpidrMpx3 C++ library to decode, read & parse the UDP (User Datagram Protocol) data stream from the SPIDR. Finally, there are external dependencies for several of these components. The commercial framework, Qt 5 underpins all Dexter classes and SpidrMpx3Lib. It provides the

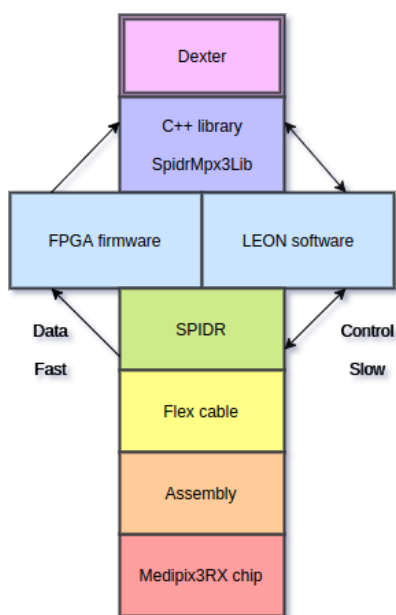


Figure A.1: A logical system overview of the detection system. Dexter controls the system as a top level application.

most extensive feature set one may need to develop a C++ application on the desktop, mobile or embedded system. OpenGL is used for displaying data in the main view; the Phidgets library is for motor control. The others are the Dlib library and a TIFF (Tagged Image File Format) handling library for saving files to the commonly used TIFF format.

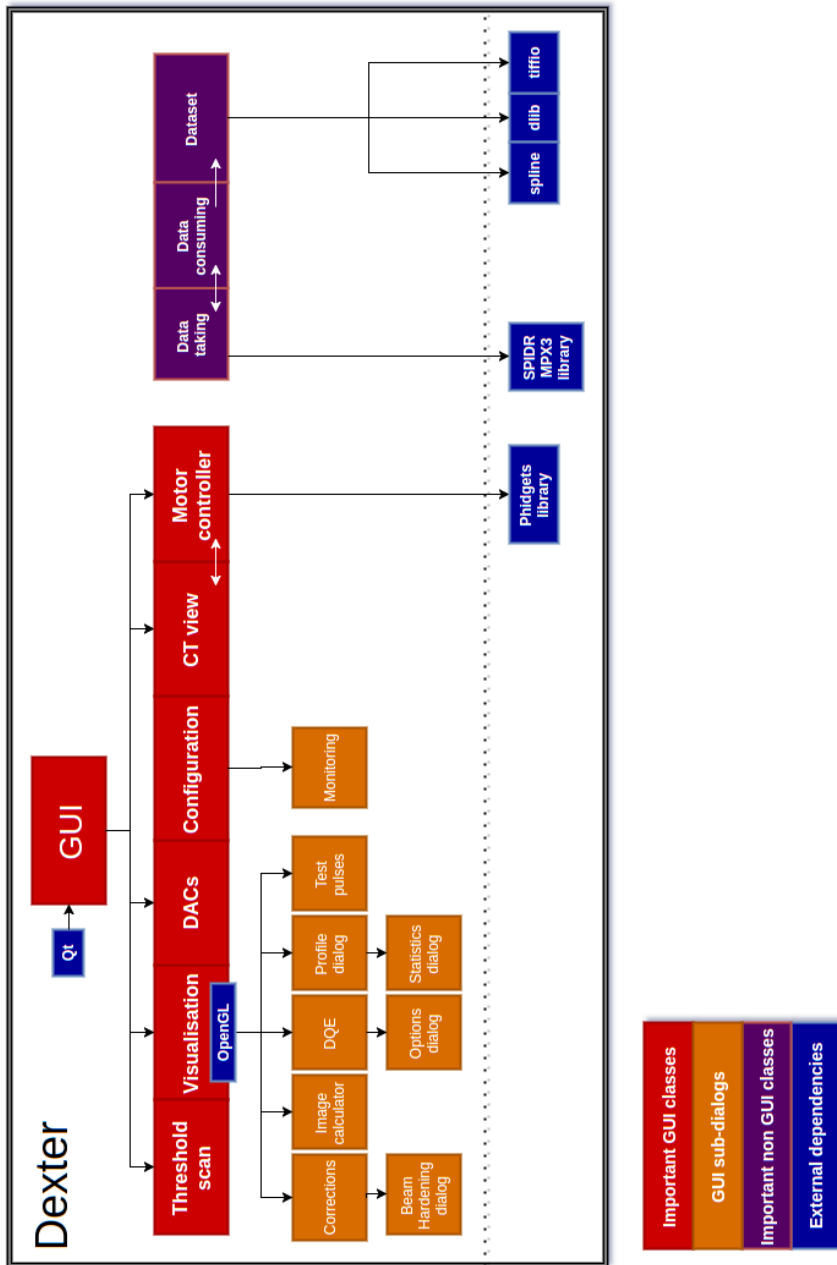


Figure A.2: A logical overview of the internal structure of Dexter. It has several main views (red) with a series of sub-views (orange); the user interacts with these exclusively. Dependencies on external libraries are blue and the important non-GUI (Graphical User Interface) classes (purple) are related to low level data handling classes.

B

Drift rates per pixel

Following on from subsection 2.6.4, the pixel drift rates for runs one, three and four are shown.

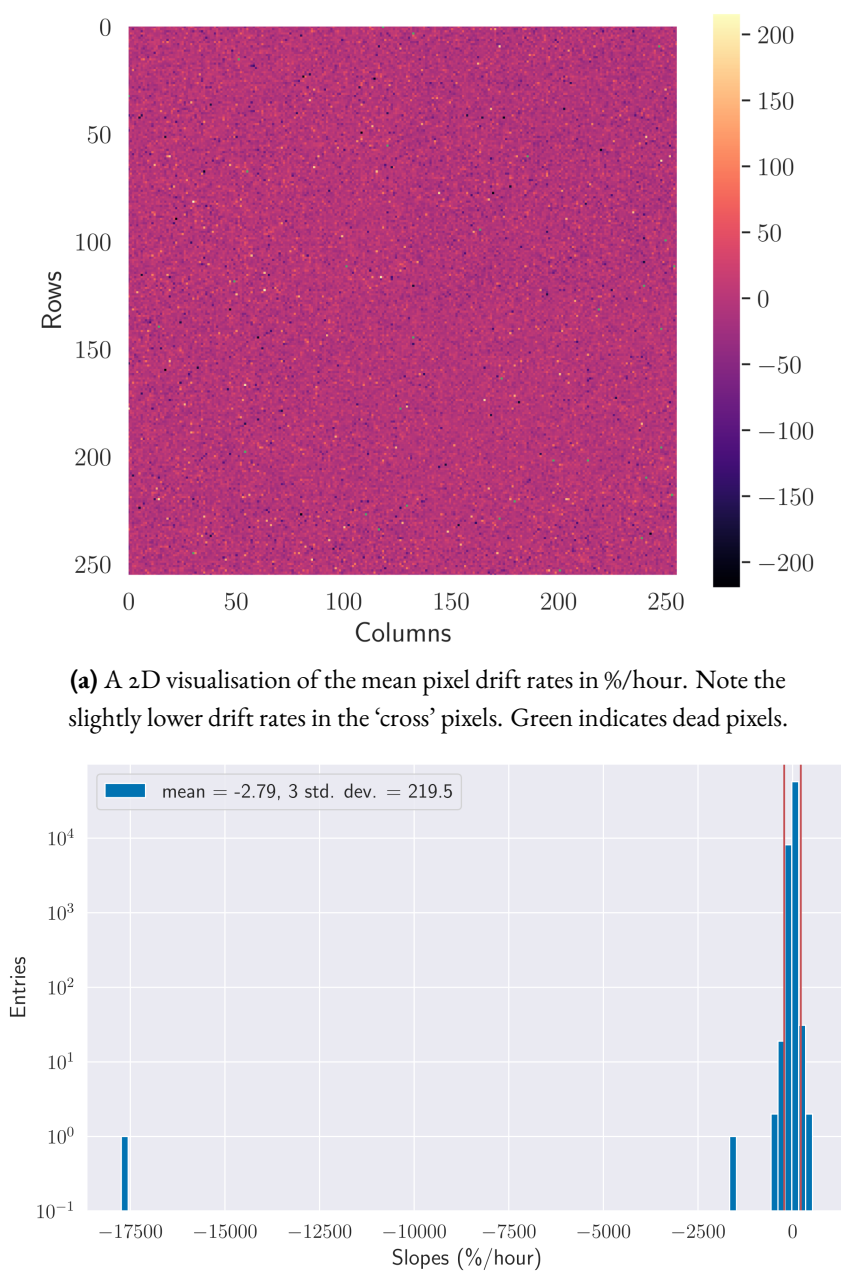




Figure B.1: The mean drift rates of all pixels for run one in %/hour.




(a) A 2D visualisation of the mean pixel drift rates in %/hour. Note the different drift rates around the ‘cross’ pixels. Green indicates dead pixels.




(b) The distribution of mean pixel drifts as a histogram with the mean plus/minus three standard deviations marked with the red vertical lines. Note the log scale on the y axis.

Figure B.2: The mean drift rates of all pixels for run three in %/hour.



(a) A 2D visualisation of the mean pixel drift rates in %/hour. Green indicates dead pixels.



(b) The distribution of mean pixel drifts as a histogram with the mean plus/minus three standard deviations marked with the red vertical lines. Note the log scale on the y axis.

Figure B.3: The mean drift rates of all pixels for run four in %/hour.



MedAustron measurement overview

An overview of all measurements is displayed in tables C.1 and C.2.

| Run | Sub-run | Particle | Energy (MeV) | Degrader (%) | Spot weight |
|-----|---------|----------|--------------|--------------|-------------|
| 1 | | Proton | 800 | 20 | N/A |
| 2 | | Proton | 800 | 20 | N/A |
| 3 | | Proton | 800 | 20 | N/A |
| 4 | Test | Proton | 800 | 10 | N/A |
| | 1 | Proton | 800 | 10 | N/A |
| | 2 | Proton | 800 | 10 | N/A |
| 5 | 1 | Proton | 800 | 10 | N/A |
| | 2 | Proton | 800 | 10 | N/A |
| | 5 | Proton | 800 | 10 | N/A |
| | 7 | Proton | 800 | 10 | N/A |
| | 8 | Proton | 800 | 10 | N/A |
| | 9 | Proton | 800 | 10 | N/A |
| | 10 | Proton | 800 | 10 | N/A |
| 6 | 1 | Proton | 800 | 20 | N/A |
| | 2 | Proton | 800 | 20 | N/A |
| | 3 | Proton | 800 | 20 | N/A |
| | 4 | Proton | 800 | 20 | N/A |
| | 5 | Proton | 800 | 20 | N/A |
| | 6 | Proton | 800 | 20 | N/A |
| | 7 | Proton | 800 | 20 | N/A |
| | 8 | Proton | 800 | 20 | N/A |
| | 9 | Proton | 800 | 20 | N/A |
| | 10 | Proton | 800 | 20 | N/A |

Table C.1: Measurements overview, 1 of 2. Spot weight has the units of numbers of particles specified.

| Run | Sub-run | Particle | Energy (MeV) | Degrader (%) | Spot weight |
|-----|---------|------------|--------------|--------------|-----------------|
| 7 | 1 | Proton | 800 | 50 | N/A |
| | 2 | Proton | 800 | 50 | N/A |
| | 3 | Proton | 800 | 50 | N/A |
| | 4 | Proton | 800 | 50 | N/A |
| | 5 | Proton | 800 | 50 | N/A |
| | 6 | Proton | 800 | 50 | N/A |
| | 7 | Proton | 800 | 50 | N/A |
| | 8 | Proton | 800 | 50 | N/A |
| | 9 | Proton | 800 | 50 | N/A |
| | 10 | Proton | 800 | 50 | N/A |
| | 11 | Background | 800 | N/A | N/A |
| 8 | 0 | Proton | 800 | 100 | N/A |
| | 2 | Proton | 800 | 100 | N/A |
| | 3 | Proton | 800 | 100 | N/A |
| | 4 | Proton | 800 | 100 | N/A |
| 9 | | Proton | 62 | 20 | 5×10^6 |
| 10 | | Proton | 148 | 20 | 1×10^7 |
| 11 | | Proton | 252 | 20 | 1×10^7 |
| 12 | | Proton | 62 | 10 | 1×10^6 |
| 13 | | Proton | 62 | 100 | 1×10^8 |
| 14 | | Proton | 62 | 100 | 5×10^7 |
| 15 | | Proton | 62 | 50 | 1×10^7 |
| 16 | | Proton | 62 | 100 | 1×10^9 |
| 17 | | Proton | 62 | 100 | 5×10^8 |
| 18 | | Carbon | 120 | 20 | N/A |
| 19 | | Carbon | 400 | 20 | N/A |
| 20 | | Carbon | 260 | 20 | N/A |
| 21 | | Carbon | 120 | 100 | N/A |

Table C.2: Measurements overview, 2 of 2. Spot weight has the units of numbers of particles specified.

D

Equalisation

Manufacturing every transistor to be identical is not possible for many reasons including resistivity variations over a silicon wafer, therefore there are process specific manufacturing tolerances from foundries producing these chips such as GlobalFoundries Inc., California. Tighter manufacturing tolerances in the process used would result in chips with more uniform pixel responses. The shaper transistors are particularly sensitive in this regard. For the Medipix and Timepix family of chips, this results in a mismatch between the shaper transistors causing a DC (direct current) offset which effectively means that each pixel has a significantly different noise baseline by default. This can be mostly corrected by using the equalisation procedure.

Equalisation reduces threshold dispersion between pixels, the per pixel noise baseline; it is not the same as the gain variation. The successor of Medipix₃, Medipix₄ is designed to compensate for the inter-pixel gain variation as a function of energy/threshold.


SPM NOISE BASED EQUALISATION PROCEDURE

An overview of the SPM (Single Pixel Mode) in fine pitch mode (FPM, 55 μm pixel pitch) noise based equalisation procedure follows, this is intended to give a flavour for the procedure so it does not include all details or variations for the sake of brevity. When a threshold scan is mentioned in this section it means that one sets a threshold at a particular value, opens the shutter for a short time and records the number of pixels counting above a certain number. 10 counts is often used to filter out cosmic radiation and low level noise. 10 counts were used because the frame times are limited to 100 ms and therefore it is exceptionally unlikely to have more than 10 counts in that time window from cosmic radiation or the very tail of the noise distribution. If one wanted to use the noise edge to calibrate with, the maximum count in a 12 bit frame could be used, which

is 4095 counts. I_DAC_DISCL and I_DAC_DISCH represent the current DAC which controls the gain of the chip; this is shared across all pixels within a chip. Each counter has its respective low (I_DAC_DISCL) and high (I_DAC_DISCH) DACs, in the following section I_DAC_DISC is used to refer to I_DAC_DISCL or I_DAC_DISCH depending on which counter is being equalised. This procedure was mostly implemented in the readout and control software, Dexter, before this work. My contribution can be summarised as taking crash and error prone code from a first implementation state to a highly reliable, stable, commercial state with all features enabled.

Figure D.1 is a visual aid for the following list:

1. Initialisation: Set OMR (operation mode register) bits;
 - Equalisation bit = 1
 - Colour mode bit = 0
 - CSM_SPM = 0
 - DISC_CSM_SPM =
 - DiscL (THL) = 0
 - DiscH (THH) = 1
2. Step one: First threshold scan;
 - I_DAC_DISC = 100 DAC
 - Set all adjustment bits to zero (00000b)
 - A threshold scan is run, an entry is added when a pixel counts more than the minimum value, resulting in a normal distribution. For speed, steps one and two scan with a subset of the matrix.
3. Step two: Second threshold scan;
 - I_DAC_DISC = 150 DAC
 - A threshold scan is run as before, one measures a normal distribution with the same amplitude as before but with a higher mean value because I_DAC_DISC controls the chip-wide gain.
4. Step three: Find the optimal I_DAC_DISC, I_DAC_DISC_optimal;
 - Now that one knows the linear relationship between I_DAC_DISC and the threshold, one can extrapolate what the optimum I_DAC_DISC value is such that: $\mu = 3.2\sigma + \text{equalisation target}$. The equalisation target is fixed at 10 DAC units as a result of early testing at CERN by the chip designers, ideally this target would be zero. This target encapsulates the noise variations of many chips and was deemed to be a



Steps 1 - 3

Find the optimal DAC_DISC (L/H) value using two threshold scans with the values at 100 and 150 DAC units respectively.

Step 4

Find good adjustment bits using two threshold scans with all adjustment bits set to 0 and 5 respectively.

Step 5

Scan with linearly interpolated or extrapolated adjustment bit values

Step 6 (optional)

Fine tune to find the optimal adjustment bits

Figure D.1: An overview of the equalisation procedure with various subplots of real scans from within Dexter (Medipix3 readout and control software).

safe number for all equalisations. A small part of the energy range is lost due to this non-zero target. The 3.2σ value can be varied based on the user's preference, either one sacrifices some pixels for better energy resolution or one makes all the functioning pixels respond with a worse energy resolution.

5. Step four: Scan for interpolation;

- Set I_DAC_DISC_optimal
- Adjustment bits are still zero (00000b)
- Run threshold scans on the whole matrix, pixel spacing is typically four
- Set all adjustment bits to five (10100b)
- Run threshold scans as before

6. Step five: Interpolation.

- For every pixel, find the closest adjustment bit value to the equalisation target (default = 10 DAC) by interpolating or extrapolating between the scans where the adjustment bits were zero and five. Adjustment bits are controlled by a per-pixel five bit DAC, so the possible values are anything between zero and 31.

To further motivate the necessity and importance of equalisation, real example values for before and after the procedure are as follows: Before (unequalised): $\mu = 80.1$, $\sigma = 15.0$ After (equalised): $\mu = 9.8$, $\sigma = 1.5$ This one order of magnitude reduction in threshold dispersion is typical for SPM+FPM.

This procedure varies slightly with colour mode where one in four pixels are bump bonded such that the effective pixel pitch is $110\text{ }\mu\text{m}$ and when charge summing mode (CSM) is active. In all three additional cases (SPM+Colour, CSM+FPM, CSM+Colour), pixel to pixel communication is introduced and so the procedure varies. For colour mode, the OMR (Operation Mode Register) colour mode bit must be set to one. For CSM, the OMR bit Disc_CSM_SPM selects the low discriminator (DISC_L) with zero and the high discriminator (DISC_H) with one, so it must be set to one. For colour mode and CSM, both respective bits must be set to one. This method is independent of gain mode (Super High, High, Low & Super Low) and readout mode (Sequential R/W and Continuous R/W). Further technical details can be found in the Medipix3 manual (24).

There are variations of this method which use test pulses or close to monochromatic x-ray beams(94), one study claims improvements of 14 % (95) by using test pulses and one using monochromatic x-rays showed an improvement of 25 %, the sigma of a distribution decreased from 1.04 to 0.83 keV using 8 keV photons. It is important to note that equalisation is temperature dependent due to higher temperatures causing higher leakage currents which can cause pixels to not respond mostly around the border and some DACs are temperature sensitive. Due to various dependencies on many factors, such as temperature, leakage current, polarity and gain, it is sensible to have equalisation settings for different operation conditions. Specifically for temperature, after approximately $10\text{ }^{\circ}\text{C}$ change, it is advisable to run a new equalisation.



Column based filtering algorithm

```
1
2 def column_based_filtering(data: np.ndarray,
   ↪ above_air_value_but_below_interesting_data: float) -> np.ndarray:
3
4     n_rows = data.shape[0]
5     n_columns = data.shape[2]
6     window_size_max = 5
7
8     # E.g. shape = (256, 600, 256)
9     for row in range(n_rows):
10
11         '''
12         Taking the mean along the axis of rotation angle
13         E.g. the shape of the following two arrays would be (256,)
14         '''
15         column_mean_array = np.mean(data[row], axis = 0)
16         modified_array = np.copy(column_mean_array)
17
18         for w_size in range(1, window_size_max):
19
20             '''
21             Iterating along the row
22             E.g. 1 to 5 to (256 - (1 to 5))
```

E. COLUMN BASED FILTERING ALGORITHM

```
23         """
24         for i in range(w_size, n_columns - w_size, w_size):
25
26             if column_mean_array[i] >
27                 ↳ above_air_value_but_below_interesting_data:
28
29                 modified_array[i] = np.mean(
30                     ↳ np.array([column_mean_array[i - w_size : i],
31                               ↳ column_mean_array[i : i + w_size]]) )
32
33     data[row] *= np.divide( modified_array, column_mean_array, out =
34                             ↳ np.zeros_like(column_mean_array), where=column_mean_array !=
35                             ↳ 0)
36
37     return data
```

Scientific output

The following is a summary of my work over my PhD project. I am responsible for the following items entirely, unless otherwise mentioned.

PAPERS

PUBLISHED

- ‘Medipix3 for dosimetry and real-time beam monitoring: first tests at a 60 MeV proton therapy facility’ - J.S.L. Yap et al. JINST 16 T11001, 2 November 2021;
 - Jacinta Yap and I had an approximately equal contribution towards this. I contributed all Medipix3 content, Jacinta contributed all of the EBT3 film content, Andrzej Kacperek with his team at the CCC facilitated this work with the cyclotron operations and access. Jacinta and I contributed equally towards the remaining sections. Javier Resta López & Carsten Welsch were Jacinta’s supervisors.
- ‘Medipix3 proton and carbon ion measurements across full energy ranges and at clinical flux rates in MedAustron IR1’ - N.J.S. Bal et al. JINST 16 C12002, 1 December 2021.
 - Claus Schmitzer and Sascha Enke were critical for the accelerator control and experimental setup. Andrea De Franco reviewed the paper and gave crucial feedback for technical details which impacted the conclusions. I was responsible for everything else.

CONFERENCE PAPERS

- ‘Beam Characterisation Using MEDIPIX3 and EBT3 Film at the Clatterbridge Proton Therapy Beamline’ - 19-24 May 2019, Proc. 10th International Particle Accelerator Conference (IPAC’19), Melbourne, Australia. DOI: 10.18429/JACoW-IPAC2019-THPMP033.

IN PROGRESS

- ‘TOPAS/Geant4 simulation modelling of the Clatterbridge 60 MeV proton therapy beamline for LET measurements with the MiniPIX-Timepix detector’ - J.S.L. Yap et al.

- My contributions included data analysis, experimental planning and execution, interpretation of data, data acquisition and detector operation.

EU HORIZON 2020 MSCA FELLOW

SCHOOLS AND TRAINING ATTENDED

The OMA (Optimization of Medical Accelerators) project received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 675265. I was a fellow within this project and I produced a number of deliverables. They were merged into later publications or this thesis with the exception of the deliverable D2.5, 'studies into different detector materials' which can be found on the EU R&I participant portal (96).

A significant part of this was the development & documentation of high performance (~7 Gb/s, 2000 FPS) control and readout software for the Medipix3RX chips (high resolution, zero dead-time pixel detector readout chip allowing spectroscopic imaging) in C++ and Qt.

All related news, media and other content is available on the OMA website:

<http://oma-project.eu/>.

- OMA Complementary Skills School - 3-7 April 2017, University of Liverpool, UK;
- 1st OMA School - Medical Accelerators - 5-9 June 2017, Fondazione CNAO, Pavia, Italy;
- 2nd OMA School – Monte Carlo Simulations - 6-10 November 2017, LMU Munich, Germany;
- 2nd OMA Topical Workshop - Diagnostics for Beam and Patient Monitoring - 4-5 June 2018, CERN, Switzerland;
- Cosylab Academy - 6-8 June 2018, CERN, Switzerland;
- 3rd OMA Topical Workshop - Medical Accelerator Design and Diagnostics - 11-12 December 2018, GSI, Germany;
- 3rd OMA School – Medical Accelerators and Particle Therapy - 1-5 April 2019, MedAustron, TU Vienna, Austria;
- International Symposium on Accelerators for Science and Society - 28 June 2019, ACC Liverpool, UK;
- Advanced Researcher Skills School and Technology Transfer Workshops - 24-27 June 2019, University of Liverpool, UK;
- OMA Researcher Careers Workshop - 2-3 September 2019, Seville, Spain;
- OMA International Conference on Medical Accelerators and Particle Therapy - 4-6 September 2019, CNA, Seville, Spain.

DELIVERABLE REPORTS

All demonstrators, pilots, prototypes, documents, reports, websites, patent fillings, videos and so on are available on the European Commission's website for the OMA project:

<https://cordis.europa.eu/project/id/675265/results>.

- D2.3 PC readout software - 25 July 2017;
- D2.4 Performance in clinical environments - 30 April 2019;
- D2.5 Studies into different detector materials - 31 October 2019;
- D2.6 Detector layout - 5 November 2019.

TECHNICAL NOTES

- Internal ASI technical note (27 pages) - CERN secondment knowledge transfer - 07/11/2018.

OTHER SCHOOLS, WORKSHOPS AND TRAINING ATTENDED

- Statistics topical lecture - 2017, Nikhef, Amsterdam, Netherlands;
- 11th Inverted CERN School of Computing (iCSC), 5-8 March 2018, CERN, Switzerland;
- Third Barcelona Techno Week: Semiconductor radiation detectors, 1-7 July 2018, Barcelona, Spain;
- BND school - 10-21 September 2018, Berlin, Germany;
- 1st Allpix Squared User Workshop - 26-27 November 2018, CERN, Switzerland;
- Dark matter topical lecture - 2020, Nikhef, Amsterdam, Netherlands.

FUNDING PROPOSALS SUBMITTED

- 'Spectral microCT system for high resolution 3D visualization of human tumours' - co-authored, not successful - NovoNordisk, 2019.

POSTERS

- ‘A versatile high-speed radiation detection platform’ - June 2017, Milan, Italy;
- ‘A versatile high-speed radiation detection platform’ - September 2017, Munich, Germany;
- ‘A versatile high-speed radiation detection platform using Medipix3’ - 4 April 2019, Vienna, Austria;
- ‘Beam characterisation using Medipix3 and EBT3 film at the Clatterbridge Proton Therapy Beamline’ - 10th International Particle Accelerator Conference (IPAC’19) - 19 May 2019, Melbourne, Australia;
- ‘Medipix3 proton and carbon ion measurements across full energy ranges and at clinical flux rates in MedAustron IR1’ - 22nd International Workshop on Radiation Imaging Detectors (iWoRiD 2021) - 10 June 2021, online only.

TALKS

- ‘Nikhef Medipix3’, Medipix Meeting - November 2017, CERN, Switzerland;
- ‘Hybrid pixel detectors’, Detector R&D Texel trip - 22 June 2018, Texel, Netherlands;
- ‘Pixel detectors in patient treatment’, 2nd OMA Topical workshop - June 2018, CERN, Switzerland;
- ‘Nikhef SPIDR readout system’, Medipix Meeting - November 2018, CERN, Switzerland;
- ‘Colour X-ray CT - Detector R&D’, Nikhef Jamboree - December 2019, Utrecht, Netherlands;
- ‘Beam and detector characterisation using Medipix3 at MedAustron IR1 using protons and carbon ions at clinical flux rates and full energy range’, OMA International Conference on Medical Accelerators and Particle Therapy - 2 September 2019, Sevilla, Spain.

TEACHING

- ‘Particle Detection’ - assisted Prof. Els Koffeman - 2018, Nikhef / Universiteit van Amsterdam, Netherlands.

Summary

Cancer causes 16 % of all deaths per year worldwide. It is the second most likely cause of death in highly and very highly developed nations.

Ideally, we would be able to prevent cancerous tumours from forming. Otherwise, we would like to detect and treat the tumours as early as possible. In reality, we cannot yet entirely prevent or detect tumours very early. Therefore, plenty of cancerous tumours slip through the cracks of our collective healthcare systems. This has significant direct and indirect societal costs including treatment costs, care, rehabilitation, missed work and premature deaths. Estimates of these costs worldwide are thought to be in the order of hundreds of billions of dollars per year, this is approximately 1 % of all global healthcare spending.

How can we improve cancer therapy?

In short: faster, better and cheaper.

Now onto the longer, more specific answer. This thesis focuses on two modalities of cancer therapy, radiotherapy and surgery. The same particle imaging device (detector) is used for both.

Radiotherapy is a cancer treatment which uses radiation to kill cancer cells. Various particles can be used to do this, including light (photons) and the particles which most matter is made of - protons, neutrons & electrons. Atoms with a net charge (ions) are also used, these are composed of protons, neutrons and electrons. Particle therapy often involves accelerating protons and ions into beams which are used to deliver radiation to the cancer cells. These medical accelerators are spin-offs from fundamental particle physics research. The most well known being the LHC (Large Hadron Collider) in Switzerland which discovered the Higgs boson (an elementary particle) in 2012.

These medical accelerators are complex and expensive machines, modern facilities cost in the order of \$200 million. They of course need maintenance and robust monitoring. Part of the monitoring involves checking beam properties. This includes checking the delivered radiation is as expected - deviations can result in missing cancer cells or irradiating healthy tissue. My work included such measurements at two clinical facilities using proton beams and carbon ion beams. It was also possible to measure unexpected behaviour over time in the beam intensities. The detector used was also tested in this relatively extreme environment, it was shown to behave at least as well as the current 'gold standard' method and particularly well at the highest beam intensities. This is encouraging evidence that this type of detector can be used to improve the quality assurance at these facilities.

Moving onto surgery, the most direct way to treat tumours is simply to cut them out of the patient. Ideally, it would be possible to cut out only the cancer cells. In reality, this is not possible and so a margin is added to the targeted volume to remove.

There are two clear problems here. Firstly, the targeted volume is identified using MRI (Magnetic Resonance Imaging) or CT (Computed Tomography) scans and then labelled by a pathologist.

If you take several pathologists, give them the same image, their answers will not be the same as each others’. The scans also can resolve features in the order of a one millimetre cube. These both can result in cancer cells being left in the patient, which go on to spread elsewhere in the body and result in an unsuccessful treatment. Secondly, healthy tissue is removed.

After the tissue is removed, it is prepared for long term storage (archived), sliced into very thin slices of a few cells thickness, a few of those slices are imaged using an optical scanner and then labelled by a pathologist. Almost all of the tissue volume is not imaged or labelled. This process takes several days to several weeks. During this time, the patient is closed up and sent home. If any tissue was found to have been left in the patient, another surgery is required. This is expensive, time-consuming and undesirable in every way.

What if it were possible to non-destructively find the boundary between healthy and tumour tissue during surgery?

My work was the first step in answering that question. The boundary between tumour and healthy tissue was matched to the current method used in clinical practice. It was even possible to observe differences within the tumour itself, better than expected! This was achieved using a custom high resolution CT scanner which also uses the colour (energy) of the x-rays, this is referred to as **‘spectral x-ray CT’**. The next step of this work is to measure many more samples, both fresh and archived.

None of this work would have been possible without a sufficiently good tool to image the particles in question. The Medipix3 is the hybrid pixel detector chip used for this work, let’s break that down. A chip in this case, means an Application Specific Integrated Circuit (ASIC). An ASIC is a combination of tiny electrical components which perform a fixed, very specific set of tasks; this one is designed to detect particles. The Medipix3 chip is electrically connected to a sensor material by a 2 dimensional grid of small metal spheres (bump bonds). Each bump bond is connected to a single pixel on the chip. Other detector chips are often physically the same part as the sensor (no bump bonds). The sensor can be a different material than the chip, so they can be optimised separately. This is why the Medipix3 is a hybrid pixel detector chip.

The Medipix3 is controlled and read by a separate system, the SPIDR (Speedy PIXel Detector Readout) was developed for this purpose at Nikhef. On top of that, specially developed software on a computer is required to communicate with the SPIDR to ultimately acquire images.

Most of the time spent for this work was spent on co-developing the SPIDR software and hardware, developing the software (Dexter) to communicate with the SPIDR, finding out how the Medipix3 actually works and optimising it. This included an effort to determine the optimum measurement time for a given system.

All combined, this work has shown promising advances for cancer treatment. I am continuing the spectral x-ray CT work further at the Danish Centre for Particle Therapy as a post-doctoral researcher.

Samenvatting

Kanker veroorzaakt 16% van alle overlijdens per jaar wereldwijd. Het is daarmee de tweede doodsoorzaak in ontwikkelde en hoogontwikkelde landen.

In een ideale wereld zouden we het vormen van kwaadaardige tumoren kunnen voorkomen. Als dat niet mogelijk zou zijn, dan zouden we de tumoren zo vroeg mogelijk willen detecteren en behandelen. In werkelijkheid kunnen we ze niet voorkomen of heel vroeg detecteren. Daardoor glipt er een heel aantal kwaadaardige tumoren door de vingers van onze gezondheidszorg. Dit leidt tot significante directe en indirecte maatschappelijke kosten, zoals kosten voor behandeling, zorg, rehabilitatie, gemist werk en voortijdige overlijdens. Dit kost wereldwijd, naast persoonlijk leed, in de orde van honderden miljarden dollars per jaar, ongeveer 1% van de wereldwijde uitgaven aan de gezondheidszorg.

Hoe kunnen we de behandeling van kanker verbeteren?

Kortom: sneller, beter en goedkoper.

En dan nu het langere, preciezere antwoord. Dit proefschrift richt zich op twee modaliteiten van de behandeling van kanker; radiotherapie en operaties. Dezelfde deeltjesdetector (detector) wordt voor beide gebruikt.

Radiotherapie is een kankerbehandeling waarbij bestraling wordt gebruikt om kankercellen te doden. Er zijn verschillende deeltjes die kunnen worden gebruikt om dit te doen, waaronder licht (fotonen) en de deeltjes waarvan de meeste materie gemaakt is - protonen, neutronen en elektronen. Elektrisch geladen atomen (ionen) worden ook gebruikt. Deze zijn gemaakt van protonen, neutronen en elektronen. Bij deeltjestherapie wordt de bestraling vaak gerealiseerd door een bundel van versnelde protonen en ionen. Deze medische versnellers zijn spin-offs van fundamenteel natuurkundig onderzoek. De bekendste is de LHC (Large Hadron Collider) in Zwitserland, waar in 2012 het Higgs deeltje (een elementair deeltje) is ontdekt.

De medische versnellers zijn complexe en dure machines, moderne faciliteiten die ongeveer \$ 200 miljoen kosten. Ze hebben natuurlijk ook onderhoud en robuuste kwaliteitscontroles nodig. Een deel van de kwaliteitscontrole is het in de gaten houden van de bundeleigenschappen. Er wordt dan ook gecontroleerd of de gegeven dosis overeenkomt met de geplande dosis, omdat afwijkingen kunnen leiden tot het missen van kankercellen of onnodige bestraling van gezond weefsel. Voor een deel van mijn werk, heb ik metingen gedaan bij twee klinische faciliteiten met protonen- en koolstofionenbundels. We konden hier ook onverwacht het verloop van de bundelintensiteit over de tijd meten. De relatief extreme omgeving was ook een test voor de detector. Het liet gedrag zien dat minstens net zo goed was als de 'gouden standaard' methode en werkte vooral goed bij de hoogste bundelintensiteiten. Dit is aanmoedigend bewijs dat dit type detector kan worden gebruikt om de kwaliteitscontrole van deze faciliteiten te verbeteren.

Chirurgie is de meest directe methode om tumoren te behandelen, door ze simpelweg uit de patiënt te snijden. Idealiter zou het mogelijk zijn om alleen de kankercellen weg te snijden. In werkelijkheid is dit niet mogelijk en moet er een marge aan het doelvolumen worden toegevoegd.

Er zijn hier twee duidelijke problemen. Ten eerste wordt het doelvolumen geïdentificeerd door MRI (Magnetic Resonance Imaging) of CT (Computed Tomography) scans, die worden gelabeld door een patholoog. **Als je meerdere pathologen dezelfde beelden geeft, dan zijn hun antwoorden niet hetzelfde.** De scans hebben daarnaast een resolutie van een kubieke millimeter. Deze factoren kunnen ertoe leiden dat kankercellen in het lichaam achter kunnen blijven, die zich door de rest van het lichaam kunnen verspreiden en resulteren in een onsuccesvolle behandeling. Ten tweede wordt er gezond weefsel verwijderd.

Nadat het weefsel is verwijderd, wordt het voorbereid om langdurig te worden bewaard (gearchiveerd). Het wordt in hele dunne plakjes gesneden ter dikte van een paar cellen. Een paar van die plakjes worden gescand met een optische scanner en door een patholoog gelabeld. Dit proces duurt enkele dagen tot enkele weken. Gedurende deze tijd moet de patiënt thuis de uitslag afwachten. Als blijkt dat er tumorweefsel in de patiënt is achtergebleven, dan moet de patiënt opnieuw worden geopereerd. Dit is duur, kost veel tijd, is onwenselijk en belastend voor de patiënt.

Wat als het mogelijk zou zijn om de grens tussen gezond en tumorweefsel te vinden op een non-destructieve manier tijdens de operatie?

Mijn werk was de eerste stap naar het beantwoorden van die vraag. De grens tussen de tumor en het gezonde weefsel werd gemacht met de huidige methode in de klinische praktijk. Het was zelfs mogelijk om verschillen in de tumor zelf te zien. Dat is beter dan verwacht! Dit was mogelijk door gebruik van een speciale hoge resolutie CT scanner die ook de kleur (energie) van de röntgenstraling gebruikt. Dit wordt ook wel **‘spectrale röntgen CT’** genoemd. De volgende stap van dit werk is het in beeld brengen van nog veel meer samples, zowel verse als gearchiveerde samples.

Dit werk zou niet mogelijk zijn geweest zonder een meetinstrument dat goed genoeg is om de deeltjes in beeld te brengen. De Medipix3 is de hybride pixel detector chip die in dit werk gebruikt is. Laten we daar verder op in gaan. In dit geval bedoelen we met een chip een Application Specific Integrated Circuit (ASIC). Een ASIC is een combinatie van microscopisch kleine elektrische componenten die een vaste, hele specifieke set taken uitvoert; het is ontworpen om deeltjes te detecteren. De Medipix3 chip is elektrisch gebonden aan een sensor materiaal door een tweedimensionaal rooster van kleine metalen bolletjes (‘bump bonds’). Elk klein bolletje is verbonden aan één enkele pixel op de chip. In andere detectoren maken de detectorchips meestal fysiek onderdeel uit van de sensor (zonder ‘bump bonds’). De sensor in hybride chips kan van een ander materiaal dan de chip gemaakt zijn. Ze kunnen hierdoor apart worden geoptimaliseerd. Dit is waarom de Medipix3 een hybride pixel detector chip is.

De Medipix3 wordt door een apart systeem bestuurd en uitgelezen, namelijk de SPIDR (Speedy PIxel Detector Readout), die voor dit doel op Nikhef ontwikkeld is. Daarbij is er ook speciaal ontwikkelde software op een computer nodig om met de SPIDR te communiceren en uiteindelijk

beelden te krijgen.

Het grootste gedeelte van de tijd van dit werk is besteed aan het mede-ontwikkelen van de SPIDR software en hardware, het ontwikkelen van de software (Dexter) om met de SPIDR the communiceren, het onderzoeken van hoe de Medipix3 eigenlijk werkt en het optimaliseren ervan. Hieronder valt een poging om de optimale meettijd voor een gegeven systeem te bepalen.

Alles bij elkaar genomen, heeft dit werk veelbelovende vooruitgang laten zien voor de behandeling van kanker. Ik ga door met het spectrale röntgen CT werk bij het Deense Centrum voor Deeltjes Therapie (DCPT) als postdoctoraal onderzoeker.

Acknowledgements

- John Idarrága-Muñoz
- Jacinta Yap
- Kevin Heijhoff
- Bram Bouwens
- Els Koffeman
- Jan Willem Palenstijn
- Alexander Kostenko
- Mathé Zeegers
- Sophia Coban
- Robbert Geertsema
- Wessel Rijk
- Rafael Ballabriga-Suñé
- Jan Visser
- Oscar van Petten
- Joop Rovekamp
- Bas van der Heijden
- Frans Schreuder
- Henk Boterenbrood