Proton Imaging with Gridpix Detectors.

Master Thesis in Physics, track Particle and Astroparticle Physics
Supervision and Review: dr. Jan Visser and prof. dr. Sytze Brandenburg

Brent Fokke Bram Huisman

August 29, 2012

CONTENTS

1 Abstract **1**

2 Introduction **2**
2.1 Structure of this document . . . 3

3 From High Energy Physics to Medicine **4**
3.1 Types of Radiation in Medicine . . . 5
3.2 Particle Beams for Hadron Therapy . . . 7
3.3 Imaging and Computed Tomography . . . 13
3.4 Modalities in Medicine . . . 17

4 Hadron Imaging **18**
4.1 Nuclear Emulsion Film Detection . . . 18
4.2 Strip Detectors . . . 19
4.3 Strip Detectors and separate Calorimeter . . . 20
4.4 GEMs and Sampled Silicon Calorimeter . . . 21
4.5 Scattered Events Beam Reconstruction . . . 23
4.6 The Nikhef/KVI approach . . . 23

5 Improving the Setup **26**
5.1 Results from November and Overview . . . 26
5.2 Towards pCT . . . 31
5.3 Calorimeter . . . 33
5.4 Calorimeter Data Acquisition . . . 36
5.5 Trigger . . . 43
5.6 Conclusions . . . 48

6 Further Thoughts **50**
6.1 Tracking . . . 50
6.2 Calorimetry . . . 51
6.3 Trigger . . . 51
6.4 Acknowledgments . . . 52

7 References **54**

CHAPTER

ONE

ABSTRACT

Hadron therapy is an upcoming cancer treatment modality with a potentially better spatial accuracy than X-ray irradiation. While an X-ray beam displays a exponential decay in intensity as it traverses matter, a hadron beam has depth profile with a sharp peak, at which point most particles will be stopped. The practical consequence of this behavior is that hadrons allow dose to be deposited more accurately than X-rays. Hadron therapy is therefore well suited to the treatment of tumors located close to sensitive organs. Imaging of the tumor and its surrounding tissue still heavily relies on X-ray CT (Computed Tomography). To construct a treatment plan for hadron therapy, the X-ray radiodensity map must be converted to a stopping power map for particle beams. This conversion introduces an inherent uncertainty of up to 3%, compromising the improved accuracy that hadron therapy potentially provides. Imaging in the same modality as the treatment would remove this conversion error. This thesis describes the Nikhef/KVI Proton Imager, a device built to measure the hadron radiodensity directly.

To construct a hadron CT image, the Detector Research and Development group of Nikhef in Amsterdam designed a prototype setup with a GridPix as a particle tracker and a crystal with PMT as calorimeter. GridPix is a miniature drift chamber atop a TimePix 256x256 pixel chip, allowing particles to be tracked in 3D. In November 2011, the first test of the setup was conducted at the KVI proton accelerator, which proved that the concept works. The calorimeter and tracker data were successfully correlated to objects placed in the beam. The GridPix readout is limited to an event acquisition rate of 120Hz, but a rate of 2Hz was attained, chiefly due to a conservative trigger design and a suboptimal calorimeter data acquisition system. In order to increase the event acquisition rate to 120Hz, these two parts of the setup were overhauled. A new calorimeter data acquisition system can read out continuously and indefinitely, greatly enhancing the previous run length of 256 events. Readout rates up to 31kHz were obtained. A new trigger design takes into account a busy signal from the GridPix readout system, and reduces a timeout of 500ms to the approximately 10ms GridPix readout time, increasing the potential event acquisition rate to at least 100Hz. The final systems were tested in a multiday run measuring cosmics and in a short run with a frequency generator. A second test at the KVI proton accelerator should confirm the performance improvements, but has not yet been conducted.

CHAPTER

TWO

INTRODUCTION

This document is the record of the graduate research project conducted by Brent Huisman from October 2011 to August 2012 for the P&AP (Particle and Astroparticle Physics) master program of the University of Amsterdam, at the Nikhef institute, in the Detector R&D (Research and Development) group. About half a year before the start of my work, a project was started in the Detector R&D group on a proton imaging system based on the Medipix family of detector chips, intended for medical imaging. My research task was the continuation and enhancement of this proton beam monitor setup. Together with the people who started the project, the proof of principle was delivered by testing the setup at the KVI (Kernfysisch Versneller Instituut) accelerator in Groningen. The test was successful, but many points for improvement were discovered. Properly charting these improvements and making a start with the implementation was my main occupation during the project.

At first glance this project may seem far removed from the topic of the P&AP master program. This field and therefore the master program is concerned with the characterization of the fundamental building blocks of matter, space and time, with which it is hoped we can develop a greater understanding of the world around us. Technology and knowledge acquired for fundamental research of this type can, and is, applied in a plethora of other fields however. One such field does so with huge success, and that field is medicine. Inspired by a lecture on proton therapy by Ugo Amaldi at the Cern Summer School of 2011, I chose the option provided by the Nikhef Detector R&D group to work on a medical proton imager. Next to constructing images for diagnostic use, this system can be used to obtain an accurate stopping power distribution necessary for treatment which is now commonly obtained via a converted X-ray CT image. Part of any radiological treatment is the construction of an image of the tissue distribution, which is used to calculate the best treatment plan. Particle beams and X-rays probe different properties of matter, so a conversion from one radiodensity (X-ray) to another (protons) has an inherent error that translates to an extra inaccuracy. To remove this conversion would be an improvement.

Proton beam imaging is more sensitive to certain differences in tissue than X-ray imaging, and delivers less radiation dose while doing so [21, 55, 56, 43, 18, 12, 57, 6]. While medical proton beam imaging is not a new idea (figure 2.1 [11, 7, 21]), a novel approach was employed in this Nikhef project. Within the Detector R&D group much research is conducted with the MediPix detector chip family, which is a programmable two dimensional charge detector. This project uses such chips as part of the proton beam monitor, whereas many alternatives use more traditional methods such as silicon strip detectors.

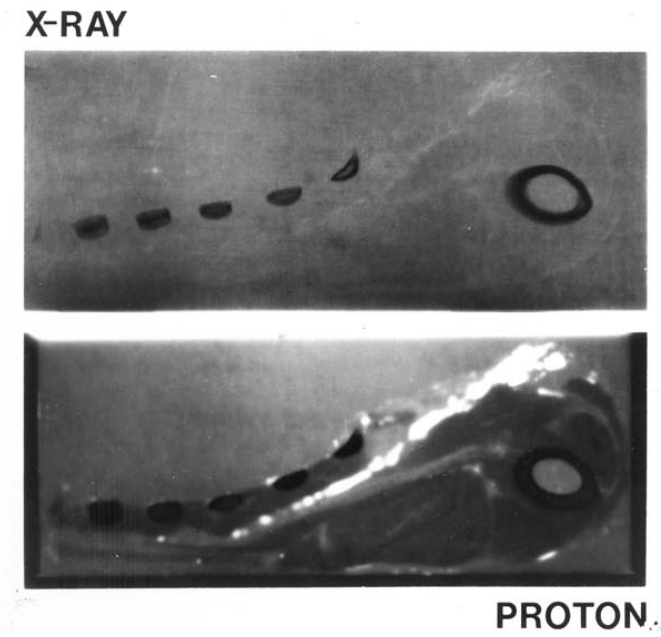

Figure 2.1: A comparison of an X-ray image (top) and a proton image (bottom) of a lamb chop, the first published proton image by Koehler [7].

Proton beam imaging is a technique that has the potential to replace or augment X-ray imaging, because protons have superior sensitivity to tissue density in the patient compared to X-rays [7, 43] and therefore allow for a more precise treatment plan when proton beam therapy is used [9, 23, 58, 30]. One of the two main reasons is that X-ray CT in practice involves a model dependent analysis, wheras proton imaging allows us to directly probe the proton-tissue interaction, omitting inherent imprecisions of the X-ray model. The other reason is that X-ray absorbtion decays exponentially as it traverses matter, whereas protons have a more or less linear absorption. The expensive nature of facilities capable of proton therapy and imaging makes that the techniques are not as common and developed as other methods, but the method is certainly gaining in popularity [6, 19, 59, 8].

2.1 Structure of this document

This document is laid out as follows: in *chapter 3* the context for proton beam imaging will be set. First, it will be illustrated how physics research has a long history of spilling over into medicine and why the application of hadrons (e.g. protons) are now upcoming. Secondly, a short recapitulation of the physics behind hadron imaging and therapy will be presented. Lastly, it will be explained where exactly hadron imaging fits in the broader world of medical imaging methods.

Chapter 4 will discuss current approaches to hadron imaging and will conclude with the Nikhef proposal for a proton beam imager. In *chapter 5* the setup will be explained in full detail, as well as the experimental results attained thus far. *Chapter 6* will summarize the accomplishments of the Nikhef setup, highlight the points of improvement and discuss the future of the project within the context of hadrons in medicine.

CHAPTER

THREE

FROM HIGH ENERGY PHYSICS TO MEDICINE

Since the Renaissance the development and use of tools in the natural sciences have allowed our species to conduct research further and further beyond the limits of our biological perceptions. These tools often proved interesting subjects of research themselves, and have spun off many useful and presently completely commonplace inventions. Such research has involved the construction of ever larger and better detectors. Legend has it that an apple caused Newton to derive the fundamental laws of motion, and in the age of Rutherford the atomic nucleus was discovered using a tabletop setup. Current forays into fundamental science require tools that range in size from the church-sized Atlas detector in Geneva to the 3000 square kilometer Pierre Auger observatory in Argentina. Early particle detection often was based on photographic film and the measurement of heat. Nowadays there are many kinds of detection methods like silicon chips, scintillating crystals, sampling calorimeters, and computers combining millions of data channels to describe a single event.

These devices are known for their use in the fundamental study of matter, but the same designs can be used to diagnose any type of material, such as a patients tissue. Knowledge of radiation, atoms, particle accelerators, magnetic confinement, photomultiplier tubes, and calorimetry are ingredients that can be taken from high energy physics and be combined with medical and biological knowledge to be used in new methods of diagnosis and therapy. This sort of crossover has happened before: the aforementioned photographic film used to discover radiation quickly found an application in medicine. X-rays would be sent through a patient and then be recorded on film, later replaced with silicon sensors. The PET-scan comes to mind as well: a marker element that decays through β^+ emission, recombines with an electron and emits two photons. These photons leave the body and point straight back to their location of creation, and so reveal the location of the marker. If the marker is designed to attach itself to a certain type of organ or tissue, the position of that certain material will be known. Both are examples of how knowledge of physics and biology is used to probe matter and therefore tissue with radiation, quite similar to physics experiments.

It is important to make the distinction between imaging and treatment. Imaging usually refers to the 2D, 3D and sometimes even 4D reconstruction of the body's tissues and the detection of abnormalities. Treatment is concerned with the alteration of abnormalities, which in the context of radiology usually means the destruction of a tumor. The biggest

practical difference is that for imaging the radiation needs to be captured after the interaction. Also, the deposited dose in the body is preferably kept as low as possible, while for treatment the dose should be as high as possible in the region that is to be treated. These different goals translate into different devices, and therefore in different physics.

3.1 Types of Radiation in Medicine

Radiation is a broad term, which has historically been subdivided in α, β and γ radiation, respectively ionized Helium, electrons and photons. Before the advent of artificial generators of radiation the most common source of radiation was nuclear decay, of which only these three types of radiation were easily detected. Nowadays however, radiation may include any species of particle traveling at sufficiently high velocity allowing it to traverse the body. The different types of radiation are in medicine referred to as modalities. Another detail is that photons generated by atomic decay fall into a certain energy range, roughly 0.1 - 100keV, historically referred to as X-rays. Photons with higher energies, in roughly the 10keV - 1MeV range, refer to photons emitted by excited nuclei. Photons with even higher energies, generated by stellar phenomena or manmade accelerators, are also most often referred to as γ-rays. Most important is to realize that these ranges are not universally applied, and that nowadays the two types of radiation are distinguished by their origin: X-rays are emitted by electrons outside the nucleus, while γ-rays are emitted by the nucleus [1, 2].

Photon imaging is well researched and now essential for modern medicine. MRI, PET-scans and X-rays all utilize photons, the difference being their energy ranges. In an MRI scanner, a magnetic field is used to align the proton spins. After the field is switched off, the spins are returned to their thermodynamic equilibrium. The relaxation causes a photon to be emitted in the $10^{-5} - 10^{-7} eV$ range. A PET machine records the emissions of photons in the 100keV range. These are generated in the decay of an unstable atomic nucleus, administered to the patient as part of a tracer fluid. The nucleus emits a positron that immediately recombines with an electron, which generates the photons. X-rays are created by accelerating electrons into a heavy-Z target, usually metal. Subsequently the photons are generated in one of two atomic effects: X-ray fluorescence or Bremsstrahlung. The former effect relies on the electrons ability to knock an orbital electron from an inner shell. As the outer electrons fill up the vacancy, they emit the X-rays. In the latter effect the accelerated electrons experience a very strong force close to the charge carriers of the target metal and especially strong near the target nuclei. This causes the electrons to bend and emit Bremsstrahlung: the X-rays. These X-rays are then channeled through the patient. Capture of X-rays has been possible since the development of photographic plates, which discolor depending on the intensity. For radiological treatment of tumors virtually always X-rays are used. With an energy ranging between 0.5 - 20MeV, these rays have enough power to seriously damage tissue. Figure 3.1 shows that they exhibit a logarithmic decay as they traverse matter, which has the practical consequence that an external beam can reach any part of the body.

The reason why α radiation has been neglected in favor of X-rays is that they have a very limited range: when created in atomic decay a sheet of paper is enough to stop them. This means they are not suited for treatment or imaging. Electrons have the problem that they

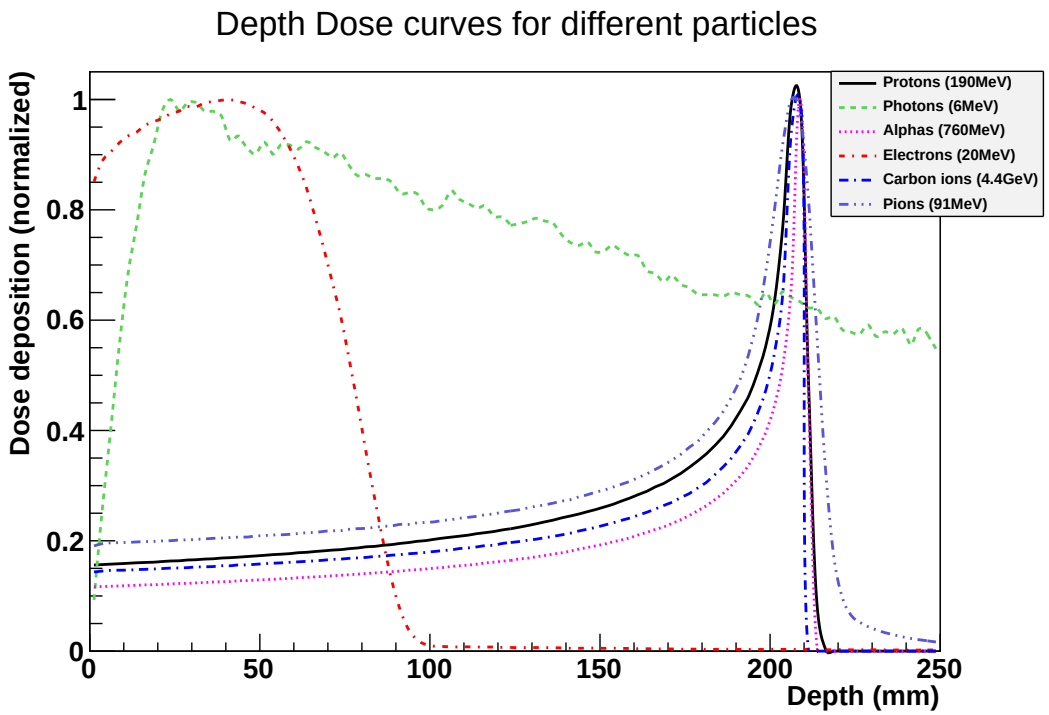

Figure 3.1: Deposited dose plotted as function of depth in tissue. Image generated using Geant4, and all curves are normalized.

scatter enormously in tissue: they are almost 2000 times lighter than protons, and even moreso than heavier ions [3]. It is therefore unlikely that they will be useful for imaging. Artificially accelerated electrons can have a higher energy and more penetrating power, and have been used for superficial treatment. Electron accelerators on a medical scale can reach up to a few tens of MeV [4], at which energy they may reach a depth of about 10cm in tissue, still not quite enough to be useful for all kinds of treatments. The NIST tables [5] show that to reach a depth of 20cm, electrons require an energy of about a GeV. Above a few tens of MeV energy losses due to synchrotron radiation dominate, making acceleration to higher energies inefficient. Two modern electron accelerators, ALBA and CEBAF attain energies of 3 and 6 GeV respectively, requiring impractically large accelerators for a hospital setup. The former has a synchrotron ring of 270m in circumference, while the latter has two linear accelerators of 1400m coupled by two half rings. All things considered, electron therapy is not practical except for superficial treatment.

The therapeutic use of alphas has not yet become common due to prohibitive costs of manmade accelerators necessary to penetrate the body. As can be gleaned from figure 3.1, at energies of about 200MeV they *can* penetrate the human body. The depth that can be reached is a function of beam energy, which becomes unpractical only at energies well above what is required in the medical context. With the advent of hadron accelerators since the mid-twentieth century their application in medicine has been studied. Helium is rather sparse, so usually protons or carbon ions are accelerated. The general term for this category is now hadron therapy, since any hadron can be accelerated with such a machine. The depth-dose curve shows a clear and sharp peak, the Bragg-peak, which means dose is delivered to a specific region as opposed to the much more even distribution of X-rays. Because physics research have made cost and size of accelerators come down,

and because of the desirable depth-dose curve of hadrons, interest has steadily increased since the first treatment of patients in 1954 [6]. In 1968 the first proton-image was made by Koehler [7], showing that there is potential for imaging using hadrons too.

3.2 Particle Beams for Hadron Therapy

An accelerator uses an electric field to accelerate charged particles (by either attraction or repulsion). Magnetic fields are used to keep the particle in orbit, in modern synchrotrons usually a circular one. As particle speed increases, confinement and control of the particles becomes very difficult: that is why an enormous engineering effort is necessary to build stable and robust particle accelerators. Of course physics research, accelerator development and hadron beam therapy developed hand in hand, taking advantage of each new development in one of the disciplines. Only in the last few decades have accelerators specifically built for medicine been feasible, and it is not hard to imagine that size and cost are deciding factors. Schardt and Elsässer [6], Amaldi and Braccini [8], Schippers and Lomax [9] provide excellent overviews of the world-wide state of affairs of medical accelerators, and Schardt and Elsässer [6] mention for example a 20 ton accelerator being installed in two hospitals late 2010. This indicates development is at full speed. Two types of accelerators are found in medical setups: the synchrotron and the cyclotron. The first type is seen in most designs intended for carbon therapy, while the second is found in most proton facilities, as well as in most standalone devices presented by companies who focus on downsizing the accelerator. While the synchrotron accelerates particles in batches, the cyclotron operates continually and gives off a current of particles.

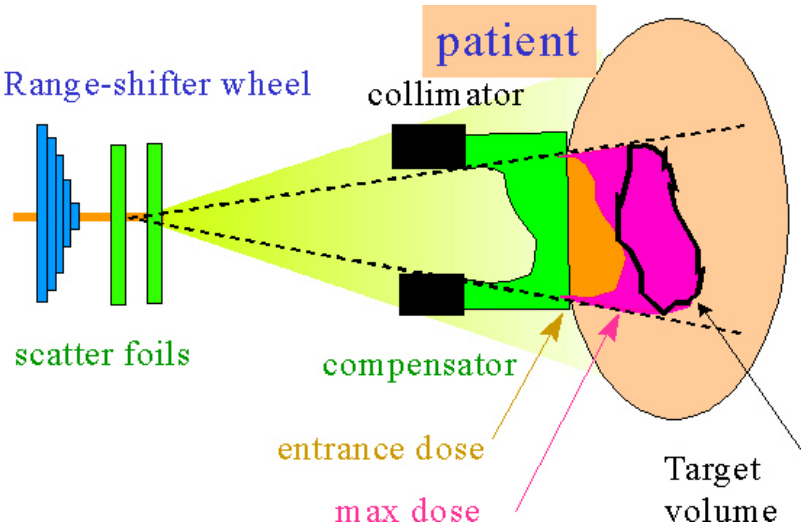

Figure 3.2: An illustration of the typical components used in passive beam shaping. The compensator or bolus shapes the beam into the tumor shape.

A treatment plan is the term for the best approach to irradiating the tumor, and is constructed on a case-by-case basis in the hospital. At the moment of discovery, the tumor will typically have a size multiple centimeters in all 3 dimensions, and be located under layers of other tissues. Each volume element, or voxel, has its specific coordinates and requires a specific dose. The treatment plan depends therefore on the beam shape,

beam energy, dimensions of the tumor and the sensitivity of surrounding tissues. The incident angle of the beam line itself can be controlled: most hospital setups feature a gantry, which is a name for the installation that can rotate the beam nozzle around the patient. Usually the table on which the patient is laid down can move. With the beam energy we can control the depth of the Bragg-peak, and by aiming the beam we control the other 2 dimensions. There are generally two ways the shape (the cross section of the beam) is controlled: passive beam shaping or spot scanning [6, 8, 9].

- Passive beam shaping is a method where the beam is attenuated with passive methods (fig. 3.2). Although it has many pieces, once set there is little or no interaction required for the beam to deliver the proper dose to the targeted areas of the body. The beam is first slightly scattered after which structures such as collimators and attenuators shape the beam into a uniform 2D distributed flux, so that a certain area can be uniformly irradiated at once. Then there is a patient specific mold, a bolus, that shapes the range such that the Bragg-peak is delivered at the proper depth. Some atomically light plastics are used to alter the range and not scatter as much as heavier materials. To alter the penetration depth of the beam, range shifters may be used. This is a wheel with plates of varying thickness which can be spun in the beamline, thereby altering the range of the protons by decelerating them, depending on the thickness of the wheel at that point.

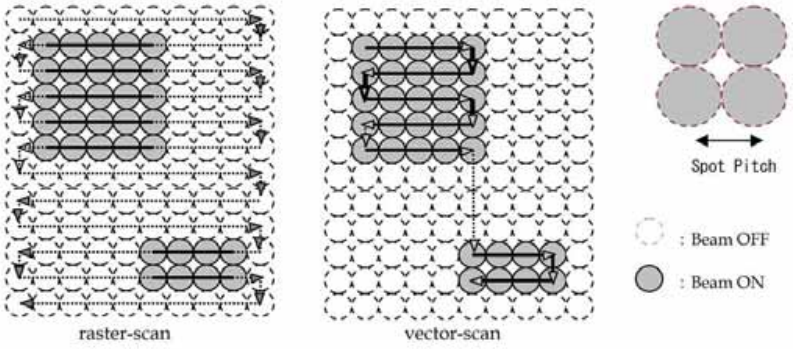

Figure 3.3: On the left raster scanning and on the right vector scanning. Vector scanning can be seen as an optimization of raster scanning, where some speed is gained in favor of a slightly more complicated treatment plan.

- Spot scanning [10], or raster scanning, moves the beam spot by a modulating a magnetic field at the accelerator nozzle. The energy can be controlled by altering the accelerator energy or using a ranger shifter wheel between the accelerator nozzle and the patient. This scanning may occur in a raster-like fashion, as found in old Cathode Ray Tube monitors, or by scanning only over the lateral projection of the tumor, sometimes referred to as vector scanning (fig. 3.3).

These aspects have their respective advantages and disadvantages. Putting materials in the beam line decreases particle rate and increases scattering, reducing the potential resolution. Moving objects in and out of the beam requires the beam to be off and is therefore dead time. Modulating the beam energy requires time as well, so the speed at which this can change has consequences for the possible treatment plans. Because the Bragg-peak is so precise, the patient is usually fixed to the table to within a few millimeters. This is

very uncomfortable, so a treatment must not last too long. Parts within the body move as well: blood flows, lungs move while breathing and bowels contract and expand. For example, the HIT (Heidelberg Ion Therapy) facility in Heidelberg takes breathing into account by switching the beam off every cycle, for a certain period of time, to ensure accuracy. Of course, this again costs time, because the amount of dose that needs to be deposited doesn't change. In other words, dose delivery is a major point of attention. Other strategies of motion compensation or mediation are being researched [6].

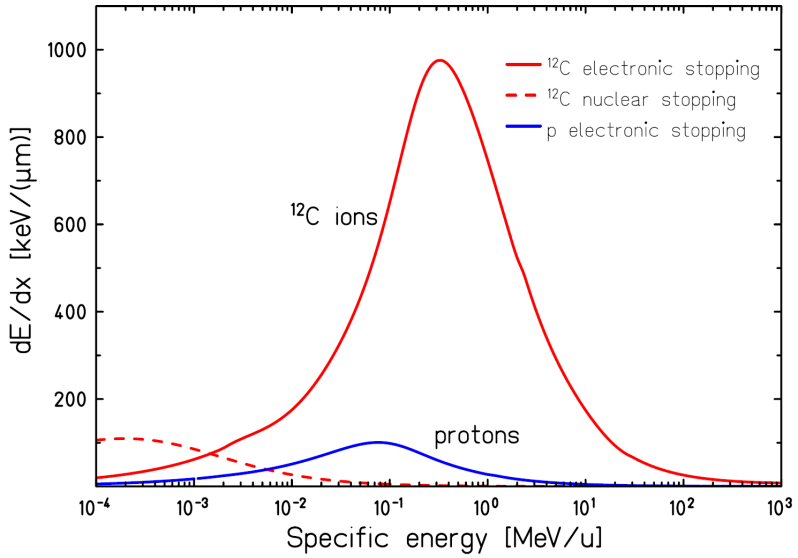

Figure 3.4: dE/dx curve, or stopping power, for protons and carbon ions as function of initial energy. Assuming the initial energy is in the range of a few 100MeV, it is clear how the Bragg-peak comes about: the deceleration accelerates itself Schardt and Elsässer [6].

Then there are the details of the interaction of hadrons with matter, not only collimators and range shifters, but also the body itself and also the air between the parts. The Bethe-Bloch equation tells us how much energy is lost per unit distance, and this curve is plotted for protons and carbon ions in figure 3.4. The deceleration accelerates itself, which is why most energy is lost over a very short distance (fig. 3.1). Wilson [11], Schardt and Elsässer [6] list the following important effects of hadrons in matter:

1. Range straggling.

   The hadron loses energy by ionizing other particles in its path. This is a statistical process, and therefore introduces smearing to the Bragg peak. The Bragg-peak size for protons is about 1% of their range in a certain material. Considering the dimensions of a tumor, this is a beneficial effect. The whole tumor needs to be irradiated, and range straggling increases the minimum voxel size and thereby speeds up treatment. Imagine that pin-pricking the tumor is precise, but takes a long time if the whole tumor volume needs to be treated.

2. Multiple Coulomb scattering.

   Multiple Coulomb scattering, or Elastic Coulomb scattering, describes interactions that cause many small angle scatterings of the hadron, which in turn cause a spread of the beam in the transverse plane. Using equation (3.1) from Molière's theory, where x is the traversed depth through a certain material and X_0 the radiation

length of that material, U. Weber of Rhön-Klinikum has constructed figure 3.5. For the case of the November testrun, assuming a proton with 150MeV, this translates to a angular spread for 3m of air in $\theta_0 = 7.4mrad$, for 20cm of water in $\theta_0 = 66mrad$, for a GridPix (2x10μm Cu and 2x50μm PCB) in $\theta_0 = 5.6mrad$ and for 5cm of aluminium in $\theta_0 = 66mrad$.

$$\theta_0 = \frac{13.6MeV}{\beta cp} z \sqrt{x/X_0} \left[1 + 0.038 \, ln(x/X_0)\right] \tag{3.1}$$

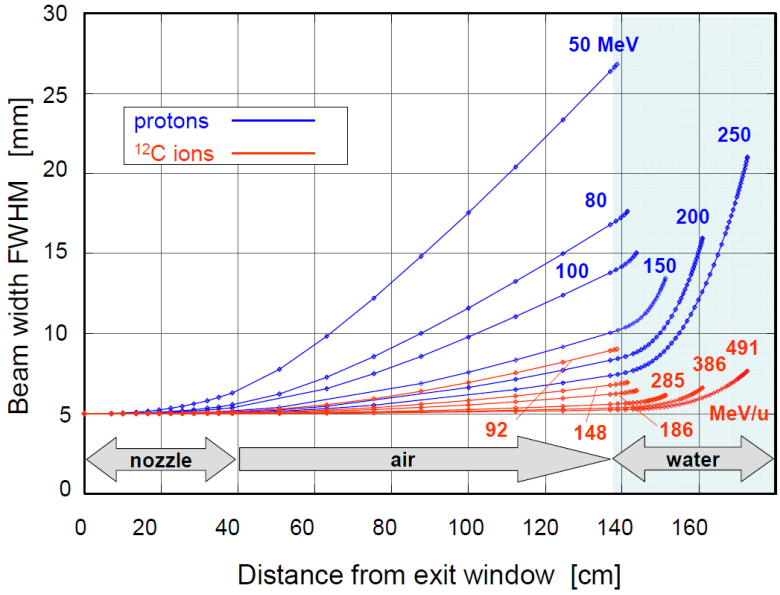

Figure 3.5: The beam spread for different types of hadrons at different energies are shown for a certain beam and a certain setup. The trend visible is that particles with higher energies and/or heavier ions provide a smaller spot size. This effect in combination with the former two results in a spread of the Bragg peak over a small volume, which is a few millimeters in depth (figure 3.1) a bit more in the transverse plane. Note that the distance from the exit window need not be this large, and the relative performance of the (lower energy) protons can be increased.

3. Nuclear fragmentation.

   Target nuclei and the accelerated ion itself may break up after a head on collision. A broken up ion translates to various new lighter nuclei traveling at a similar speed to the parent ion. Lighter ions have a longer range (see figure 3.4) and therefore contribute to a tail after the Bragg-peak of the parent ion. Heavier ions tend to break up into heavier parts, which in turn can break up again. Figure 3.6 shows how for ^{20}Ne this results in a forwardly smeared depth-dose distribution. While usually seen as a disadvantage, Kraft [12] describes it in the case of carbon ions as a potential advantage: typically beta emmitting ^{10}C and ^{11}C are produced that could be used for direct verification with a PET scanner. Because protons can not break up, they don't suffer from nuclear fragments and the associated downsides. Amaldi and Braccini [8] argue however that lighter nuclei, such as He, Li and B, have a sharper distal edge to the Bragg-peak and suffer less from lateral scattering because of the increased mass.

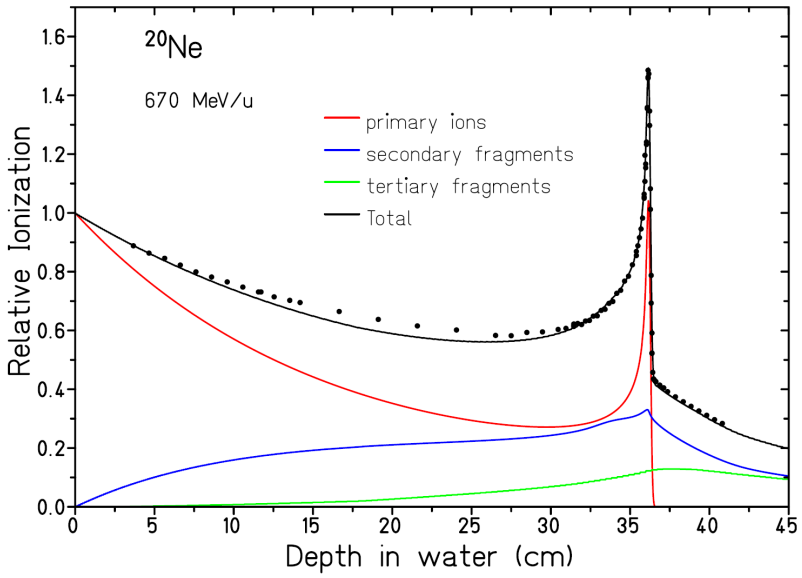

Figure 3.6: Dose curve for 670 MeV/u ^{20}Ne ions in water measured at GSI and calculated contributions of primary ions, secondary and tertiary fragments. Image courtesy of Schardt and Elsässer [6].

4. Neutron dose.

   When protons and host nuclei collide, neutrons may be produced. They keep the speed and direction of the incident proton. Passive beam shaping produces more neutrons than an active beam shapers. Fast neutrons seem to have a small contribution (<1%) to the effective dose [13]. Neutrons do have a higher RBE (Relative Biological Effectiveness) than protons, and they are linked to induced secondary cancers. Brenner and Hall [14] argue that since most neutrons are created in passive beam shaping objects, any adverse effect is easily remedied by using better accelerator designs and active scanning techniques.

5. δ-rays.

   J.J. Thomson noticed that the recoil of target nuclei can cause secondary ionisation. He referred to this radiation as δ-rays, in the spirit of the other three types known at the time. These δ-rays, electrons, have a typical energy in the range 10 to 500 eV, and don't travel more than a few millimeters. The number of such δ-rays depends greatly on the particle beam species. It is well known that heavier particles have a bigger effect on target nuclei and therefore generate much more δ-rays than protons (figure 3.7). The primary ions also produce δ-rays directly. These δ-rays have a great disruptive power by directly or indirectly causing a double strand break in DNA. The δ-rays generate many free radicals that increase the double strand breaking tremendously [6]. In dose calculation the effect of the free radicals is often modelled as a multiplication factor in the radial dose distribution, which in turn relates damage to distance from Bragg peak [15]. Tumorous tissue has less regenerative capacity than healthy tissue, because cell division is increased and that is when radiation sensitivity is the highest. Slightly mitigating the positive effect is that normal tissue is more sensitive to free radicals than tumors, due to the typically low concentrations of oxygen in tumors. A double strand break often

results in the host cell being unable to repair itself and die, which is why this effect is an important consideration.

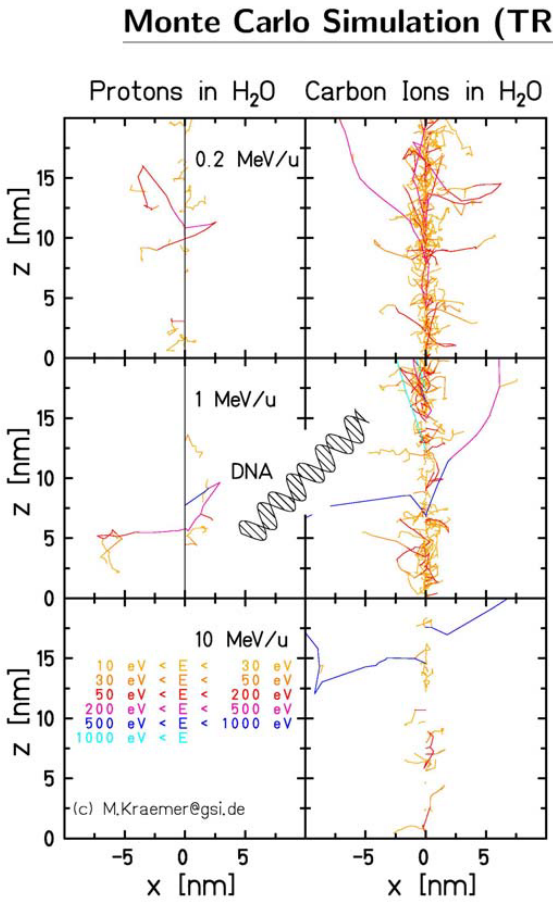

Figure 3.7: Carbon ions and protons have a very different ionisation potential along their paths. Because carbon has much densely ionized clusters, the disrupted volume reaches the dimensions of a DNA strand, which then can be broken clean in two, greatly disabling the host cancer cells ability to recuperate. Image courtesy of Dieter Schardt.

Another drawback of hadron therapy is shown in figure 3.8. The Bragg-peak delivers its highest concentration of damage in a volume of a few millimeters cubed. A tumor that manages to be discovered is typically quite a bit larger than that. In a treatment plan the volume of the tumor is subdivided in voxels of about the size of the Bragg-peak volume of damage. The beam is moved and altered to as to place a peak in all voxels. The unfortunate side effect is that the radiation in tissue before the tumor is a sum of many curves. Compared to the exponential decay of an X-ray curve, this is still an improvement.

Species of hadrons considered for medical purposes are types such as H, He, Li, Be, C and Ne, of which the proton is by far the most used. The upside to heavier hadrons is the increased RBE compared to protons, predominantly due to increased ionization (δ-rays). The downside is fragmentation: the debris stops before or after the Bragg peak and makes the dose deposition less accurate. Target cells can be however killed only once, and ions heavier than oxygen create more damage than what is required. Some research has been done on pion radiation [16, 17] that concluded that the performance has no clear

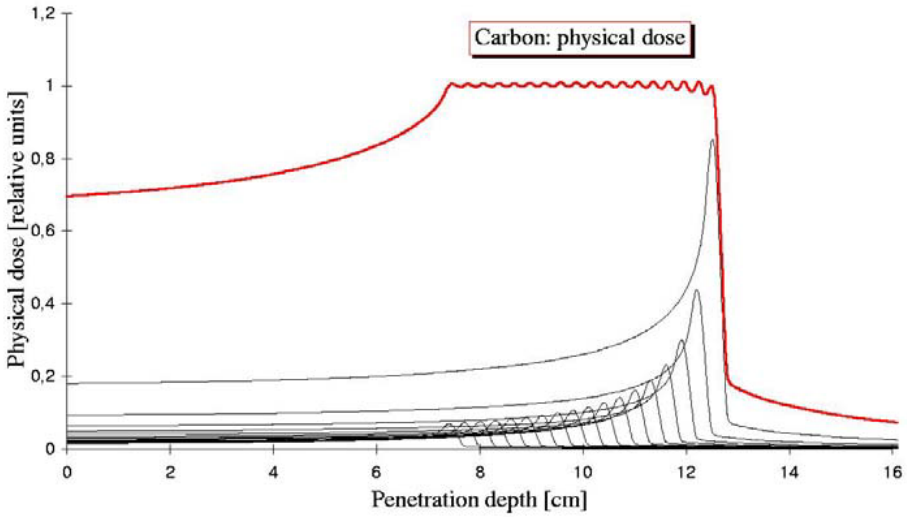

Figure 3.8: In order to irradiate over the tumor volume, multiple hadron energies are summed. This has the unfortunate side effect of an increased dose before the tumor, to healthy tissue. Image courtesy of Dieter Schardt.

advantages nor disadvantages compared to protons or carbon ions. Even less research has been performed on antiprotons which is summarized by Amaldi and Kraft [18] as having 'non-obvious advantages'. Both have one clear disadvantage: the production of these exotic particles is costly compared to normal matter. Without going into detail, the biological effects of different types and energies of particles are rated with an RBE rating, which converts deposited dose into effective dose. This rating depends on many factors, and is used to quantify the specific characteristics of a certain type of radiation into a single number that can be easily used in treatment plans.

Taking these considerations into account, on paper hadron therapy compares favorably to X-ray therapy. A comparison of treatment plans may be seen in figure 3.9. The advantage of less dose before the tumor is also used to reduce the number of directions the beam is pointed. Changing beam position (or patient position) consumes time, so this time is saved. Two excellent papers that cover the current state of affairs in much more detail are Rossi [19] and Schardt and Elsässer [6]. The former includes the specific details of the construction of a facility, while the latter mostly deals with technical and biological considerations.

3.3 Imaging and Computed Tomography

Previously, we saw that imaging and treatment have distinct goals and distinct demands on the type of radiation used. We've also established that the X-ray is currently the default modality for both imaging and treatment. In treatment, as well as diagnosis, an accurate view of the patients internals is required, and this view is obtained with γCT. This method makes use of the fact that different tissues have different X-rays absorbtion coefficients. Imaging is a collective term for any and all types of images, including simple 'photographs'. A CT scan composites multiple flat images into a 3D model, in which

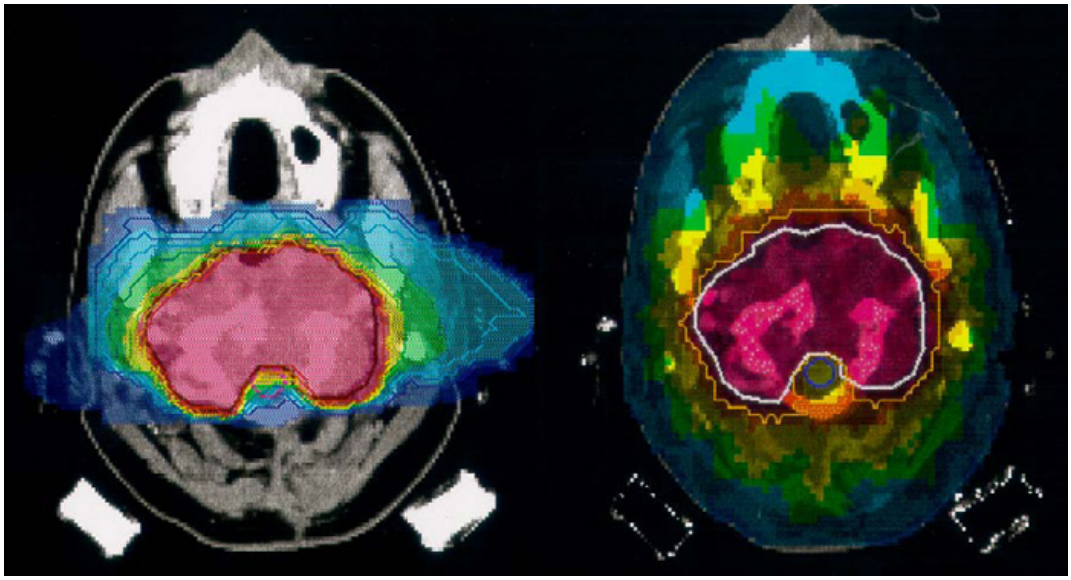

Figure 3.9: On the left a carbon irradiation using 2 fields. On the right a treatment plan for an X-ray irradiation using 9 fields. Using fewer fields the carbon plan manages a much more confined dose delivery. Note that in addition the carbon plan deposits almost no dose in the optical nerves, which is a well known type of tissue extremely sensitive to radiation. Source: Amaldi and Kraft [18].

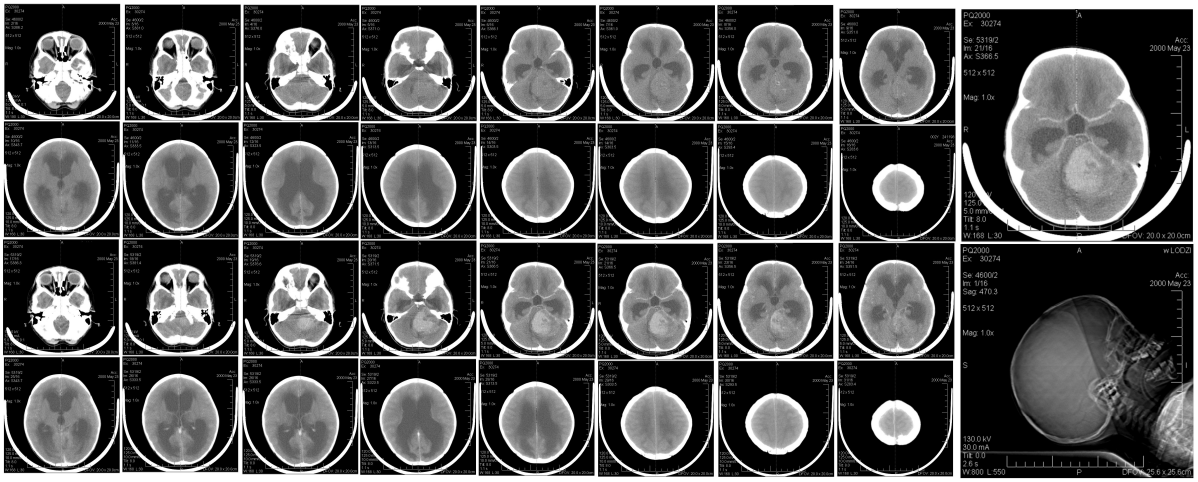

Figure 3.10: An example of a X-ray CT scan [20].

certain tissues are identified and optionally highlighted and viewed in the direction that is desired (figure 3.10).

Using the fact that the intensity loss of the X-ray beam can be calculated from the known intensity of the beam and the intensity recorded in the camera, an image of (a region of) the body can be created by taking such pictures from all around the patient. Godfrey Hounsfield received the Nobel Prize in 1979 for having this idea, and the scale used in radiology to measure the radiodensity is referred to as the Hounsfield-scale with unit HU. Equation (3.2) shows how a HU relates to radiodensity (expressed as $\mu()$) which in turn relates to the energy loss of the X-rays by way of Lambert-Beer's law.

$$HU = [[\mu(X) - \mu(water)]/\mu(water)] * 1000 \tag{3.2}$$

To give a feel for these numbers: air is -1000, bone +700, and fat, water, blood and muscle are respectively -84, 0, +30 to +45, and +40. The major drawback of γCT is that many types of soft tissue, of which the body is primarily composed, have very similar Hounsfield values, and therefore are hard to distinguish. This distinction is not only hard to make visually, but also when creating the best treatment program. Sparing the surrounding healthy tissue is difficult when the physical limits of the tumor are not well known. The radiodensity for hadrons is usually expressed as stopping power. Hanson et al. [21] wrote a pioneering comparison of the two modalities, and he expresses radiodensity in terms of linear stopping power compared to water for protons and the X-ray attenuation coefficient for the γ modality. He also shows that for a given contract ratio, protons incur an order of magnitude lower dose on the body while imaging, compared to X-rays.

In the introduction to this chapter the use of tracers in PET scans was briefly mentioned. The imprecision of γCT is typically mitigated with tracers as well; high contrast materials that are administered to the region of interest. The disadvantage of tracers is that they require time and effort to administer, it is not trivial to place them precisely at the region of interest, and that metabolic processes in the body move and disperse the tracer. Particle beams introduce two alternative options: In-Beam PET and Hadron CT.

- In-Beam PET.

  Particle irradiation can introduce β^+ sources in the body. These can be produced in fragmentation of particles due to the beam, or by using β^+ emitters as the beam particle. Amaldi and Braccini [8], Schardt and Elsässer [6] mention that at GSI (Gesellschaft für SchwerIonenforschung) ^{11}C and ^{15}O were used to this extent. The article mentions further that the detection of 'prompt' gammas is particularly interesting. These prompt gammas are created in the interaction of the incoming particle and bodily material, and originate therefore certainly at a point of damage. If the emitter is the incident particle, or a secondary emitter generated by an interaction, there is a certain decay time, which means the particle will have moved due to metabolic processes and thereby decrease precision. The main benefit is therefore that there is no separate tracer needed, and that imaging is a 'free' bonus to therapy. Imaging is perhaps an overstatement: this method makes the irradiated volumes directly visible and can be used as a check of a treatment plan but not for diagnosis. The downside is that the emissions of photons have no preferred direction, so a detector should enclose all angles, which in turn is expensive and impractical. If the detector is in another room, prompt emissions are not detectable at all. The decay times of ^{11}C and ^{15}O are respectively about 20 and

2 minutes, which may be longer than the time it takes to move the patient to the imaging room. An extra complication is that the treatment itself has a duration of this order, which means that there will be a significant difference in activity between regions irradiated at the start and the end of the treatment.

- Hadron CT.

  The hadrons themselves could be tracked and absorbed in order to create an image. Using energies higher than for therapy, placing the Bragg peak behind the body, the ionizing damage could be minimized further. Apart from sparing tissue compared to X-ray irradiation (figure 3.1), there is a second benefit to hadrons for imaging: an increased density resolution in comparison to X-rays. Figure 3.11 shows how Ryu et al. [22] makes that clearly visible. More contrast is seen between the different types of material, and, using different calibration curves, the contrast may be adjusted as a doctor sees fit. Also visible is the difference in sharpness: because protons scatter much more than photons, the precision in the plane perpendicular to the beam is lower. For imaging, it seems hadron-CT could augment the γCT in the depth dimension. Moreover, conversion of radiodensity of the different modalities (γCT Hounsfield numbers versus hadron stopping powers) is a relatively imprecise method, introducing an error of up to 3% [23]. Using hadron data sidesteps this problem completely. Producing usable hadron images therefore is desirable even if they will not be used for diagnosis.

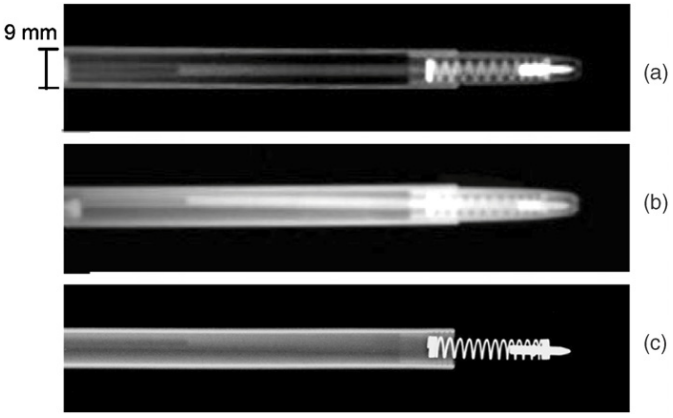

Figure 3.11: The top two images show proton beam images for two different dose-depth calibration curves. The bottom image shows a standard X-ray scan. Note the missing features in the X-ray image and the difference in sharpness. Image courtesy of Ryu et al. [22].

Going from simple 2D projections to a fully reconstructed CT scan is a nontrivial matter, so let us start with the simple case of the 2D image. For every pixel one residual intensity value is measured, and therefore the opacity of the body is summed over all tissue behind that pixel. The spatial distribution in the depth direction is therefore lost. An accurate 3D model reconstruction would aid in the establishment of opacity, and thus radiodensity, for each volume element. The radiodensity is tissue-specific, so full knowledge of spatial distribution in the depth direction would be reconstructed and the structure of the the patients body would be known. With conventional X-ray imaging, the common geometry in γCT is a camera rotating about an axis parallel to the length of the patient. A series of 1D line integrals are obtained by moving the camera around the patient. As the camera

moves along the body, new sets of 1D line integrals are obtained, each corresponding to a slice of tissue: tomography. The line integral of the absorption is now known for many different positions and angles, and the reconstruction software reconstructs for each voxel in a slice the local absorption coefficient, resulting in a 2D image: computed tomography. These typical 2D 30cm squared images are what most doctors are used to. As available computational power increased, the 3D reconstruction of the body in the detection volume became possible. For hadron CT, or pCT, as protons are the most commonly used hadrons, rotating the camera and the beam around the patient requires an expensive gantry. An extra constraint for pCT could therefore be that algorithms work well without a gantry; it can cope with images from a reduced number of angles. One could consider changing the position of the patient table instead of the gantry.

3.4 Modalities in Medicine

This chapter has given the reader a flavor of the field of hadron treatment and imaging. To briefly summarize the benefits of hadrons as an *imaging* modality is that it is sensitive to structure in the dimension that γCT is not sensitive, so a potential combination of the two seems promising. Secondly, an inherent advantage for hadron imaging is less ionizing damage incurred on the body. Lastly, it is preferable to image in the same modality as the treatment to sidestep radiodensity conversions for an increased accuracy for dose deposition. The following table is supposed to summarize the advantages and disadvantages of the different modalities.

Table 3.1: Comparison of Modalities

Modality	**Ease of use**	**Imaging**	**Treatment**
γ	Cheap to generate	Easy to record	Can penetrate whole body
β	Cheap to generate	Limited range Extreme scattering because of atomic electrons	Limited range
Hadron	Expensive to generate	Hard to record Good density resolution	Accuracy of dose deposition at specific depth
Neutron	–	–	Ill defined dose distribution
Pion	Very expensive to generate	–	Instability with undesirable decay products

CHAPTER

FOUR

HADRON IMAGING

My research focuses on a new medical hadron imaging setup. This section goes into detail on some of the other approaches to hadron imaging that are currently conducted. At the end of the chapter our approach will be briefly summarized and compared to the competition. The next chapter will treat our setup at length.

4.1 Nuclear Emulsion Film Detection

The oldest method of proton imaging [7] is still of interest to some. Cost and ease of use are cited as advantages. Braccini et al. [27] proposes a setup where a "brick" of interleaved absorbers and nuclear emulsion film is used to obtain an image, akin to a sampling calorimeter found in some high energy physics experiments. The brick is made up of two parts: first the proton will go through 30 emulsive sheets interleaved with 1.5mm thick absorber plates. Second the proton goes across 0.5mm absorbers interleaved with 40 emulsive plates. The beam will be tune to deposit the Bragg peak in the second part, which has increased precision compared to the first.

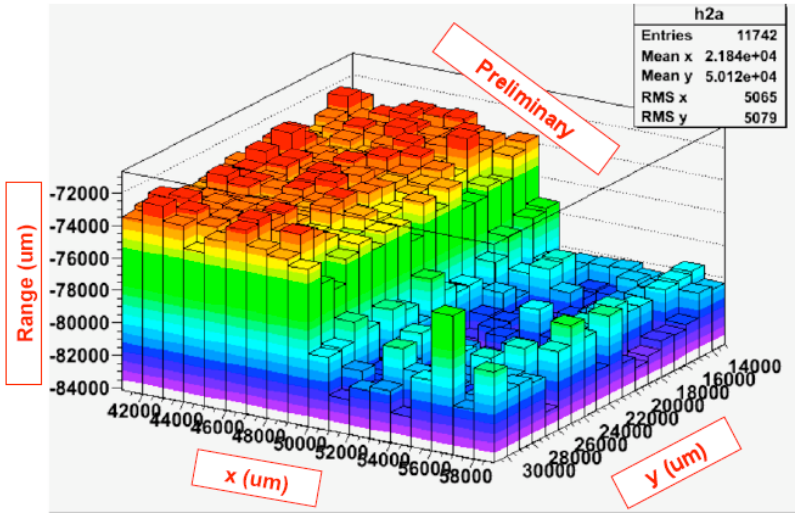

Figure 4.1: Reconstructed image of a stepped PMMA block. A clear distinction between the 1 and 4cm thickness can be made. [27]

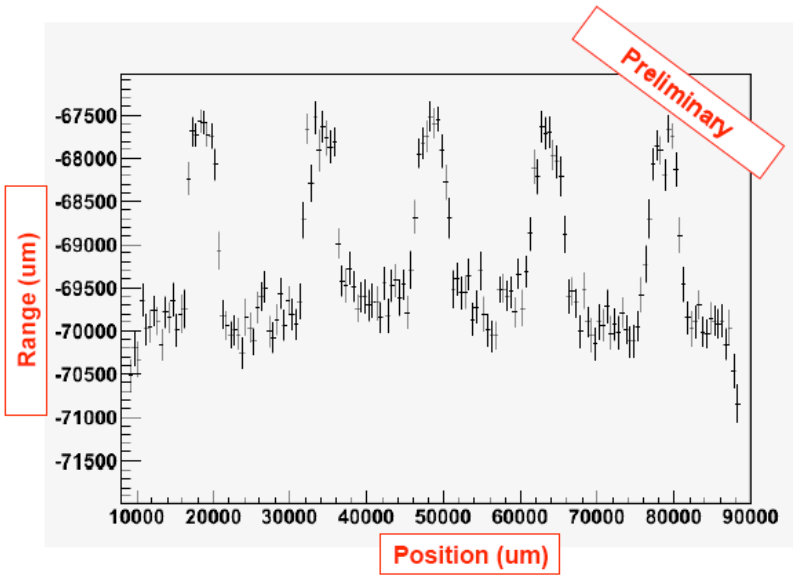

Figure 4.2: Reconstructed image of the PMMA block with aluminum rods. [27]

The paper describes the Monte Carlo simulation of the setup and the results of a test run at the Paul Scherrer Institute in Switzerland. 10^6 protons per $10cm^2$ are required for a complete image, and a low intensity beam is used to spot-scan the object. Two phantoms are scanned: an object of homogeneous substance (PMMA) with a thickness of 4cm on one side and 1cm on the other. The second object is made up of a block of the same material but with five embedded aluminum rods.

While this particular research appears to be in its initial stages, the reconstructed images in figure 4.1 and figure 4.2 are convincing. However, no numbers for position or energy resolution are provided.

4.2 Strip Detectors

Another detection paradigm found in high energy physics is the strip detector. Johnson et al. [28] consider their use for proton imaging. Four plates with 400μm thick silicon strip detectors (SSDs) are used for both tracking and calorimetry. Two groups of two plates at an angle of 90 degrees are placed some space apart to provide two points for the track. One 2-plate plane is placed directly after the phantom and the other 25cm farther away. The time over threshold (TOT) of a signal in a strip is used to give energy loss information. Given the known dimensions and composition of the strip, there is a correlation between deposited charge and signal, although the resolution is not quite as high as dedicated calorimeters.

The setup is put to an experimental test, as well as simulated using Geant. A phantom consisting of a hollow aluminium cylinder is successfully imaged (figure 4.3). It is noted that the group of SSD plates farthest removed from the phantom provides no discernible features of the cylinder, and was therefore not used. A cross-section of the cylinder at its widest is used to obtain some more specifics of the measurement. The number of proton events per pixel varies from approximately 100 to 500, owing to the non-uniformity of the

beam. Around the edges of the aluminium the reconstructed image is fuzzier than the Monte Carlo of the setup written with Geant4. This discrepancy is attributed to both a calibration accuracy of 15% in the energy conversion electronics, to multiple Coulomb scattering of the protons (although the simulation takes this into account), and divergence of the degraded proton beam. A more specialized simulation of the energy distribution of the proton beam indeed shows a larger deviation (RMS) around the interfaces between air and aluminium. The flat response of the plates farthest away is also attributed to the multiple scattering. Although both energies are found in the data, they appear nearly fully decoupled from position.

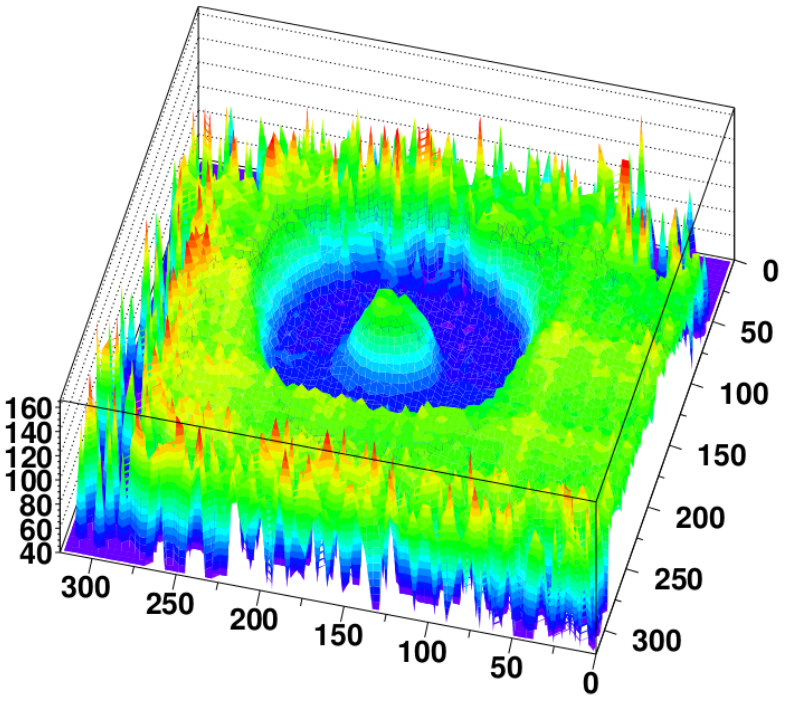

Figure 4.3: Reconstructed aluminium cylinder using silicon strip detectors for both tracking and calorimetry [28].

To improve the position resolution, cuts on exit angles are recommended. Johnson et al. [28] ends with the remark that their setup gives an energy resolution of approximately 20%, "which could mean that the image contrast at acceptable dose levels is insufficient". Exploring different techniques for the measurement of energy is recommended and even promised.

4.3 Strip Detectors and separate Calorimeter

Much research on hadron imaging is conducted in Italy, often in the context of the PRIMA collaboration. Talamonti et al. [29] describe a similar setup as in the previous section, but this time coupled with a separate calorimeter to improve the energy resolution drastically. The calorimeter used was made up of a CsI crystal attached to a PMT. This group uses four strip detector planes, again each made of two strip-plates at an angle of 90 degrees from each other. Two planes were placed just before and after the phantom, and two

between a bit farther up- and downstream. This setup was built to be able to cope with a proton beam frequency of 1MHz.

In addition two phantoms are tested. Both are complex compared to the ones used in the previously mentioned studies, and consist of different materials and complex shapes suitable to a detailed study of detector performance. Phantom "A" is a block of PMMA with holes of varying radius, while phantom "B" is more complex, with finer structures and composed of a large variety of materials. The spatial coordinates obtained with the four planes are used as input for a most likely path (MLP) fit, for which a method proposed by Williams [30] and verified by Cirrone et al. [31] was used.

Again the detection plane farthest removed from the phantom shows very little detail. The four planes coupled with the MLP track reconstruction method did however recover some of the detail present in phantom "A", indicating the MLP method is able to mitigate some multiple scattering. With phantom "B" it becomes clear that the setup is not sensitive to small features yet: a spatial resolution of 1.1mm is calculated. The energy resolution is put at 4%, and among the limiting factors "the number of proton histories" is mentioned. Not enough tracks have been recorded, so the statistics couldn't be better than this.

In a paper published earlier the same year, Cirrone et al. [32] uses a similar setup with YAP:Ce and YAG:Ce calorimeters. Energy resolutions of 3% and 2.7% are found, respectively. Although the paper concludes that further studies will be carried out with the YAG:Ce calorimeter, it is not explained in the following paper on what basis another, third option is chosen.

Civinini et al. [33] also combines strip detectors and a calorimeter. This time a YAG:Ce crystal is used, and a silicon strip thickness of 200μm is mentioned to be a good compromise between scattering and detection efficiency. The energy resolution of the calorimeter is claimed to be about 1%.

4.4 GEMs and Sampled Silicon Calorimeter

Project AQUA, an initiative that tries to advance the research on hadron therapy, makes use another kind of setup, named PRR10 [34]. This group uses a pair of Gas Electron Multipliers (GEMs) to track protons and a stack of thin silicon plates to measure the range and therefore energy of the proton (figure 4.4). The whole device is rather compact, and the stated purpose is real-time imaging with sub mm accuracy and a range accuracy of a "few" percent, using 10^6 protons per image. Currently the data acquisition for the GEMs was limited at 10kHz, but an improved system managing 1MHz is foreseen. Quite a few lines are spent on the data acquisition, and the online and automated calibration is an impressive achievement.

Two phantoms are investigated. The first a relatively simple one, a slab of plexiglass from which some cylinders of varying length and radius are removed, resulting in the left pane in figure 4.5. The second phantom is constructed so as to be more similar to the more homegeneous composition of the human body. It consists of three plexiglass cylinders with varying radius immersed in a volume of water. Its reconstruction is found in the right pane of figure 4.5.

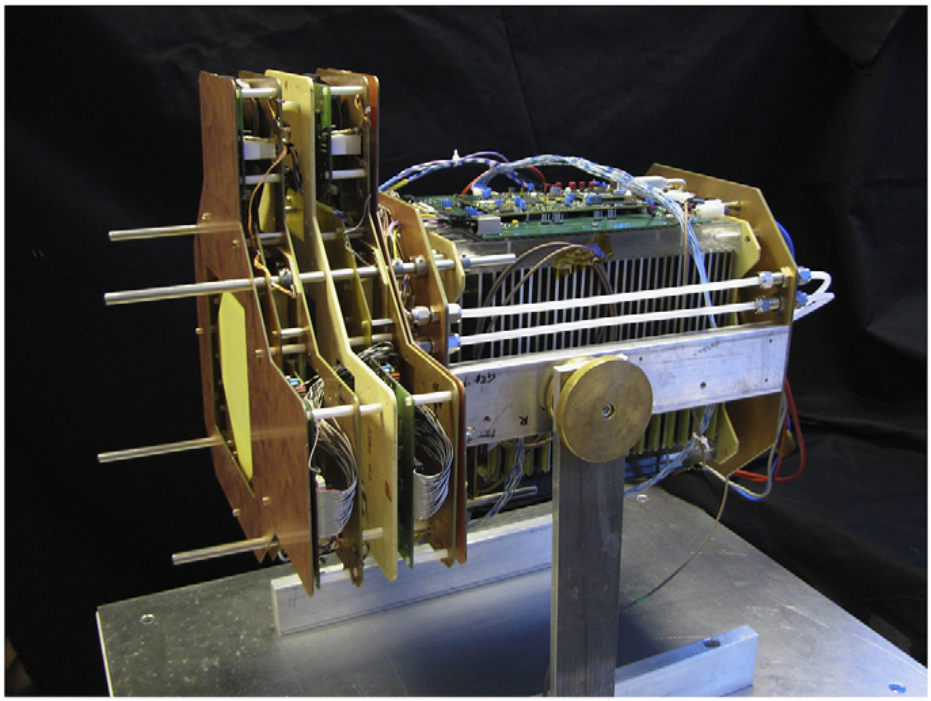

Figure 4.4: Photo of the segmented Aqua setup. [34]

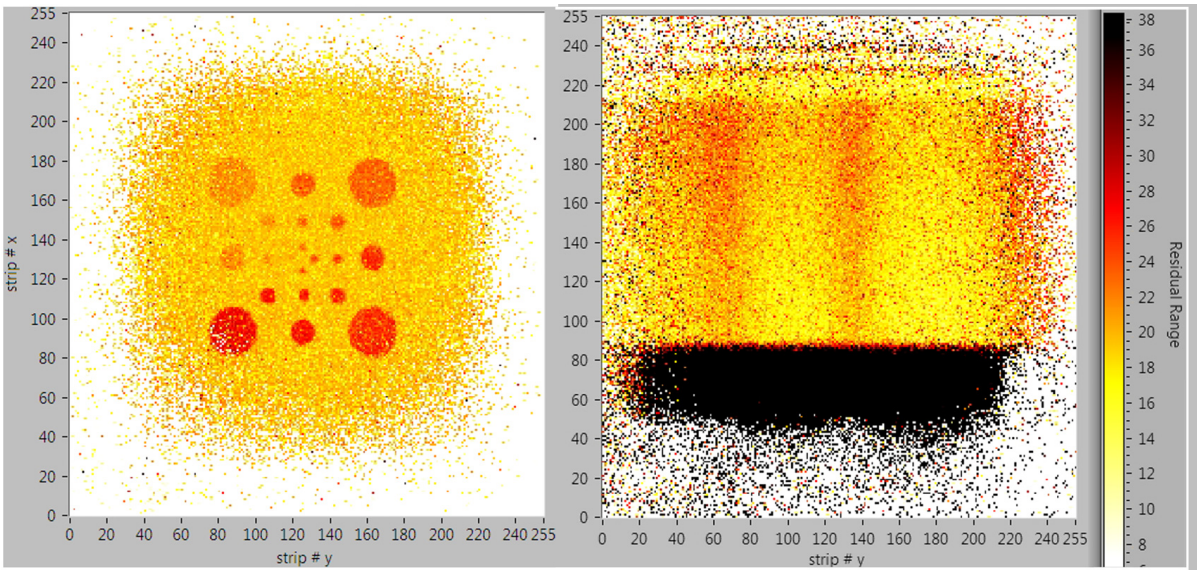

Figure 4.5: On the left, an AQUA reconstruction of a slab of plexiglass with some cylinders of varying length and radius removed. Right: a reconstruction of three plexiglass cylinders with varying radius immersed in a volume of water. [34]

The spread in proton beam spot introduced by copper diffusers is specifically mentioned as an important point of improvement for clinical setups. Active raster scanning methods should improve the energy resolution. Results were obtained and analyzed on-line, and compared to a Monte Carlo off-line. The claimed theoretical accuracies are obtained, although the energy resolution is not elaborated upon further. The paper concludes with one remark in particular, namely that the GEMs prove to be a cheap, low weight alternative to traditional silicon strips and can easily be increased in size.

4.5 Scattered Events Beam Reconstruction

With the argument that objects in the beam scatter beam particles and generate secondary ions, Jakubek et al. [35] in Prague propose a detector placed outside the beam to capture some of these stragglers. The problem of scattering is turned into a detection mechanism, because at the interface of materials with different densities, there is an increase in scattering. Using a stack of 2D silicon detection grids, TimePix chips, a 3D tracker is made. It is proposed that with such a setup the whole beam line can be imaged (figure 4.6), primarily to use for real-time beam verification. An energy measurement is not reliable since energy information is often lost in the scattering. Henriquet et al. [36] suggest that therapy with heavier ions might be monitored with a scatter detector, because heavier ions generate secondary protons especially at the interface between bone and other tissue.

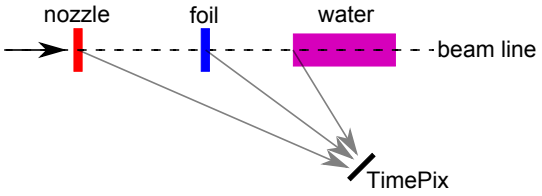

Figure 4.6: The TimePix chip stack records stragglers, predominantly originating from interfaces between materials.

4.6 The Nikhef/KVI approach

From the examples seen above we can conclude that medical hadron imaging is in its early stages. Most approaches are not ready for practical medical use. This means that the recent start at Nikhef for a hadron imager is not a tremendous disadvantage. Currently the purpose for the setup is to lead to a clinical device for use in a future proton treatment facility in Groningen, where the KVI is closely involved. To circumvent the conversion from Hounsfield number to proton stopping power, a proton image is desired. A feasible time schedule for a single treated patient session is about 30 minutes, of which the total treatment time is about 2 minutes. Similarly, imaging should take up not more than 2 minutes to be practical, and would be a last check on the morphology of the patient. It will depend on the final design if the treatment plan allows for last minute adjustments: in the case of passive beam shaping, the construction of a new bolus would mean that treatment must be postponed. Considering how large a typical tumor can be, let's say

10cm cubed, the viewing window, or detector size, should be of similar size. That said, our current setup has a GridPix tracker and a BaF_2 calorimeter.

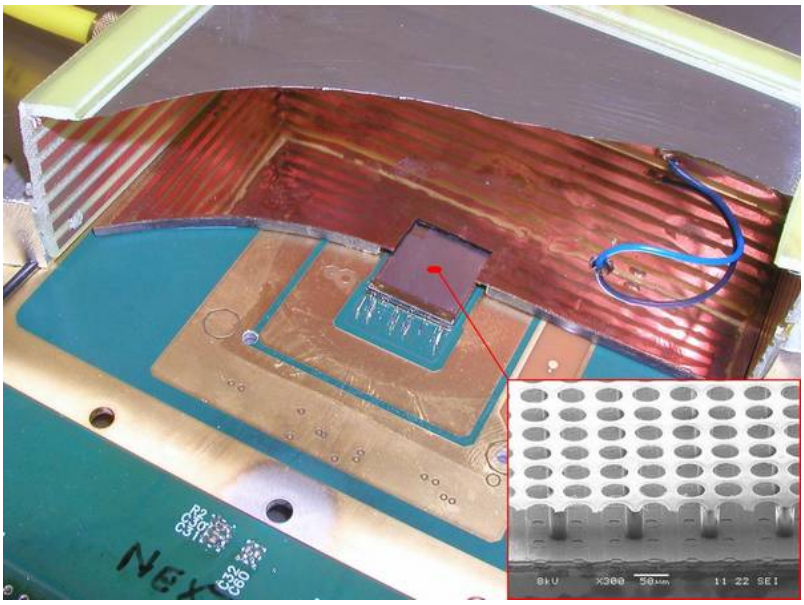

Figure 4.7: A GridPix cut-away view. In the center is the chip, surrounded by a drift cage, containing the drift gas. The inset shows a magnified view of the Ingrid on top of the TimePix chip. The drift cage used in our setup was smaller than what is shown here: at first 1.5cm cubed and later replaced with a 4 by 4 by 1.5 cm^3 volume.

- The GridPix consists of a small gas container attached to a GridPix chip. The GridPix chip is a TimePix chip, with the addition of an Ingrid. An Ingrid is a 1 μm thin perforated aluminium foil suspended 50 to 75 μm above the silicon by a nonconductive spacer (figure 4.7). Two HV potentials are applied, one between the top of the drift cage and the grid, and the other between the grid and the chip, which is grounded. The gas container volume above the grid function as a drift volume, where ionizing particles will break up gas molecules into ions and electrons [37]. These will drift because of the potential. When electrons arrive at the grid, they enter the avalanche region. Between the grid and the chip the field is kept at $10^4 V/m$, which causes the electrons to accelerate. As their speed increases, so does the chance of ionizing gas particles and so freeing more electrons. This multiplication effect results in an avalanche of electrons directed towards the chip. This amplified charge then induces a charge in the chip, which is then registered as a hit. The TimePix chip consists of 256x256 pixels, each 55 μm squared. Each pixel is able to measure the Time of Arrival (ToA), which means that given a certain shutter signal, it can measure the time between the start of the shutter and the time of the event (for each pixel). Given a known drift velocity, the ToA for each pixel can be converted to a height and a 3D image of a passing charged particle can be reconstructed. The chip is digitally connected to a standard computer, where Pixelman and RelaxDAQ software are used to configure, calibrate and read out the chip. More detail on GridPix and TimePix may be found in Tsopelas [38], but the main advantage over SSDs is that gas scatters much less than silicon.
- The calorimeter is much simpler, and consists of a PMT (Photo Multiplier Tube)

attached to a BaF_2 crystal, that generates an amount of light with a proportional relation to the absorbed energy of an incoming charged particle. The PMT registers this light, after which an ADC (Analog to Digital Converter) converts it to a digital signal. We use an ADC borrowed from the HiSparc project, which is a small box with a two channel 400MHz sampler, connected to a computer via USB.

- The trigger system is built with scintillators roughly the size of a side of the drift cube. This fast scintillating material ensures as little delay as possible. The trigger logic is constructed such that the tracker and calorimeter operate in acquisition mode by default. A signal from the trigger system stops the measurement, writes data to the computer, and resumes acquisition. This way we hope to achieve as little dead time as possible. In a future iteration the particle accelerator itself may provide the trigger.

Typical considerations for the eventual use of the setup are:

- How many protons do we need for an image?
- What is the event acquisition rate that is required?
- How large is the beam spot?
- What is the required beam current?
- How many tracks can be correlated (before and after patient)?
- Scattering in strip detectors versus GridPix: is GridPix advantageous?

These points are interdependent to a certain extent, but it is still informative to see what are key points to think about when constructing the setup. As for size we are not quite there yet: our drift volume and therefore viewing window is 1.5cm squared. Our aim is to accommodate the desired clinical use case as optimal as possible. In the next chapter more detail will be given on the detectors, in particular some of the current deficiencies and attempts at improvement. Data obtained from test runs at the KVI accelerator and from laboratory sources are discussed to demonstrate the performance.

CHAPTER

FIVE

IMPROVING THE SETUP

The first months of my time spent on the Nikhef/KVI Proton Radiography project overlapped with the previous student working on it, Tsopelas [38]. The first test run with protons was in November 2011 at the AGOR (Accelerateur Groningen-ORsay) cyclotron at KVI. The results from this test run are the basis of my work on the project. Any problems or issues were the things I would work on and improve in my time at Nikhef, which will be discussed in this chapter.

5.1 Results from November and Overview

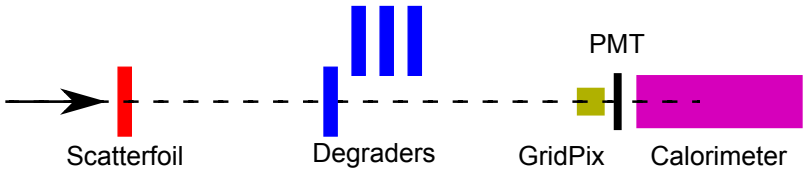

Figure 5.1: A schematic of the proton detector setup as installed at AGOR. From left to right: the beam enters the experimentation room at 150MeV. A lead scatterfoil of 1.44mm thickness broadens the beam and reduces the energy to 145.5MeV. In total the beam passes through about 3m air, reducing the energy down to 143MeV. The aluminium degraders can be remotely moved into or out of the beam. On the right hand side the GridPix tracker, the trigger scintillator (PMT) and the BaF_2 calorimeter are represented.

Apart from establishing the efficacy of the setup in the situation of a proton accelerator, the performance of the GridPix was profiled and the calorimeter calibrated as well as profiled. The main objective of the endeavor is to image objects. To test this, we collected data in the following configurations:

- Changing the full beam energy with aluminium absorbers in order to provide an energy calibration for the calorimeter. The experimentation room provided remotely movable aluminium plates, which we varied between between 20, 40, 57 and 61mm, corresponding to an energy (at the detector) of 118, 88.3, 54.8 and 44.4MeV.
- The aim of the setup is to reconstruct objects in the beam, so a very simple object was placed in the beam: up to four copper plates covering half of the detector. The four plates were each 2.7875mm thick.

- Two copper wedge shapes with different angles provided a more complex test of our ability to reconstruct objects. The dimensions were 30x30x12mm and 30x30x24mm.

A thorough analysis of the GridPix, and a preliminary analysis of the calorimeter during these runs may be found in Tsopelas [38]. We attained an event acquisition rate of about 2Hz, due to a frequent loss of synchronicity between the tracker and calorimeter acquisition systems.

The calorimeter performance in particular presented us with some challenges, as we found discrepancies between the runs. Using the aluminium absorbers we intended to calibrate the energy response of the calorimeter. Also, the full beam should be calibrated precisely, because the energy deposited in the object is calculated by subtracting the calorimeter energy from the full beam. Throughout the day the full beam (143.6MeV) was measured with the calorimeter, but unfortunately the measurements did not always correspond with each other. Tsopelas mentioned that a possible cause could be temperature, but after further study it was found that the full beam calorimeter response variation correlates well with changes in proton rate. The AGOR proton beam operates at a fixed 55MHz frequency, but allows tuning the number of protons per bunch. Because initially synchronicity was lost after only a few events, the effective proton rate was lowered to about 0.5 - 1 kHz, meaning that most bunches would be empty, and the nonempty bunches would be filled according to some Poisson distribution. After discussion with dr. Tjeerd Ketel, a Nikhef calorimeter expert, the rate sensitivity of the calorimeter was deemed the most probable cause of the shifts in the calorimeter response. The shifts were found to be the same for peak position and peak integral. To reconstruct the wedges, the data taken from the copper plate runs were therefore used to calibrate the calorimeter response, as these runs were performed at a similar proton rate. The calibration curve is seen in figure 5.3.

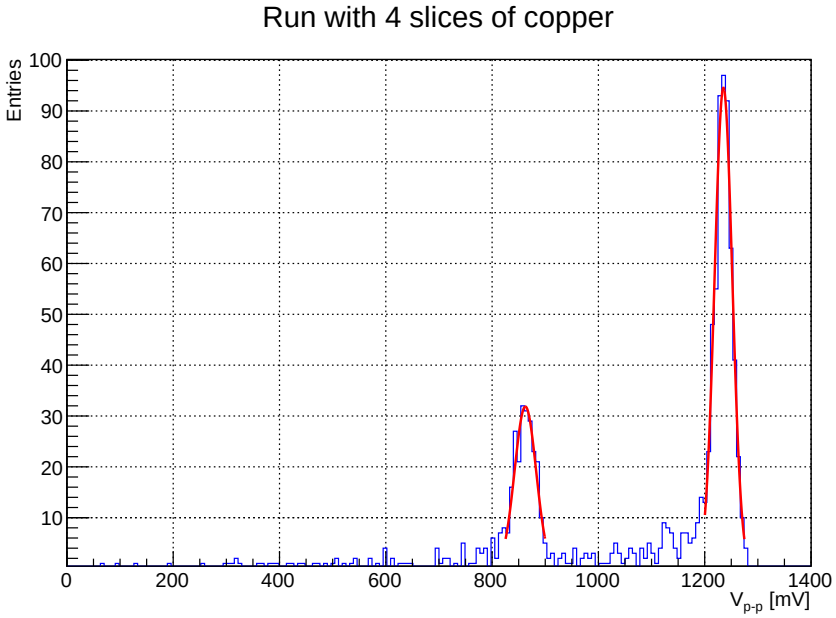

Figure 5.2: The calorimeter response for the 4 copper slice run. Two peaks are visible, the left corresponding to the protons that went through the copper and the right corresponding to the full beam. Both peaks are fitted with a gaussian, and correspond to points in figure 5.3.

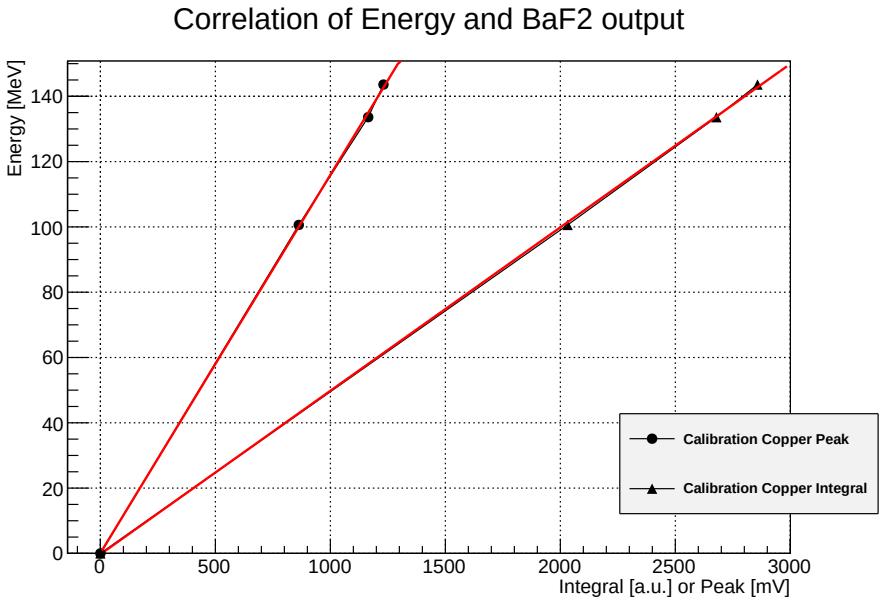

Figure 5.3: By fitting a gaussian over a histogram of all data points with 1 plate of copper, 4 plates of copper, and full beam, the calorimeter response for that particular energy was obtained (see figure 5.2). The fits were tested for stability by altering the range to which the gaussian would fit. This fit is only accurate for a capture rate of ~1kHz. The fit is linear, as expected.

About 30 tracks for each wedge each were selected by hand, because an algorithm for automated event selection was not available at the time. Event selection is required, because we have seen problematic events such as double tracks, δ-electrons in the drift volume, and curved tracks. For a proper analysis, these need to be filtered, and either fitted differently or be discarded. For now, some events were selected where the projection of the track in the plane transverse to the beam was a clean and compact blob: straight single proton tracks without complications. Using the NIST tables [5] for protons in copper, the wedges were reconstructed. Fitting the data was done using both peak position and peak integral, and both with a free first order polynomial and one with a fixed slope, based upon the known dimensions of the wedges. The peak position fitted data agrees best with the known slope.

These results (figures 5.4, 5.5, 5.6), together with the analyses of the GridPix and calorimeter, prove the concept works. The limited event set, due to the lack of event selection, precludes us from doing a meaningful statistical analysis, so error bars are omitted for now. The main issues discovered are loss of synchronicity, low event acquisition rate, a small loss of hits up to 1mm above the grid, supposedly due to trigger delays (figure 5.7 in [38]), and an unknown sensitivity to rate of the calorimeter. Most of my time was spent in solving these problems and it was decided to focus on four points:

- An improved calorimeter, with more stability. A small literature study is performed to establish whether or not such a calorimeter exists and is available for purchase.
- A new trigger design that can keep synchronicity better and perhaps solve the missing hits 1mm above the grid. The GridPix readout system, RelaxD, with the associated software, RelaxDAQ, has a readout limit of 120Hz. A trigger system

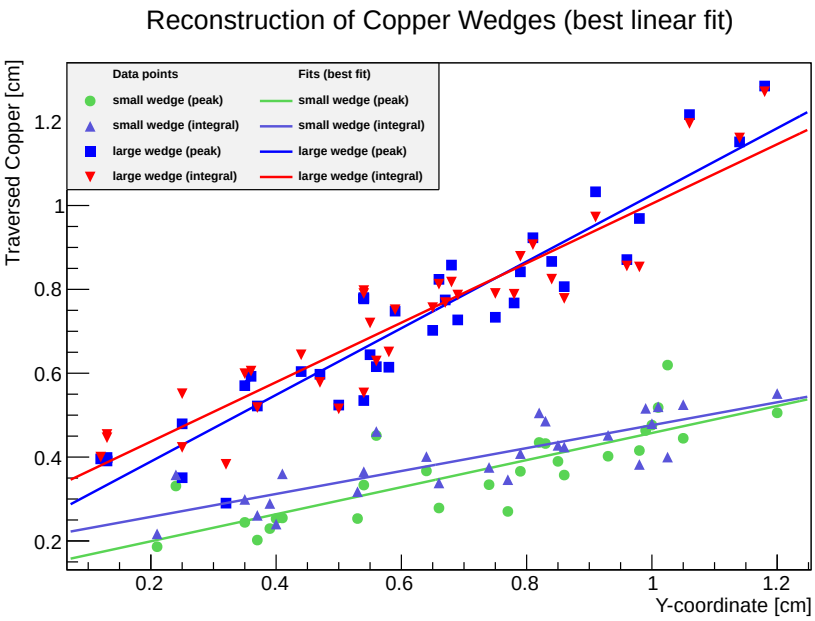

Figure 5.4: The analysis of the reconstruction of the wedges. The plotted points are traversed copper paths as function of transverse coordinate. These are reconstructions using either the integral of a calorimeter response curve of simply its peak height. Also plotted are the best linear fits. It appears the integral derived wedge reconstruction has a less steep slope than the peak derived wedge reconstruction. The slopes are for the large wedge 0.79 (peak) and 0.71 (integral), and for the small wedge 0.32 (peak) and 0.27 (integral).

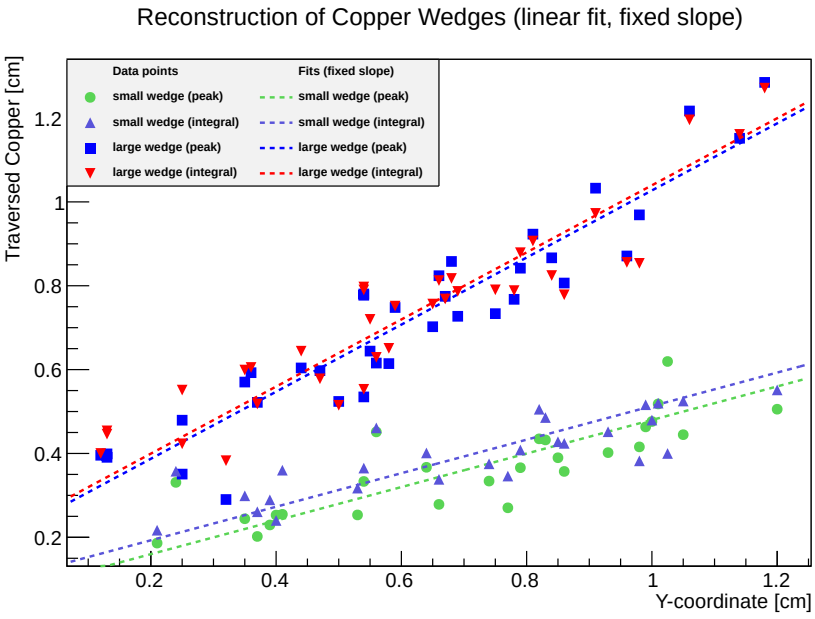

Figure 5.5: The analysis of the reconstruction of the wedges. The plotted points are traversed copper paths as function of transverse coordinate. These are reconstructions using either the integral of a calorimeter response curve of simply its peak height. Here fits with a fixed slope are plotted, taken from the actual slope of the wedges. Now the integral derived reconstruction results in slightly thicker wedges. The actual slopes are 0.40 and 0.80 for the small and large wedges respectively.

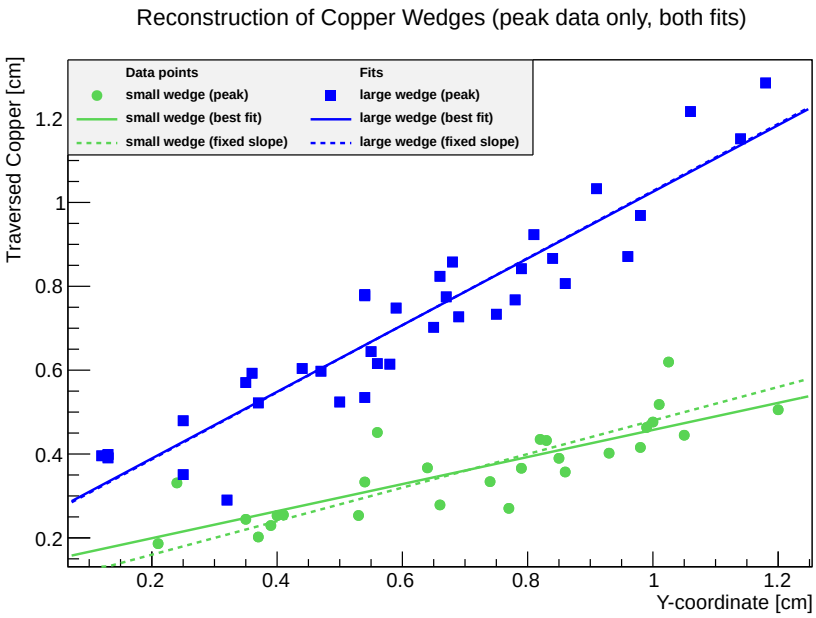

Figure 5.6: The analysis of the reconstruction of the wedges. Only the peak reconstructed points are shown. Straight lines are best linear fits, dotted lines are fits with the slope set to the known value. The slope of the free fit corresponds best with the known slope for the peak derived points. Only for the thin wedge are the slopes different between the two fits. The large wedge fitted slope (0.79) agrees very well with the real slope (0.80), while the small wedge fitted slope (0.32) diverges a bit from the actual value (0.40).

that can keep up with such a rate is therefore investigated. Usage of the RelaxD busy signal is planned so that dead time during readout is minimized further.

- Timestamps on both the tracking data and calorimeter data would help, because in case one system missed a trigger and the other not, the timestamps could be used to filter the incomplete events and still allow the remaining events to be used. Although not a must-have, it would make the data acquisition more fault tolerant. Also, new software for the RelaxDAQ readout already provides nanosecond timestamping.
- An improved data acquisition system for the calorimeter. We had been using a digital oscilloscope, which due to limited buffers could only store a small number of events, resulting in very short data taking runs. Nanosecond precision timestamping would provide an independent method of ensuring synchronization.
- An increased number of GridPix chips, so as to increase the drift volume and thereby the tracking volume. In addition, to improve the fitting accuracy and therefore object reconstruction performance, it would be advantageous to have separate trackers before and after the object.
- Event selection, so that the analysis and statistics can be improved. Problematic tracks should be filtered and handled separately.

5.2 Towards pCT

Our 1.5cm cubed tracker does not live up to the demands of medical professionals quite yet. As mentioned in *chapter 3*, recreating the typical 30cm squared projection would require us to scan over that area in some fashion, or keep the tracker aligned with the beam. Because this is time consuming or not very practical, a larger tracker is the desired alternative. Another improvement is to add a second tracking plane. As seen in *chapter 4*, most setups feature multiple tracking planes. Since the GridPix tracks in 3D, there is no reason to have multiple planes close to each other. Far apart is however again advantageous, in particular before and after the patient or object, to correct for scattering inside the imaging volume. For our setup this means that we want to put a number of GridPixes alongside each other inside the same gas volume, which then would not be cubical anymore but extended in the transverse dimension. This 'slab' should be repeated, once before and once after the patient. The increased tracking volume results in a larger detection area and thus a faster imaging process. Unfortunately, our GridPix broke down, so a new one had to be constructed. A new batch of TimePix chips that arrived at the Detector R&D group proved to be of lesser quality, and most had to be discarded, which meant that the multi-GridPix was not possible at this time. One single-chip GridPix replacement was created, with a gas volume with the same height as before but wider dimensions around the chip. In the sides of the GridPix casing metallic leads serve as field shapers. Close to the field shapers the field is not quite uniform, so putting the walls of the gas volume outwards improves the fields uniformity above the chip, but keeps the effective drift volume the same. The gas volume changed from about 1.5cm cubed to 3cm x 3cm x 1.5cm. The new TimePix chip has been confirmed to work, but because of the bad batch was assumed to be sensitive to strong HV fields. Previously a $CO_2 - DME$ (50%-50%) mixture had been used, which is known to provide a reasonable middle ground between drift velocity and diffusion. It requires a comparatively high potential between anode and cathode to ionize however, so another gas mixture was selected that allows for lower potentials at the cost of more diffusion: Helium-Isobutane (80%-20%). The minimum track-generating drift potential was experimentally obtained by placing a beta emitter close to the GridPix and increasing the potential until tracks became visible. Using Magboltz, a CERN software package to simulate the transport of electrons in gas mixtures, the drift velocity was obtained.

5.2.1 Fitting curves

Fitting curves to the recorded proton tracks is a part of the analysis not yet performed. The procedure can be broken up in two parts: fitting tracks within the GridPix drift volume and fitting through the imaged object.

- Fitting GridPix.

  In the November test run, the GridPix was aligned with the beam line, and any tracks that were not straight were discarded. This resulted in a spot on the projection on the transverse plane, which circumvented any fitting difficulties. A proper analysis, especially when scattered, non-aligned tracks are reconstructed, requires a robust fitting method. Within the Detector R&D group, the master thesis of De

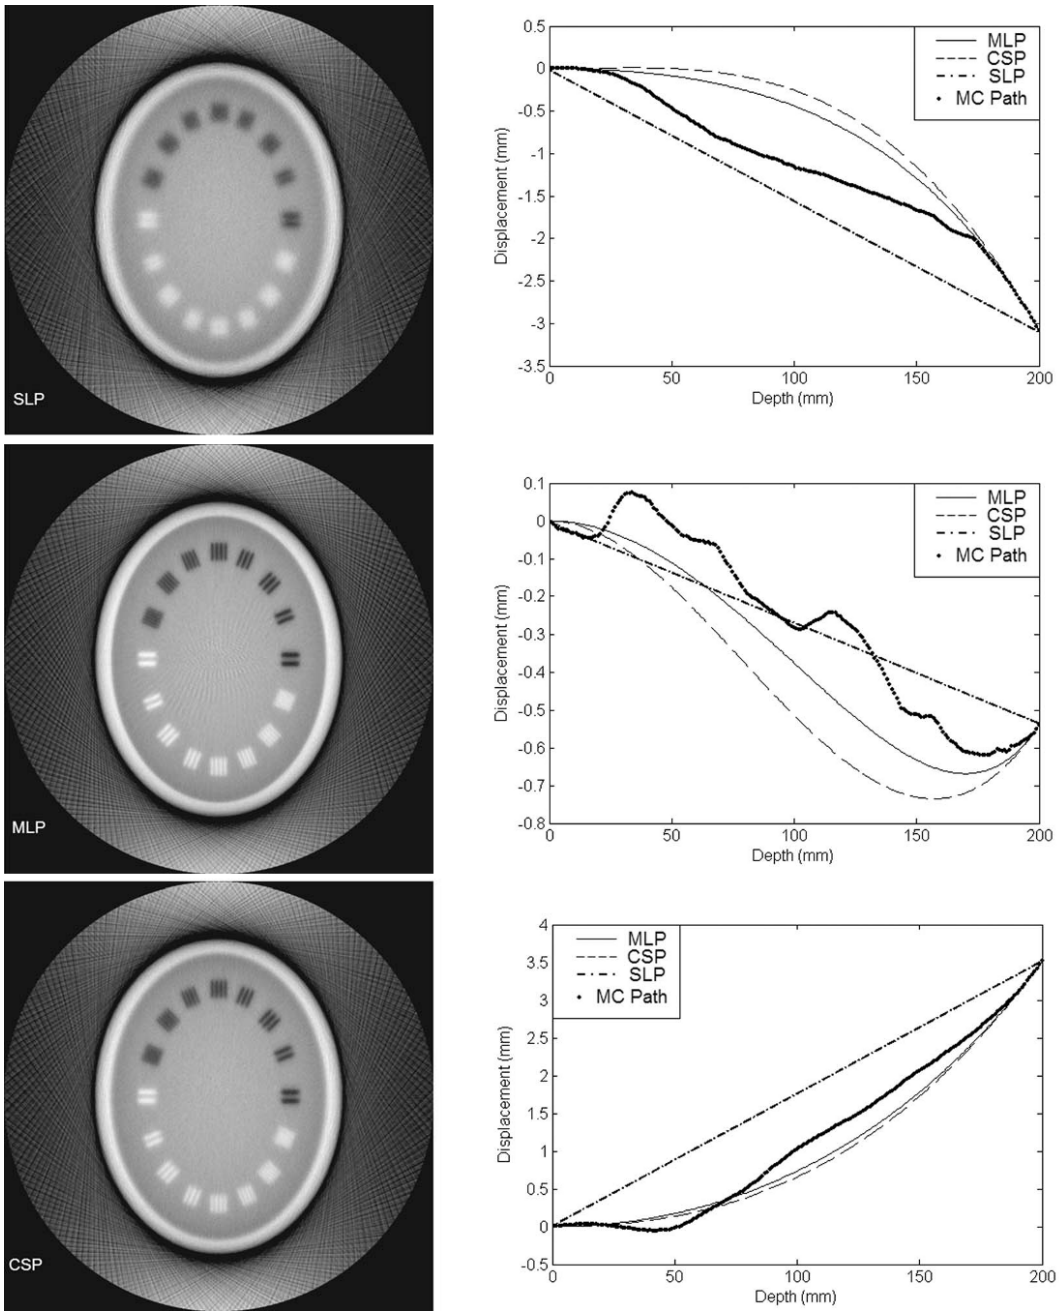

Figure 5.7: Left column: Fully reconstructed phantoms using three algorithms. In the center left MLP is used, and this algorithm is often seen as the most successful for academic purposes. Right column: Different fitting algorithms are compared to the simulated track of a proton (thick line). Algorithm abbreviations: SLP (Straight Line Path), MLP (Most Likely Path), and CSP (Cubuc Spline Path). Images courtesy of Li et al. [39].

Nooij [40] may serve as a starting point for an actual implementation of a reconstruction algorithm. Of particular interest will be the alignment between multiple GridPix planes, as that is proving rather complicated elsewhere in the R&D group.

- Fitting objects.

  Literature shows that most groups dealing with proton imaging use a Most Likely Path method [41, 39, 23, 31, 42, 32, 43, 29], which involves known entry and exit points, and knowledge of the interaction of the particle with tissue (figure 5.7). In X-ray CT reconstructions are often implemented using FBP (Filtered Back Projection) or IR (Iterative Reconstruction) methods. FBP is based on Radon transformations, which is a method of deconvolving the projection that the calorimeter measures into an series of coefficients (integral), which correspond to the different tissues that the beam passed through. For these coefficients a model of the body or organ is used to provide an ansatz for the transformation, and so a possible solution is derived. In IR this process is repeated in order to decrease noise and increase medical information as much as possible. The trade off is between speed and precision. For pCT most groups speak of an ART (Algebraic Reconstruction Technique), which is similar to IR approaches. This stage of data processing was never reached so no implementation has been built as of yet, but the references may be used to see where the competition is.

5.3 Calorimeter

In the November test run we had placed two objects in the setup, with the intention of attempting to reconstruct them from the data. These objects were two triangular wedges of different size. The way we planned to reconstruct the wedges was by correlating the pixel data perpendicular to the beam and the calorimeter output. In calibrating the calorimeter it became apparent that the output was sensitive to changes in beam rate. The beam rate had changed a few times, and after repeatedly trying to create a calibration curve that agrees with common knowledge on PMTs, it was clear that the datasets taken at different beam rates are not compatible: one cannot (easily) compare between the two. The known temperature dependence was considered [44], but since the environmental temperature was monitored and found to be stable, it was rejected as a significant factor in performance.

Because of the above reasons, and because the calorimeter is an old device that was not used and maintained for over 10 years, the use of another calorimeter was considered. The point of the calorimeter is to measure the residual proton energy, and BaF_2 does not compare well to other crystals on its energy resolution. At Nikhef an F101 crystal was found available. This crystal was created for the HERMES experiment, for which Nikhef supplied the PMTs attached to these crystals. In order to see if the new crystal could be suited for our setup, a study was done on the penetration depth of protons using simulation software Geant4. Protons of 190MeV were found to reach a depth in F101 and BaF_2 to within a few percent of each other. The F101 crystal was large enough to capture such protons, but about its use in proton capture no information was found. At KVI a spare PMT was found, the XP2242B, and I looked into coupling it to a $LaBr_3$ crystal, as at KVI the combination had been used with satisfaction. Figure 5.8 shows the

resulting light yield of such a combination. Connecting the PMT to the crystal turned out to be problematic, because it was hygrofobic. Moreover, again not much on proton capture performance was known, making it unclear whether this combination would be an improvement over the BaF_2.

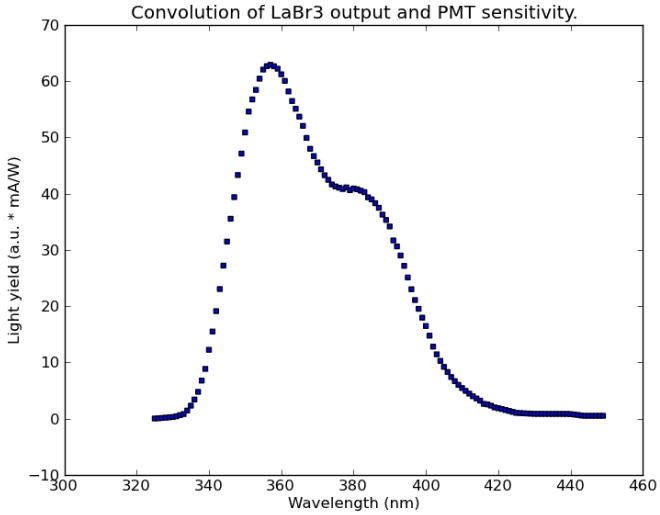

Figure 5.8: The convolution of the XP2242B PMT and $LaBr_3$ response. Sources: Photonis [45], Shah et al. [46].

Since the sensitivity to particle rate is a main concern with the BaF_2 calorimeter, we decided that we needed to be sure about such details before any calorimeter purchase. For $LaBr_3$ we could not find such data, putting that crystal on the list of uncertain performers. We could find data on Helium capture performance with Cesium doped YAP/YAG (Yttrium Aluminium Garnet/Perovskite). At this point, the decision was made to move the focus towards the calorimeter acquisition system, as that could be independently researched from the calorimeter. Also, the fast output of the BaF_2 crystals was a trait not found in any other crystal, and is advantageous because it reduces the number of required secondary scintillators in the beam for a coincidence trigger, reducing the problem of scattering. For the record, *table 5.1* contains all data found on crystals in calorimetry.

For a future search, the following points must be taken into consideration:

- Decay constant: how fast is the signal? For higher rates, the decay constant must be shorter. Also, a secondary fast peak like BaF_2 displays is very useful as input for a trigger system.
- Energy resolution. The point of the calorimeter is energy measurement, so the error should be as small as possible. Ideally, the resolution is known for protons in the 50-200MeV range.
- Performance dependence on factors such as rate, temperature, etc.
- Emission wavelength is ideally in a range easily captured by PMTs.
- Radiation hardness.

Table 5.1: Comparison of materials for calorimetry.

Material	Density [g/cm3]	Emission maximum [nm]	Decay constant	Conversion efficiency	Energy resolution	Remarks
NaI (Tl)	3.67	415	0.23 ms	100		Radiation sensitive, Hygroscopic
CsI (Tl)	4.51	550	0.6/3.4 ms	45	0.6% @He 45MeV	Radiation sensitive, Hygroscopic
CsI (Na)	4.51	420	0.63 ms	85		Hygroscopic
CsI	4.51	315	16 ns	4 - 6		
CaF_2 (Eu)	3.18	435	0.84 ms	50		
^6LiI (Eu)	4.08	470	1.4 ms	35		Hygroscopic
^6Li glass	2.6	390 - 430	60 ns	4 - 6		
CsF	4.64	390	3 - 5 ns	5 - 7		Hygroscopic
BaF_2	4.88	315, 220	0.63 ms, 0.8 ns	16, 5	~9%	
YAP (Ce)	5.55	350	27 ns	35 - 40	2.7%	
GSO (Ce)	6.71	440	30 - 60 ns	20 - 25	2.3% @He 45MeV	Resolution for 1H similar
BGO	7.13	480	0.3 ms	15 - 20	2.9% @He 45MeV	Rad. Hard., Res. for 1H similar
$CdWO_4$	7.90	470 / 540	20 / 5 ms	25 - 30		Radiation hard
Plastics	1.03	375 - 600	1 - 3 ms	25 - 30		
$F101$	3.86			see [47]	6.6% @1GeV	Radiation hard
YAG (Ce)		50	70 ns		3.00% @He	Res. for 1H similar
$LaBr_3$ (Ce)	5.06	360 - 380	25 ns		2.7-3%	Temp. insensitive, Noisy, Hygroscopic

Sources: [48, 47, 32, 46, 44, 49, 50].

5.4 Calorimeter Data Acquisition

Using an oscilloscope for data acquisition, even an advanced one, turned out to be sub-optimal. The main problem with our device was the limited internal buffer that required the operator to reset the measurement every half minute. A system able to continuously readout the calorimeter is desirable. The other requirements are that it be able to cope with acquisition rates of at least 100Hz, and that it has a history buffer allowing for look back at the moment of trigger. Of course the device has to support an external trigger. A feature that would be nice to have is a timestamp with better than microsecond precision.

The device suitable as CaloDAQ (Calorimeter Data AcQuisition) was found in the HiSparc project [51]. HiSparc is a network of cosmic ray detectors in the Netherlands, placed at high schools throughout the country. The number of HiSPARC detection sites in the Netherlands is approximately 100, and they are used to reconstruct the particle showers caused by cosmic rays. The hardware is designed to be cheap and consists of a set of scintillators coupled to a computer. One single device digitizes the scintillator outputs, tests for a coincidence and sends the sampled signal to a computer, on which the results may be viewed and uploaded to a central server. The project is coordinated from the HiSparc group at Nikhef, and both hardware and software are developed at Nikhef. The device used to hook up the scintillators to the computer is called the HiSPARC-box, and this is the device we chose to use as well.

5.4.1 The HiSparc ADC

The HiSPARC-box version 3 has the following characteristics:

- Two ADCs continuously sampling at 200MHz are run in antiphase to give an effective sampling rate of 400MHz. Each sample corresponds to a time period of 2.5ns and is sampled at 12 bits. Given the 0-2.5V input range of the ADCs, this results in a precision of $1.25 * 10^{-12}$ V/bin, which falls well below the noise threshold observed in the data in figure 5.9: $\sigma = 0.0135V$. To get an estimate on whether the sampling rate will suffice, the peak height is studied as function of the sample rate. Figure 5.10) shows that changing to a sampling rate of 400MHz would introduce an error smaller than the error due to other factors.
- The two ADCs need to be aligned because uncalibrated they will have a different response due to being different chips.
- A dynamic range of about 2.5V. Amplitudes from the BaF_2 calorimeter were between 0 and -1.5V, so this is a perfect fit.
- The interface between the computer and the HiSparc-box is a USB2 connection and occurs via a custom message protocol. Software on the computer can send messages to the USB driver, which then transmits them to the HiSparc-box: a write action. Conversely, a read action can be performed by the computer that causes the USB driver to request the contents of the device buffer. The structure of the custom message protocol is defined in a manual.
- A buffer sized at 40k samples, translating to 10µs of history. The length of an event

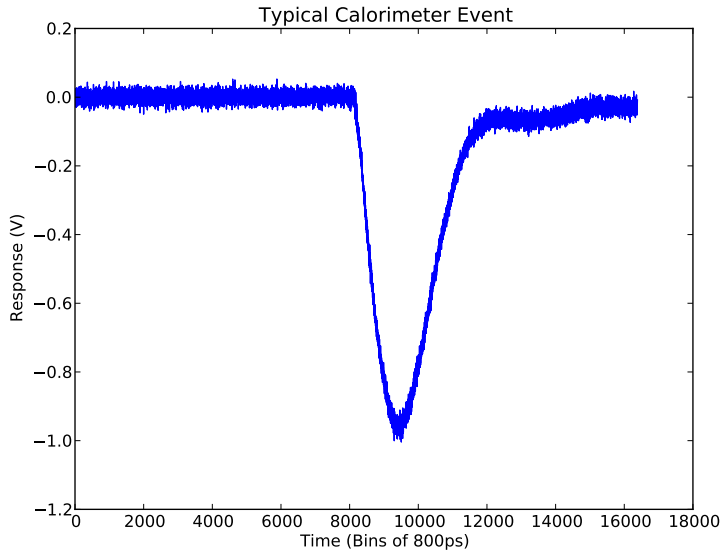

Figure 5.9: A typical curve recorded with the digital oscilloscope in the test run at KVI in November 2011. This oscilloscope samples at 1.25GHz, resulting in a time period per sample of 800ps. The event window is 13μs. This data was resampled to see the effect of lower sample rates on the error in peak height, see figure 5.10. The noise on the baseline before the peak is seen as the convolution of the ADC quantization error and the noise of the calorimeter itself. Assuming a gaussian distribution, this noise is calculated at $\sigma = 0.0135V$.

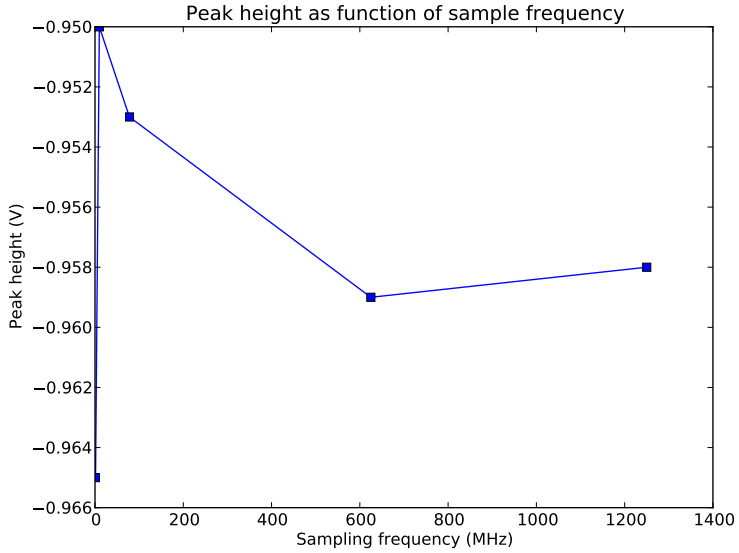

Figure 5.10: Here the peak height is plotted as a function of sampling frequency. The curve in figure 5.9 is resampled to obtain the various frequencies plotted here. The peak height is taken from a second degree polynomial fit around the top of the peak. It can be seen that the variation even at 1.2MHz is 0.7% of the peak height at the original 1.25GHz sampling rate, and less than the convolution of the ADC quantization error and the noise of the calorimeter itself, which is $\sigma = 0.0135V$. For a sampling rate of 400MHz the difference is of the order of 0.1%.

can be set to anything that fits in this window. The device will store event data in the buffer, according to definable parameters such as trigger condition and the number of samples that constitutes an event. This event window is set in terms of a number of samples pre-trigger and post-trigger. Upon a request for the buffer contents, the contents is moved out of the buffer and into the computer. If no readout requests are made, the buffer will fill up and the device will halt storing events until memory is freed.

- Although the device has two channels available for analog to digital conversion, only one is currently used. A future possibility is to use analog electronics to integrate the peak signal and use the second channel to capture this integrated signal. Potentially this could result in a higher acquisition rate, as presumably less samples are required to sample the integral compared to the full curve.
- Because of the inclusion of the buffer, reads and writes may occur simultaneously. This means that a readout does not imply dead time. A read action remove that data from the buffer, freeing space for new acquisitions. If data is written (acquired) faster than it is read, the device will block new acquisitions until read actions free some memory, resulting is an effective dead time. Reasons for this to occur may be high trigger rates or a large event window.
- In the HiSparc experiment the trigger is a coincidence between the two channels, but an external trigger is available. Theoretically the only limit to the trigger rate is the sampling frequency, so there is no reason to assume the desired 100-120Hz is a problem.
- Because of the distributed nature of the HiSparc experiment and the need to couple shower data between stations, each HiSparc-box has a GPS unit which provides a calibration point every second to the internal HiSparc clock. An extra nanosecond clock inside the unit is reset every second when the GPS signal is received, resulting in the internal HiSparc clock to achieve nanosecond precision for the timestamps it attaches to every event message. This provides the extra check for synchronicity.

5.4.2 Development of custom software

We have seen that the hardware specifications of the HiSparc-box are equal or better than what we are looking for, but the software requires discussion as well. In the HiSparc project the box is accompanied by an extensive software package. A lot of analysis is performed on site, at each station, by this software. This allows for both a local view for high school students and saves bandwidth when uploading the data to the central server. While this software is comprehensive, it is fairly slow and relies on commercial software for which the Detector R&D group has no license. The event rate of a typical HiSparc station is 1-4Hz, and the software can keep up with an acquisition rate of approximately 10Hz. This is mostly due to the single-threaded nature of the software and the amount of computation required for analysis. Also, the software is designed with an always online internet connection in mind, to upload the data. Only some statistics are saved locally. These were reasons enough to not use this software, but to write a custom program.

A major part in the development of our own custom software was the driver: in the standard HiSparc program it had been implemented inside the runtime of the commercial

LabView package, which connected directly to the USB device. I started the software task with an investigation into rewriting this driver without using the commercial runtime. Being coded in C using LabView libraries, initially such a port was deemed feasible. Because it was designed to work on the Windows operating system, the first attempt was to get the code to compile within Microsoft Visual Studio. After some attempts this was abandoned in favor of a Python binding for the official manufacturers driver for the chip inside the HiSparc-box. Although the message protocol would have to be reimplemented, this option was deemed more suitable to exploratory programming and less complicated to achieve. A HiSparc derived device called Muonlab, built for muon detection by high school students, already had a software implementation built with the Python binding. Apart from the (undocumented) requirement to load the firmware first, the HiSparc-box could be interfaced with the computer in a similar fashion. Because the manufacturer provides drivers for both Windows and Linux, this provided the option to use either. The rule of thumb is that RelaxDAQ attains higher rates under Linux due to a more suitable scheduler, making the potential move to Linux advantageous.

Encouraged by this success, the HiSparc group at Nikhef revived their own old plan to have a Python-environment for testing. They proceeded to write their own implementation of the HiSparc communication protocol, but now based on a combination of libusb [52] and PyFTDI [53], a driver for the chip written in Python itself, interfacing directly with the USB device via libusb. This approach, titled PySparc [54], turned out to be equally successful, and a comparative study on the performance of both implementations yielded similar acquisition rates at similar event window settings (10μs versus 13μs), indicating that the hardware was the limiting factor and not the Python-environment, known for ease of use but not numerical performance.

So, two basic implementations were produced, comparable in performance in first instance. In addition, the HiSparc group wrote a comprehensive message parser, a ADC calibration routine and a firmware-loader. These last two have Windows-only equivalent tools, built by the HiSparc hardware group.

5.4.3 Online message parsing and speedup

The point of the new CaloDAQ is synchronized acquisitions at a high rate, so both implementations were tested in order to determine the best option. Time was spent on the separation of readout and data-processing into separate threads or processes. Transferring data out of the HiSparc-box buffer into the computer is a matter of calling a read-function as often as possible; a classic example of a pull architecture. Its result is a stream of bytes which need to be decoded into messages according to a message protocol. Using the HiSparc message parser to process the acquisition data in the same thread as receiving the acquisition data consistently led to slowdowns: the HiSparc-box buffer was filling up faster than the software could read, causing the HiSparc-box to start missing events. The continuous readout loop should therefore be as light as possible and do not much more than transfer the data to computer memory where another process or thread can take as long as it needs to decompose the bits into meaningful data. Since Python does not support true multithreading (spreading threads over multiple CPU-cores) but only emulates it, a multiprocessing variant was investigated for both implementations. Multiprocessing in Python works by simply starting multiple Python

interpreter processes tasked with whatever function is delegated to each process. The API (Application Programming Interface) is fortunately very similar between threading and multiprocessing. It was quickly unconvered that my own implementation could not make use of multiprocessing due to the Python binding using Ctypes to interface with the manufacturers driver. This driver is simply a compiled library (.dll or .so), which can't be shared by the multiple processes created by Python. Python however seems to enforce access to all resources for all processes, even if they don't use them, but it can't fork the single USB resource. The PySparc implementation also lost control over the readout process with multiprocessing. Within the HiSparc group a suitable explanation could not be found, but it was firmly established that it didn't work.

A secondary benefit of multiple threads or processes is the ability to spin out the readout loop from the main program, allowing it to accept user interrupts upon which the readout could be stopped. This is why the (emulated) multithreading API was used in the final implementation. The other major design decision was to simply write all data directly and continuously to disk, allowing for offline processing by the PySparc message parser. A few extended data taking runs with both implementations showed that the implementation with the manufacturers driver resulted in extraneous bits and missed hits (compared to a separate counter). A modified PySparc implementation was therefore chosen as our default. Three side-effects are that PySparc includes a routine to align the ADCs and one to upload the firmware, voiding the use of two Windows-only utilities. The third side-effect is that the HiSparc group experienced problems running the code on Windows, although the creator of the Python bindings reports that there should be no objections. It successfully runs on MacOSX and Linux, which is good enough, considering the better supposed performance of RelaxDAQ in Linux. In summary:

- A slight modification of the PySparc program is the default implementation.
- It uses the Python threading API to allow a measurement to run a predefined length of time, and be cleanly interrupted.
- It allows the user to preload ADC alignment data instead of redoing the alignment every measurement.
- Data is directly streamed to disk, although realtime message parsing is left in for debug purposes.
- A small program relying on the HiSparc message parsers is used to process data offline. Currently an event counter, timestamped output of event-graphs, a histogram of all events, and the display of base/peak data are implemented and configurable for output.

5.4.4 Verifying the HiSparc ADC

The HiSparc-box provides no busy-out signal, making a verification of an acquisition on the trigger level impossible. The HiSparc-box does however provide concurrent acquisition and transmission without dead time, if it can keep up with the event rate. Coupled with the desire of a higher acquisition rate, a study on effective acquisition rate was performed. A frequency generator was hooked up to both the external trigger input and channel one, and was varied at frequencies between 50Hz and 10kHz. At the default event window

setting, the acquisition rate seemed to be topping out at 150Hz. This is promising, as it is above the desired minimum rate of 120Hz. There is a certain bandwidth budget between the HiSparc-box and the computer, so it seemed reasonable to assume the acquisition rate could go up if the event window would go down. A range of event windows was therefore set to see the effect on the acquisition rate. The maximum event window is 10μs (this corresponds to 12kB of data, multiplied with 2 channels, plus some overhead yields something close to the HiSparc-box hardware buffer of 36kB). With the 0.63ms decay time of the BaF_2 and 400MHz sampling rate in mind, an event window of 1000 samples, corresponding to 2.5μs, would be plenty (figure 5.9). *Table 5.2* shows the results. Note that the Event Window is set with bins of 5ns, but retrieved in bins of 2.5ns. Both values are given respectively. The HiSparc hardware design team stated that the theoretical limit of the data rate should be 1250kB/s, which agrees with the observed results to within a factor 2. With the event window at 2.5μs a rate just below 300Hz was obtained. This means the design target of 120Hz has been reached and surpassed, making the new calorimeter data acquisition a success.

Lastly the response of the calorimeter-HiSparc combination was tested using cosmic rays. Since nothing is known about the internals of the calorimeter, it has not been calibrated for cosmic rays. The uncalibrated data can still give an impression on the performance of the new CaloDAQ. Figure 5.11 shows examples of the four typical event curves encountered. First, a 'proper' curve is shown, fitting nicely within the dynamic range of the ADCs. The second curve is clipped: the energy deposited by the passing cosmic ray surpassed the dynamic range of the ADCs. From previous cosmic tests, we know that the BaF_2 can output signals up to -20V due to cosmics. The third curve shows clearly the two ADCs running in tandem. The height of the peak is comparatively shallow, which allows the difference between the ADCs to become visible. The automated ADC alignment procedure can be improved, or the ADC configuration could be tuned manually. The last curve shows a strange defect visible in about 1-2% of the events recorded from cosmics. The kink usually appears at the same location in the curve: at about half height on the fall of the signal. A potential cause has not yet been identified. A possibility is a second low energy hit, or a noise effect of the PMT.

Table 5.2: Comparison of data rates with different Event Window sizes.

Overhead (bytes)	**Event Window** (samples)	**Event Window** (ns)	**Event Window** (bytes)	**Frequency** (Hz)	**Data rate** (kB/s)
25	3 / 6	15	18	31644	1361
25	11 / 22	55	66	11628	1058
25	21 / 42	105	126	6499	981
25	51 / 102	255	306	2793	925
25	101 / 202	505	606	1433	904
25	201 / 402	1005	1206	726	894
25	501 / 1002	2505	3006	293	888
25	801 / 1602	4005	4806	183	884
25	1201 / 2402	6005	7206	122	882
25	1601 / 3202	8005	9606	92	886
25	2000 / 4000	10000	12000	73	878

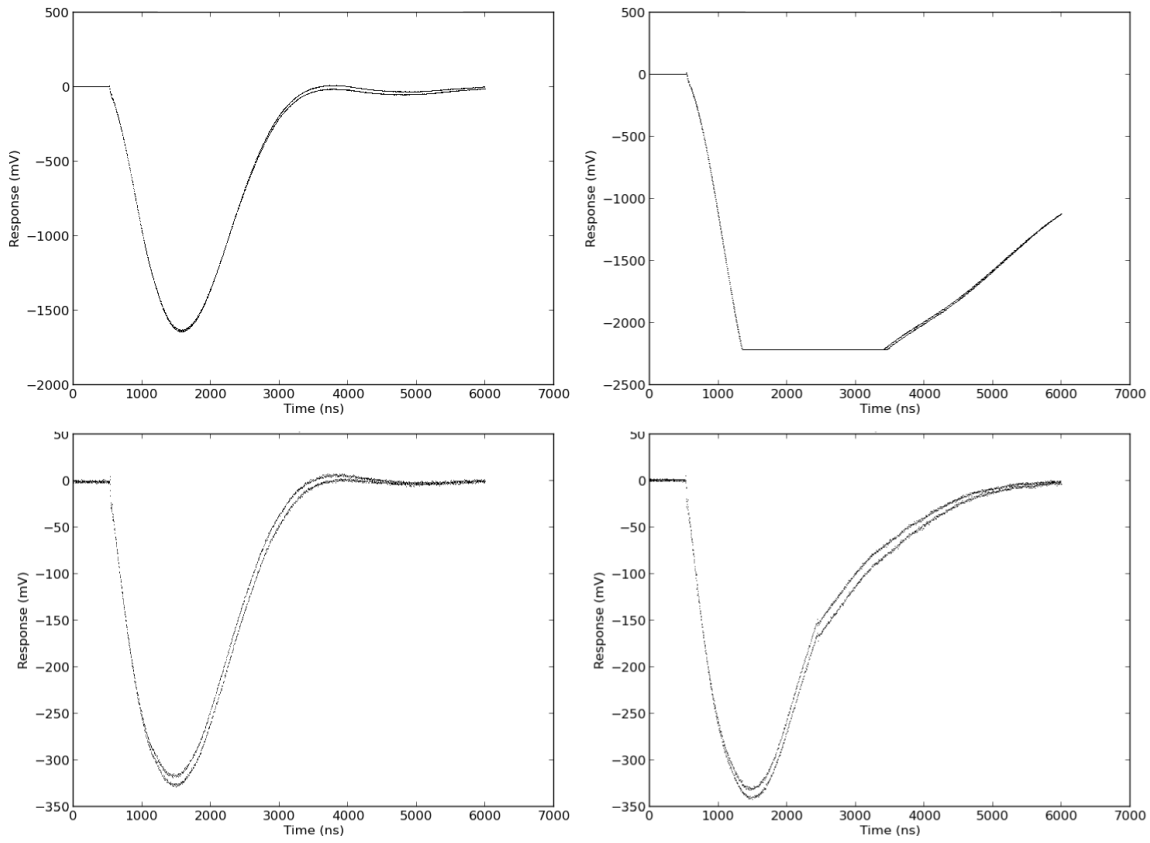

Figure 5.11: Four typical event curves corresponding to cosmics. Clockwise, starting top left: a regular curve, a clipped curve, a shallow curve where the two ADCs become distinguishable and a strange kink. Explanations may be found in the text.

5.4.5 Summary

The results on the new CaloDAQ system can be summed up as follow:

- The event acquisition rate exceeds the desired 120Hz by a wide margin. Using analog signal preprocessing so that only a few samples are necessary for an accurate representation of the energy captured by the calorimeter, rates up to 30kHz are a possibility.
- The software performs very robustly at higher rates and extended data taking runs, although a serious test of the combination of the two still needs to be performed.
- An offline message parser outputs various kinds of graphs, easily configurable.
- The dynamic range and quality of the acquired data seems in order. The ADC alignment could be better tuned however, and the exact source for the kink must be discovered.
- Changes in the HiSparc firmware may increase the acquisition rate even further. Switching off the second channel and lowering the sampling rate seem easily implemented.

5.5 Trigger

Now that the calorimeter data can be acquired at sufficiently high rates, it is time to optimize the trigger so that it can actually can deliver a reliable rate in the vicinity of 120Hz. Previously, when the busy signal of the RelaxD was not used, a timer set to ~500ms was used to both give the RelaxD time to fully complete the readout of an event and to compensate for trigger delays. Three important things were changed:

- The trigger delays would be shortened as much as possible by using short cables.
- The default state of the system would be acquisition mode on, instead of off.
- The busy out of the RelaxD board would be incorporated in the trigger as the signal that dictates the start and end of a readout, and therefore no acquisitions.

The latter two are discussed below.

5.5.1 Changing the default state

Within the Detector R&D group, work was being done on a trigger incorporating a GridPix that was built with the only (theoretically) limiting factor: the GridPix readout process at 120Hz. As has been seen in the previous section, the calorimeter DAQ was designed with a continuously sampling ADC. The device keeps a 10μs buffer, which provides plenty of history to select the appropriate window for our uses. Moreover, it allows data taking without dead time. The GridPix has a certain dead-time (7-15 ms) when it receives the end of a shutter signal. After the end of a shutter, all pixel data is unloaded from the TimePix chip and transmitted to the computer by the RelaxD board. During this time, the chip can not acquire new event data, so a trigger should take

this into account. Secondly, there is time lost in executing the trigger logic between the passing particle and the commencement of the GridPix readout. As stated in Tsopelas [38], this translates to a small offset in track height. This is easy to correct for, but the data not recorded during the time it takes to execute the trigger logic is lost. The default state of the GridPix was therefore changed from off to on. By default, the GridPix is collecting data, and only when a trigger is received is data collection ceased and readout started. This trigger is then delayed to allow the electrons and ions to drift and induce a charge in the TimePix chip. The drift time is fortunately much larger than the time it takes to execute the trigger logic, so no data is lost and no offset should be introduced.

5.5.2 Utilizing the Busy signal

The second improvement is the incorporation of the RelaxD busy signal. This signal precisely determines the dead time during readout and will be used to block new hits during this period. Previously the busy signal was not used, but instead a relatively large delay with a timer. This delay was about half a second, which agrees nicely with the observed 2Hz acquisition rate in November. Occasionally the RelaxD board fails to correctly switch to acquisition mode, and will not recover by itself. Fortunately, when this happens, the busy signal is never activated and the RelaxD is reset by switching the shutter back on. A fallback timer, analogous to the delay timer as used in November, creates the automatic reset for that scenario, and could be set to a similar value as before: 500ms. The CaloDAQ trigger is for this reason connected to the busy out of the RelaxD: it never transmits an event to the computer if the RelaxD doesn't, so that we don't have to filter for incomplete events with only calorimeter data.

The new trigger design makes full use of the busy signal to decrease acquisition dead time as much as possible for a large majority of all incoming triggers. The dead time in a successful readout action is 7-15ms, so one can see how this is a major improvement over 500ms. For comparison: if we want to be able to maintain a 120Hz acquisition rate, one readout event must take no longer than 8.3ms.

5.5.3 Trigger logic

Commencing a data taking run starts with switching the system into acquisition mode. Currently, this is done by waiting for the first event, which can be generated artificially. The end of this event results in the shutter to be opened by default. The triple coincidence unit has a window of 40ns. If this condition is met, the drift timer (t_{drift}) is started, to allow the electrons generated by the proton to drift towards to chip and induce charge in the TimePix chip. When the drift timer elapses, a second timer ($t_{fallback}$) is started. Upon starting this second timer, it closes the shutter for the RelaxD board, which should cause it to start readout immediately. The busy output of the RelaxD will switch on: it should always be the opposite of the shutter. The start of the busy is used to trigger the CaloDAQ, and the end of the busy is used to reset the fallback timer. The fallback timer reopens the shutter when either its timer elapses or the end of the RelaxD busy resets it. It then releases the veto on the coincidence unit so that data acquisition is resumed. While two timers are running, the system is either waiting for the electrons to drift or

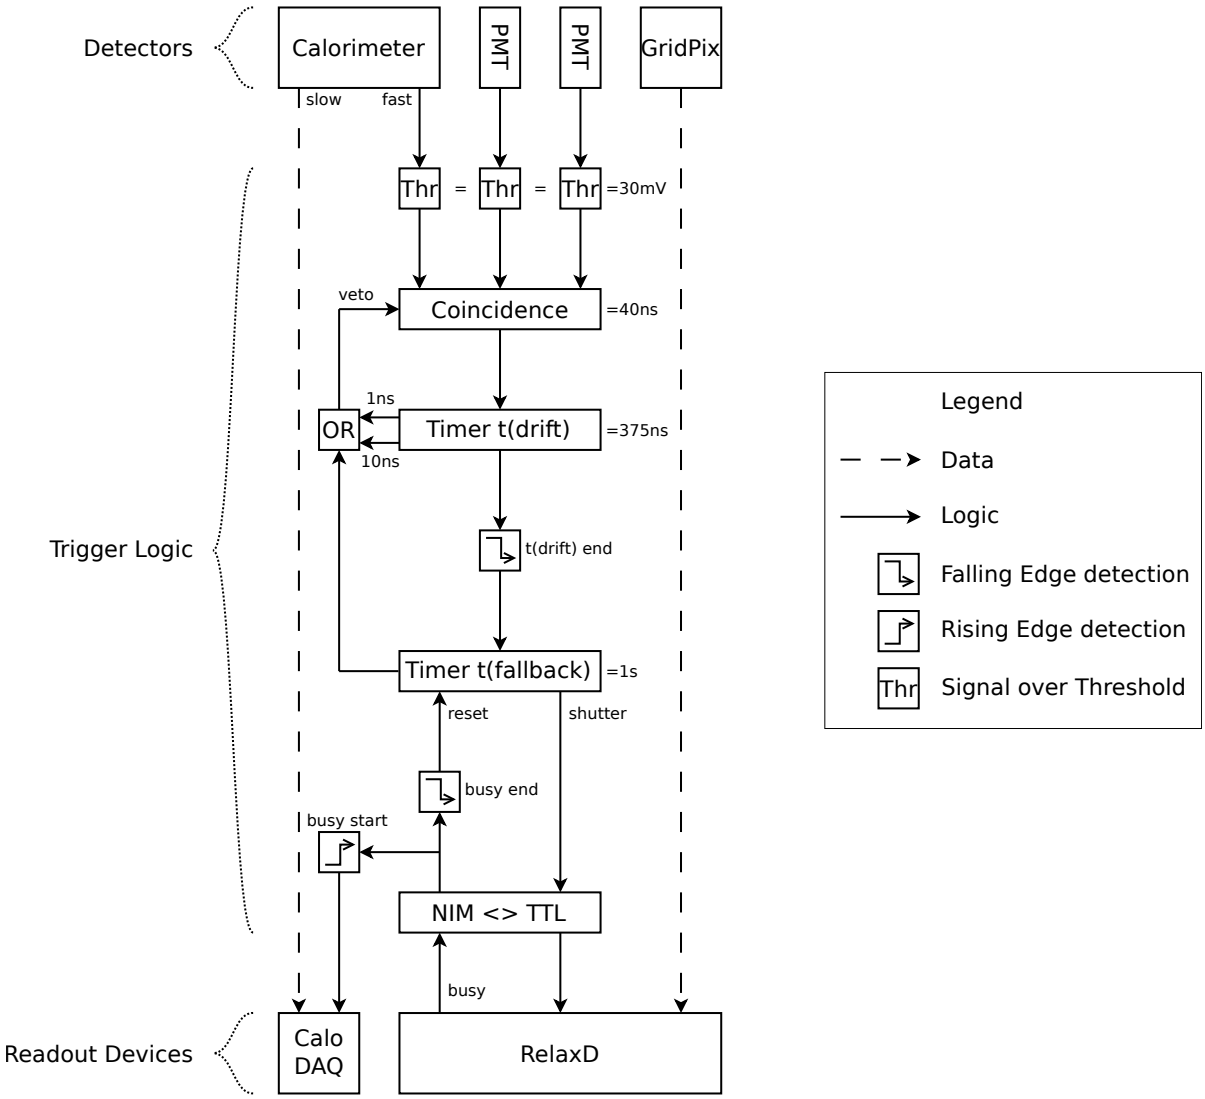

Figure 5.12: Overview of the trigger logic and its implementation. The Falling Edge detection between the timers is implemented by using the End-port of the drift timer. The other two Edge detections are implemented by a single coincidence.

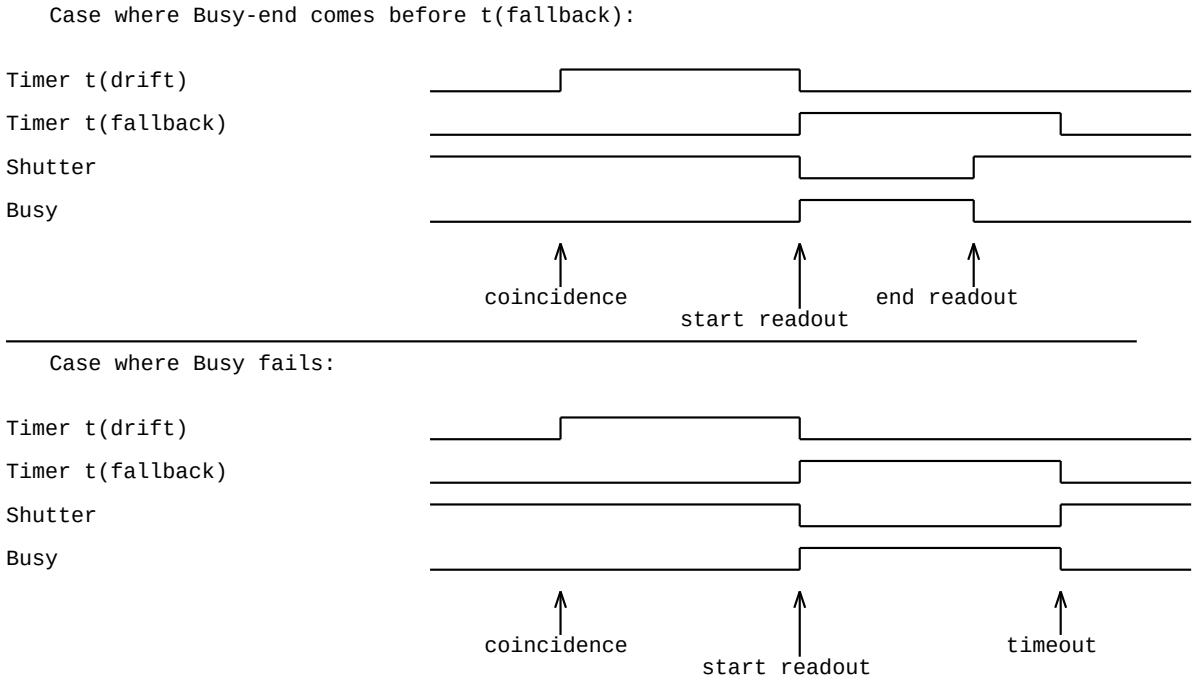

Figure 5.13: Trigger timing. The drift timer allows the ionized gas and electrons to drift, and is set to the top-to-bottom theoretical drift time based on Kopperts simulations. The readout stops when either the RelaxD board stops sending out a busy signal, or when the fallback timer has elapsed. The fallback timer was inserted to reset the system after occasional problems with the RelaxD that cause it to 'hang'. The fallback timer prevents us from down time and an otherwise incomplete acquisition, by forcing the system out of this state if it lasts too long.

for the data acquisition to finish. These two timers therefore veto the coincidence unit and thus the acceptance of new events.

For the final design we decided a triple coincidence between two small scintillators and the BaF_2 fast out would be the best way to exclude noise hits. Another detail is that the drift timer is connected to the OR-port twice: with a short (1ns) cable and with a longer 10ns cable. This is done to bridge the ~6-8ns it takes for the signal to arrive in the fallback timer and to start sending out its signal to the OR-port. Since the coincidence rates can be quite high, this small window was enough to accept new events which upset the proper acquisition of events.

In the figure 5.12 the trigger schematic is shown. The timing diagram is visible in figure 5.13.

5.5.4 Trigger verification

Initially the trigger was tested on a single frequency generator. This is only a partial test: first it will not generate jitter and second it does not test the triple coincidence. It was selected nonetheless, because a frequency generator that also could produce jitter was not available, and no two or three frequency generators were available or properly configurable for a reliable coincidence. The single frequency generator did however show that the trigger logic could be stably operated with rates up to 100kHz, which puts us far above the previous rate of 2Hz and far above the maximum acquisition rate of the GridPix. The triple coincidence was tested on a lower rate using cosmic rays: both scintillators and the calorimeter provided the proper coincidence. The acquisition rate in the calorimeter DAQ and the hit counter in the electronics rack kept proper synchronization during multiple multi-day test runs, and generated no noise hits. This test run was also a test of the calorimeter data acquisition quality: curves with the correct shape and magnitude that can be expected from cosmics were produced. The last test the trigger was subjected to was the single coincidence of the noise of the calorimeter. This produced event rates of about a 100Hz, and would test everything except the coincidence unit. Unfortunately an as of yet unresolved bug in the RelaxDAQ software did not allow for an extended test run at high rate: the software would crash on our machine after having recorded about 9000 events. By prematurely stopping the acceptance of new events, synchronicity could be tested for high rates during 2-3 minutes. During these very short high rate runs synchronicity between RelaxDAQ and CaloDAQ was kept; a promising result. Once the software bug in RelaxDAQ is resolved, an extended high rate run should be the last test before a new run at the proton accelerator at KVI.

5.5.5 Summary

The efforts on improving the trigger system have resulted in the following:

- The full setup has been tested with cosmics for continuous multiday runs with 100% synchronicity and success.
- The trigger with calorimeter data acquisition has been successfully tested up to an acquisition rate of 31kHz.

- The full setup has been tested with 2-3 minute runs at an acquisition rate of ~70Hz with 100% synchronicity. 120Hz was not attained, and after discussion with the RelaxDAQ authors it was deemed a problem with the Linux build that should be fixable. The Windows build manages 120-130Hz for the RelaxDAQ authors.
- The Linux build of RelaxDAQ must be fixed so that it does not crash after ~9000 events. The problem may partly lie with the RelaxD board as the firmware is in flux at the moment. Once fixed, an extended high speed run (an hour at 100Hz+) must be conducted to test whether or not the trigger keeps the two data channels in sync.
- The design does not correct or block events with multiple particles. This means that bunch occupancy in the beam must remain as close to 1 as possible and the particle rate should still be kept within an order of magnitude or so of the desired acquisition rate (120Hz). The trigger provides only a logic block on new acquisitions, but not on new particles.

5.6 Conclusions

At the end of *Chapter 4* some considerations were laid out for the eventual use of the setup in a medical context. Now we revisit these questions with our current abilities in mind.

- How many protons do we need for an image?

  Amaldi et al. [34] write that 10^6 usable proton tracks should be enough for a 100 by 100 mm^2 image. Scaled to the dimensions of our current GridPix drift volume, we would need 227k tracks.

- What is the event acquisition rate that is required?

  Amaldi et al. [34] also write that the desired time for a full scan is in the order of a few seconds. If we take 10 seconds for example, this translates to an event acquisition rate of 27kHz. While the GridPix does not come close, a future version might. The TimePix version 3 is designed with a target pixel event acquisition rate of 40MHz which translates to an event acquisition rate of roughly 120kHz (a typical track had 100-300 pixel hits in our test run in November). Our new HiSparc based calorimeter ADC can manage an event acquisition rate of 31kHz, if a way is found to put an event with sufficient energy resolution in 6 samples.

- How large is the beam spot?

  While not investigated with the November data, for imaging a large spot is advantageous as speed is a more important consideration than spatial precision. Using a tracking plane before and after the patient, all the spatial information that is required should be provided.

- What is the required beam current?

  This depends greatly on the ability to handle multiple particles in one event. Since our calorimeter can measure only one residual energy value, any multi-particle

event must be discarded. Tsopelas [38] shows that a dependence of the number of ionizations per unit track length on the particle energy is visible, so perhaps this information can be used to select the track corresponding to the residual energy in the calorimeter in some cases. For the time being there is no experience with such a selection mechanism yet, which means that the beam current must translate to a proton rate no higher than the maximum event acquisition rate that the system can handle. At the moment this could be as high as 120Hz if the Linux RelaxDAQ problems are solved.

- How many tracks can be correlated (before and after patient)?

  Unfortunately multiple GridPix chips were not available, so correlation between before and after the patient is not possible at this time.

- Scattering in strip detectors versus GridPix: is GridPix advantageous?

  It is widely accepted that gas trackers scatter incident particles much less than solid state trackers, but no quantitative comparison has been performed at this time in the context of medical imaging.

CHAPTER

SIX

FURTHER THOUGHTS

The performance of the Nikhef/KVI Proton Imager was investigated. Using the data and experience from a proof of concept test conducted at the proton accelerator at KVI, points for improvement were identified. In particular, a calorimeter data acquisition system (CaloDAQ) and a new trigger were built. Preliminary tests indicate the improvements reach or exceed the target event acquisition rate of 120Hz. Timestamping of both data stream give a secondary handle on the synchronicity. Continuing problems with the tracker hard- or software have made a real test impossible at this time. Some recommendations and considerations for future development follow.

6.1 Tracking

Although the TimePix family of chips are used in many applications and to great satisfaction, the experiences with the two particular chips in this setup have not been great. We suffered from a broken chip, a new batch that was of questionable quality, and the ongoing development of RelaxD and RelaxDAQ. Hopefully these will all be resolved so that operating of the GridPix chip will be routine, and that the number of chips can be expanded to provide a greater tracking volume and a second tracking plane. Track fitting and 3D reconstruction will require a considerable effort.

For applications in medicine higher event acquisition rates are required, much higher than the 120Hz of the TimePix chip. A future version, TimePix 3, is currently being designed to attain a 40MHz pixel event acquisition rate, translating roughly to a 40kHz+ event if it were to be used in this setup. This gets the performance in the realm of medical feasibility and such an improvement is therefore eventually essential.

Dr. T. Tanimori gave a presentation at KVI on his research on large cubical drift chambers. While not using TimePix chips, his 3D tracker has a design similar to our GridPix: a 2D detection surface to one side of a cube, with precision time registration of each signal. Dubbed μPix, this is then used to recreate hits inside the drift volume. He reported successful performance of the device when applied in Dark Matter searches. While this research causes much lower pixel event rates than medical imaging, the array can register pixels at a rate of 10MHz. The experience with large drift chambers could serve as an example for when our drift volumes are scaled up [24].

A potential replacement for TimePix related chips could be SiPM (Silicon PhotoMultipliers). These chips integrate a small solid state photomultiplier above each pixel, allowing the detection of single photons per pixel with excellent timing. If each pixel is connected to a light conducting fiber, which are then fanned out and configured in a typical silicon strip detector plane, tracking hadrons could be a possible application. These fibers are still advantageous compared to silicon strips, because they can be of low-Z materials so that scattering and energy loss is minimized, as is the case with gas detection in GridPix [25, 26].

6.2 Calorimetry

First the calorimeter itself should be replaced. The current device is old and therefore most likely discolored and no longer perfectly light tight. Once a new device is available, with known dimensions of the crystal, lab tests with energy reconstruction will be easier to conduct and performance will improve and will be better defined.

As for the calorimeter data acquisition system:

- The strange kink seen in the last curve of *figure 5.7* must be investigated and resolved.
- An energy reconstruction must be performed at the earliest opportunity, to test the real-world performance.
- Pre-shaping of the calorimeter output could shorten the required event window, so that the effective event acquisition rate can be increased. It has been established that the device can acquire events up to 30kHz if the event window is small.
- A small firmware change in the HiSparc-box could disable the second channel, doubling the effective buffer space.
- The communication between the USB interface chip and the memory chip is currently capped at 1250kB/s. It seems that increasing this bandwidth and thereby acquisition rate would be fairly straightforward. The USB interface of the chip, the FT2232H, supports a data rate of 10MB/s so is not limiting us yet. Within Nikhef there already is a spin-off of the HiSparc-box, Muonlab, so there is some experience with application specific implementations of the same basic device.

6.3 Trigger

A small improvement that can be implemented with the recently developed Python bindings for RelaxDAQ is a deeper coupling to the CaloDAQ. It should be fairly straightforward to implement a unified starter script that prepares both tracker and calorimeter for data acquisition. A next step could be a post-acquisition script that verifies the number of acquisitions and timestamps for both tracking data and calorimeter data. The HiSparc message parser as of yet is too slow for such a task (as an example: processing the data in a 30 second run at a 31kHz event acquisition rate takes about an hour).

In a future application the accelerator itself may provide the trigger. This would remove the need for scintillators in the beam line, improving energy resolution and decreasing scattering, and provides a known signal for when particles should be seen, allowing for a better performance analysis of the setup. At this point an effort can be made into simulating the setup with software such as Geant.

6.4 Acknowledgments

I would like to thank Nikhef, Jan Visser and Sytze Brandenburg for the opportunity to work on this project: a crossover between disciplines and one of the few master projects in particle physics that are close to a real and lofty application. It was also a project that entailed almost any type of scientific activity:

- literature study: not only a bit on particle physics and the competing alternatives in hadron imaging, but also a tiny bit of medical knowledge,
- physics analysis: how we can combine our knowledge of physics with experimental data to fully understand what is going on,
- software development: apart from processing gathered data, the CaloDAQ driver had to be written which was my first experience with performance critical coding,
- clean room construction of the replacement GridPix,
- the data taking run at the KVI accelerator and all the excitement that goes with a live experiment,
- attendance of a 4-day conference in Darmstadt on the topic of hadron therapy, allowing me to speak with one of the persons cited quite often in this thesis,
- bugging many people in the Detector R&D group and elsewhere in Nikhef to help a novice out,
- setbacks: the broken GridPix or the missing documentation on the HiSparc-box, to test how difficult keeping a schedule can be,
- lab work: testing my abilities to solder, plug a cable or built a frame on the fly.

The breadth of tasks was astonishing, and I am very glad to have had my time as a master student on a project like this. In no particular order, a few names deserve to be mentioned. Vincent van Beveren was an essential help in both combating RelaxDAQ oddities and in getting the HiSparc-drivers to work. Bas van der Heijden was the go-to guy for all RelaxD problems, and didn't mind reviewing some cabling every now and then. Arne de Laat, David Fokkema and Hans Verkooijen, the Hisparc guys, for helping me understand the hardware design of the HiSparc-box and developing most parts of the current driver. The two Martins, van Beuzekom and Fransen, taught me basically all I now know of electronics. I had not had any experience with cable latency or trigger logic before, and without their hands-on work there would be no trigger. Wilco Koppert, for his detailed knowledge on GridPix, gas and drift properties, track reconstruction and clean room construction of GridPixes. Panagiotis Tsopelas for giving me a running start, and Jan Visser for helping me keep running. Francesco, Mary, Enrico, Marten, Rolf,

Matteo for all their support and the R&D group as a whole for showing me what the life of a researcher is like. Lastly, I want to thank my supervisors Jan Visser and Sytze Brandenburg for their thorough support with this thesis and for sharing their experience and knowledge in performing this project.

I hope this setup will soon see a successful test at the KVI accelerator and that the project will someday result in a competitive apparatus which can help cure cancer victims. I want to thank the reader for reading!

CHAPTER

SEVEN

REFERENCES

1) Claus Grupen. Astroparticle physics. 2005.

2) RP Feynman, RB Leighton, M Sands. The Feynman Lectures on Physics; Vol. I. *American Journal of Physics*, 1965.

3) Michael Hale. PHYC 6421.03 Radiological Physics Lecture. *Dalhousie University*, http://myweb.dal.ca/halem/phyc6421/,

4) H J Hoekstra, W F Sindelar, T J Kinsella, J Oldhoff. History, preliminary results, complications, and future prospects of intraoperative radiotherapy. *Journal of surgical oncology*, 36:175–82, 1987.

5) National Institute of Standards and Technology. NIST PSTAR Tables. *http://physics.nist.gov/PhysRefData/Star/Text/PSTAR.html*, Accessed: December 11, 2011.

6) Dieter Schardt, Thilo Elsässer. Heavy-ion tumor therapy: Physical and radiobiological benefits. *Reviews of Modern Physics*, 82:383–425, 2010.

7) A.M. Koehler. Proton radiography. *Science*, 160:303–4, 1968.

8) U. Amaldi, S. Braccini. Present challenges in hadrontherapy techniques. *The European Physical Journal Plus*, 126:1–15, 2011.

9) Jacobus M Schippers, Antony J Lomax. Emerging technologies in proton therapy. *Acta oncologica (Stockholm, Sweden)*, 50:838–50, 2011.

10) Adrian Staab, Hans Peter Rutz, Carmen Ares, Beate Timmermann, Ralf Schneider, Alessandra Bolsi, Francesca Albertini, Antony Lomax, Gudrun Goitein, Eugen Hug. Spot-scanning-based proton therapy for extracranial chordoma. *International journal of radiation oncology, biology, physics*, 81:e489–96, 2011.

11) R.R. Wilson. Radiological Use of Fast Protons. *Radiology*, 47:487, 1946.

12) G Kraft. Tumor therapy with heavy charged particles. *Progress in Particle and Nuclear Physics*, 45:2000.

13) K Gunzert-Marx, H Iwase, D Schardt, R S Simon. Secondary beam fragments produced by 200 MeV/u 12C ions in water and their dose contributions in carbon ion radiotherapy. *New Journal of Physics*, 10:075003, 2008.

14) David J Brenner, Eric J Hall. Secondary neutrons in clinical proton radiotherapy: a charged issue. *Radiotherapy and oncology : journal of the European Society for Therapeutic Radiology and Oncology*, 86:165–70, 2008.

15) Thilo Elsässer, Michael Scholz, Thilo Elsa. Cluster Effects within the Local Effect Model Cluster Effects within the Local Effect Model. 167:319–329, 2007.

16) L Wisser. Pion treatment of prostate carcinoma at Paul Scherrer Institute (formerly Swiss Institute for Nuclear Research (SIN)) from 1983 to 1992. *Cancer radiothérapie : journal de la Société française de radiothérapie oncologique*, 8:88–94, 2004.

17) G Goodman, L Skarsgard, G Thompson, R Harrison, G Lam, C Lugate. Pion therapy at TRIUMF. Treatment results for astrocytoma grades 3 and 4: a pilot study. *Radiotherapy and Oncology*, 17:21–28, 1990.

18) Ugo Amaldi, Gerhard Kraft. Radiotherapy with beams of carbon ions. *Reports on Progress in Physics*, 68:1861–1882, 2005.

19) S. Rossi. The status of CNAO. *The European Physical Journal Plus*, 126:2011.

20) Wikipedia. File:CompT.jpg. *http://commons.wikimedia.org/wiki/File:CompT.jpg*, Accessed: May 29, 2012.

21) K M Hanson, J N Bradbury, T M Cannon, R L Hutson, D B Laubacher, R J Macek, M a Paciotti, C a Taylor. Computed tomography using proton energy loss. *Physics in medicine and biology*, 26:965–83, 1981.

22) Hyungjoon Ryu, Eunsuk Song, Jaeki Lee, Jongwon Kim. Density and spatial resolutions of proton radiography using a range modulation technique. *Physics in medicine and biology*, 53:5461–8, 2008.

23) S. N. Penfold, a. B. Rosenfeld, R. W. Schulte, K. E. Schubert. A more accurate reconstruction system matrix for quantitative proton computed tomography. *Medical Physics*, 36:4511, 2009.

24) Michiaki Takahashi, Shigeto Kabuki, Kaori Hattori, Naoki Higashi, Satoru Iwaki, Hidetoshi Kubo, Shunsuke Kurosawa, Kentaro Miuchi, Kiseki Nakamura, Hironobu Nishimura, Joseph D. Parker, Tatsuya Sawano, Atsushi Takada, Toru Tanimori, Kojiro Taniue, Kazuki Ueno. Development of an Electron-Tracking Compton Camera using CF4 gas at high pressure for improved detection efficiency. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 628:150–153, 2011.

25) Philips Press Release. Philips Digital Photon Counting. *http://www.research.philips.com/initiatives/digitalphotoncounting/news/backgrounders/091008-photon-counting.html*, Accessed: June 1, 2012.

26) Elizabeth Clements. A new kind of detector technology that could lead to discoveries in particle physics may also lead to better 3D images of the human body and help cancer patients. *http://www.symmetrymagazine.org/cms/?pid=1000942*, Accessed: June 1, 2012.

27) S. Braccini, A. Ereditato, I. Kreslo, U. Moser, C. Pistillo, S. Studer, P. Scampoli. Nuclear Emulsion Film Detectors for Proton Radiography: Design and Test of the

First Prototype. 626–630, 2010.

28) L. Johnson, B. Keeney, G. Ross, H.F.-W. Sadrozinski, a. Seiden, D.C. Williams, L. Zhang, V. Bashkirov, R.W. Schulte, K. Shahnazi. Initial studies on proton computed tomography using a silicon strip detector telescope. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 514:215–223, 2003.

29) C. Talamonti, V. Reggioli, M. Bruzzi, M. Bucciolini, C. Civinini, L. Marrazzo, D. Menichelli, S. Pallotta, N. Randazzo, V. Sipala, G.a.P. Cirrone, M. Petterson, N. Blumenkrantz, J. Feldt, J. Heimann, D. Lucia, a. Seiden, D.C. Williams, H.F.-W. Sadrozinski, V. Bashkirov, R. Schulte. Proton radiography for clinical applications. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 612:571–575, 2010.

30) D C Williams. The most likely path of an energetic charged particle through a uniform medium. *Physics in Medicine and Biology*, 49:2899–2911, 2004.

31) GAP Cirrone, Giacomo Cuttone, Giuliana Candiano. Monte Carlo studies of a proton computed tomography system. *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, 54:1487–1491, 2007.

32) G.a.P. Cirrone, G. Candiano, G. Cuttone, S. Lo Nigro, D. Lo Presti, N. Randazzo, V. Sipala, M. Russo, S. Aiello, M. Bruzzi, D. Menichelli, M. Scaringella, S. Miglio, M. Bucciolini, C. Talamonti, S. Pallotta. The Italian project for a proton imaging device. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 576:194–197, 2007.

33) C. Civinini, M. Brianzi, M. Bruzzi, M. Bucciolini, G. Candiano, L. Capineri, G.a.P. Cirrone, G. Cuttone, D. Lo Presti, L. Marrazzo, E. Mazzaglia, D. Menichelli, S. Pieri, N. Randazzo, V. Sipala, C. Stancampiano, C. Talamonti, M. Tesi, S. Valentini. Towards a proton imaging system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 623:588–590, 2010.

34) U. Amaldi, a. Bianchi, Y-H. Chang, a. Go, W. Hajdas, N. Malakhov, J. Samarati, F. Sauli, D. Watts. Construction, test and operation of a proton range radiography system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 629:337–344, 2011.

35) J. Jakubek, M. Martišíková, C. Granja, K. Gwosch, B. Hartmann, S. Pospíšil, O. Jäkel. 249 Investigation of the Timepix Detector for Beam Range Verification in Ion Beam Therapy. *Radiotherapy and Oncology*, 102:S128–S129, 2012.

36) P Henriquet, E Testa, M Chevallier, D Dauvergne, G Dedes, N Freud, J Krimmer, J M Létang, C Ray, M-H Richard, F Sauli. Interaction vertex imaging (IVI) for carbon ion therapy monitoring: a feasibility study. *Physics in medicine and biology*, 57:4655–69, 2012.

37) F Sauli. Principles of operation of multiwire proportional and drift chambers. *Nuclear Instruments and Methods*, 162:92 p, 1977.

38) Panagiotis C Tsopelas. Master Thesis. 2011.

39) Tianfang Li, Zhengrong Liang, Jayalakshmi V. Singanallur, Todd J. Satogata, David C. Williams, Reinhard W. Schulte. Reconstruction for proton computed tomography by tracing proton trajectories: A Monte Carlo study. *Medical Physics*, 33:699, 2006.

40) Lucie De Nooij. Data Analysis of Measurements on a GridPix detector. 2009.

41) R F Hurley, Campus Street, Loma Linda. Water-equivalent path length calibration of a prototype proton CT scanner. 39:2438–2446, 2012.

42) Vladimir Bashkirov, Reinhard Schulte, George Coutrakon. Development of proton computed tomography for applications in proton therapy. *APPLICATION OF ACCELERATORS IN RESEARCH AND INDUSTRY*, 2–5, 2009.

43) Nicolas Depauw, Joao Seco. Sensitivity study of proton radiography and comparison with kV and MV x-ray imaging using GEANT4 Monte Carlo simulations. *Physics in medicine and biology*, 56:2407–21, 2011.

44) P Schotanus, CWE Van Eijk. Temperature dependence of BaF 2 scintillation light yield. *Nuclear Instruments and ...*, 238:564–565, 1985.

45) Photonis. XP2242B. 7–9,

46) K S Shah, J Glodo, M Klugerman, W M Higgins, T Gupta, P Wong. High Energy Resolution Scintillation Spectrometers. *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, 51:2395–2399, 2004.

47) M Kolsteinb, A Airapetiand, N Akopovd, M Amariand, V Garibiand, S Taroiand, A Simonf, B Brayg, M Doetsb, T Henkesb, R Avakiand, A Avetissiand. Performance of F101 radiation resistant lead glass shower counters. *Nuclear Instruments and Methods in Physics Research A*, 378:155–161, 1996.

48) Scionix. Crystal Data. *http://scionix.nl/crystals.htm*, Accessed: May 29, 2012.

49) R Novotny. Performance of the BaF2-calorimeter TAPS. *Nuclear Physics B-Proceedings Supplements*, 137–142, 1998.

50) R Novotny, W Doring. A plastic-BaF2 phoswich telescope for charged/neutral particle and photon detection. *Nuclear Science, ...*, 1996.

51) HiSparc. HiSparc. *http://hisparc.nl*, Accessed: May 1, 2012.

52) Libusb authors. Libusb. *http://www.libusb.org/*, Accessed: July 30, 2012.

53) PyFTDI authors. PyFTDI. *https://github.com/eblot/pyftdi*, Accessed: May 30, 2012.

54) David Fokkema, Arne Laat. PySparc. *https://github.com/HiSPARC/PySPARC*, Accessed: June 30, 2012.

55) Gerhard Kraft. Tumor Therapy with Heavy Ions.

56) D Schulzertner, a Nikoghosyan, C Thilmann, T Haberer, O Jakel, C Karger, G Kraft, M Wannenmacher, J Debus. Results of carbon ion radiotherapy in 152 patients1. *International Journal of Radiation OncologyBiologyPhysics*, 58:631–640, 2004.

57) P Zygmanski, K P Gall, M S Rabin, S J Rosenthal. The measurement of proton stopping power using proton-cone-beam computed tomography. *Physics in medicine and biology*, 45:511–28, 2000.

58) Eike Rietzel, Dieter Schardt, Thomas Haberer. Range accuracy in carbon ion treatment planning based on CT-calibration with real tissue samples. *Radiation oncology (London, England)*, 2:14, 2007.

59) M. Benedikt, a. Wrulich. MedAustron—Project overview and status. *The European Physical Journal Plus*, 126:2011.