Development of a Novel Transverse Beam Profile and Emittance Monitor for the CERN Proton Synchrotron

A thesis submitted to The University of Manchester for the degree of Doctor of Philosophy in the Faculty of Science and Engineering.

2020

Hampus Sandberg

Department of Physics and Astronomy
in the School of Natural Sciences

CERN-THESIS-2020-236
//2020

Contents

Word Count: 54984

List of Figures

List of Tables

Abstract

Beam profile and emittance monitoring is essential to understand the dynamic behaviour of the ensemble of particles in an accelerator beam. A non-invasive measurement of the beam profile ensures reliable measurements. One such device is the beam gas Ionisation Profile Monitor (IPM), which relies on detecting the ionisation products from the interaction between residual gas molecules in the vacuum of the accelerator and the beam particles.

Traditional detectors in IPMs have limited the instruments reliability and performance. A novel approach using Timepix3 pixel detectors is explored in this thesis project which enables direct detection of ionisation electrons with a precise time resolution. A radiation hard readout system was developed to allow the maximum detection performance of the pixel detectors and beam measurements were recorded during 2018.

A beam profile reconstruction method was developed that takes advantage of the information recorded by the Timepix3 detectors to identify and only select events that are associated with ionisation electrons. From these events, a beam profile was reconstructed and the beam size calculated by fitting a Gaussian model to the beam profile data or by calculating the RMS beam size directly.

During 2018 a prototype IPM with pixel detectors was installed in the Proton Synchrotron at CERN. Beam profile measurements recorded with this instrument demonstrated the ability to measure the beam profile continuously throughout the beam cycle in the accelerator. Expected beam dynamics effects such as adiabatic damping and oscillations during transition crossing were observed with the instrument. The time resolution of the pixel detectors enabled bunch-by-bunch measurements of the beam profile by integrating the recorded events separately for each bunch over multiple turns.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright Statement

1. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the "Copyright") and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
2. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
3. The ownership of certain Copyright, patents, designs, trademarks and other intellectual property (the "Intellectual Property") and any reproductions of copyright works in the thesis, for example graphs and tables ("Reproductions"), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
4. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=24420), in any relevant Thesis restriction declarations deposited in the University Library, The University Library's regulations (see http://www.library.manchester.ac.uk/about/regulations/) and in The University's policy on Presentation of Theses.

Acknowledgement

This PhD journey would not have been possible without the help and expertise from a large number of people at CERN and in the scientific community, which I am grateful to now be a part of. First and foremost I would like to thank my two supervisors, James from CERN and Will from the University of Manchester, whose comments and suggestions made this thesis a reality. Their patience has been invaluable during this PhD and their feedback and suggestions in times when I got stuck helped bring this journey to an end. I am grateful for the time you spent teaching me physics concepts I had never heard of before and for listening and asking all kinds of questions when I presented my ideas to you.

I also want to acknowledge all the people who played a part in making the PS-BGI instrument a reality and who are a testament to the success of a project when multiple disciplines of engineering and science come together to achieve a common goal. This includes among others; Swann for designing the detector, Gerhard for managing all the mechanical aspects of the instrument, Dominique for the magnet design and finally Kenichirou and Mariusz for their expertise in all aspects of IPMs.

Finally, I want express my gratitude to all the people at CERN who I have had the opportunity to meet and exchange ideas with; Xavi, Jerome and Michael from the Medipix group, the beam physicists Alex, Matt and Alex, the PS machine operators and all the people in the BE-BI-BL section where I have spent these past four years of my life. For those I forgot to mention, thank you also!

1 | Introduction

Beam instrumentation is essential in particle accelerators to allow the operators to observe, monitor and tune various properties of the particles in the beam and the accelerator machine in general. One such property is the distribution of the particles in the beam, the emittance, which can be observed as a 1D projection, the beam profile, from which the beam size can be extracted. The shape and density of the beam is important to quantify and monitor as it limits the probability of observing a physical process with a given cross section in the experiments, which the accelerator delivers beam to. A dense beam (low emittance) increases the chance of an interaction in a given space and time and therefore results in a higher luminosity.

At the European Organization for Nuclear Research (CERN) particle accelerators have been in operation since 1957, including different types of beam instrumentation. In the next section of this chapter the historical context of the thesis will be outlined, after which requirements for the next generation of instruments for transverse profile monitoring and currently available instruments will be discussed in section 1.2. To meet these new requirements, a synergy between modern microelectronics and traditional beam instrumentation design is explored in this project, starting with a brief historical view of hybrid pixel detectors in section 1.3.

In chapter 2 the basic theory of beam diagnostics and beam dynamics will be detailed, with an emphasis on the transverse distribution of particles in the beam, i.e. the transverse beam profile and size. At the end of the chapter, a beam profile monitor based on beam gas ionisation will be introduced and the novel synergy with hybrid pixel detectors will be detailed.

Chapter 3 introduces the specifics of the implemented beam profile monitor, going over aspects of the mechanical design, the detector construction and the readout electronics and software. Specific consideration of the expected response of the monitor with the pixel detectors will be detailed, including expected performance

metrics.

The interpretation of the output from the pixel detectors in the context of a beam profile monitor will be discussed in detail in chapter 4. A step by step procedure to select and process the events that are generated by the detectors will be shown, highlighting the strengths and weaknesses of this novel implementation.

In chapter 5 examples of beam dynamic studies, using the implemented beam profile monitor, will be presented. Six different measurements will demonstrate the unique capabilities of the monitor and discuss how it can provide better insight to beam diagnostics.

In chapter 6 potential future applications of the instrument will be discussed, focusing on the expected improvements but also potential issues that need to be solved. Finally, suggestions for future development of pixel detectors will be discussed based on the experience of using pixel detectors in a beam profile monitor application.

1.1 Historical Context

The CERN accelerator complex, shown in fig. 1.1, is composed of multiple accelerators that are connected to each other and to different experiments. The Large Hadron Collider (LHC) is the largest accelerator with the highest energy and is the last in the chain at the moment. Particles are first accelerated in one of two linear accelerators, Linac 2 for protons and Linac 3 for ions. In the proton case, they are then injected into the Proton Synchrotron Booster (PSB) which further accelerate the protons before delivering them to experiments or injecting them into the Proton Synchrotron (PS). The PS is the oldest accelerator in operation at CERN and also the most versatile in terms of the different types of beams it can produce. Ions are injected into the PS from the Low Energy Ion Ring (LEIR) and the PS can alternate between accepting ions from the LEIR and protons from the PSB from one beam

cycle to the next, which can last from 1.2 s to 3.6 s. From the PS the particles can be delivered to different experiments, as seen in fig. 1.1, or they can be injected into the Super Proton Synchrotron (SPS). The SPS accelerates the particles even further and delivers them to more experiments or to the LHC. Two counter circulating beams are injected to the LHC, which are further accelerated and then collided at centre of mass energies of up to 13 TeV.

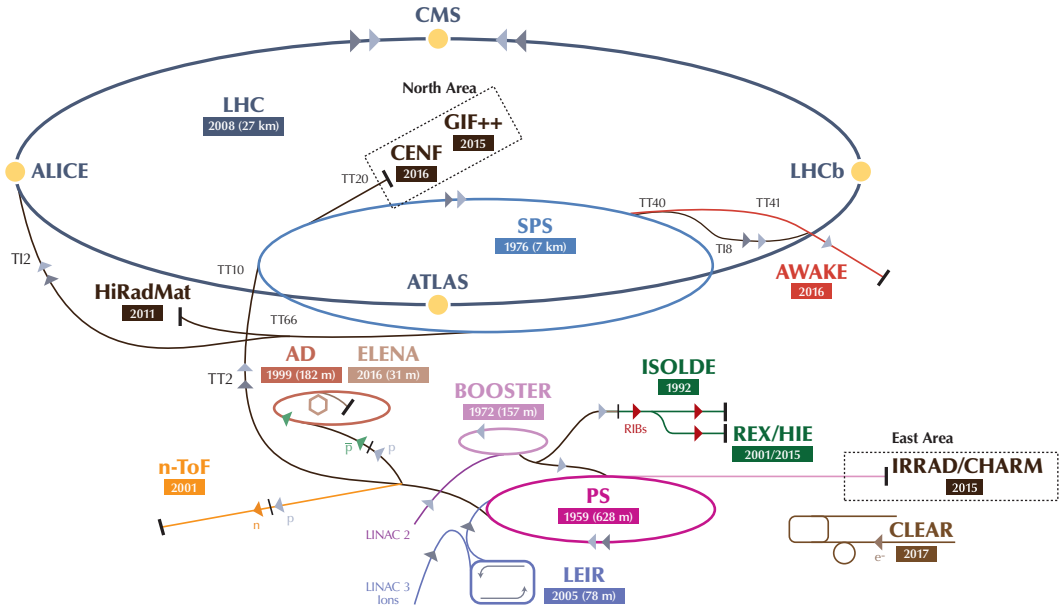

Figure 1.1: Overview of the CERN accelerator complex. Image from [1].

The Linacs, PSB, PS and SPS make up the LHC injector chain and they are referred to as the LHC injectors. Major upgrades are planned for the LHC in the 2020s to increase the luminosity by a factor of five [2]. This new configuration of the LHC is therefore called the High Luminosity LHC (HL-LHC). The current LHC requires performance improvements in the injector chain to reach the goals of HL-LHC. A parallel project was therefore initiated called LHC Injectors Upgrade (LIU) where the aim is to produce beams in the injectors with double the intensity and more than double the brightness [3]. One major part of the LIU is the replacement of Linac 2 with a new Linac 4 that will inject protons into the PSB at an energy of 160 MeV compared to 50 MeV for Linac 2.

Another subject the LIU project aims to address is to improve the beam diagnostics through upgrades of currently available beam instrumentation and introducing new devices that address limitations in the latter. If the operators of the accelerators can measure the beam, the expected beam parameters can be verified. Additionally, reliable measurements enable the operators to tune the machine and therefore improve the quality of the beam. The beam emittance and size are important measures and in the LIU design report it states that "continuous beam size monitoring during the entire machine cycle is one of the priorities for PS operation" [3]. The currently available instruments to measure beam size in the PS are limited and cannot provide such continuous monitoring during the cycle. They are also invasive measurements, where the instruments interact with the beam and cause disturbances or result in instrument damages.

Non-invasive monitoring of the beam profile and emittance throughout the accelerator cycle in the PS is the goal of this thesis project.

1.2 Transverse Beam Profile and Emittance Monitoring

The beam emittance describes properties of the ensemble of particles that make up the beam in an accelerator. It has six dimensions, two along the direction of the beam (longitudinal) and four perpendicular to the beam (transverse). The transverse emittance can be separated into the horizontal and vertical emittance, where each can be described by a two-dimensional phase space with particle position on one axis and particle momentum on the other. If we project the transverse emittance onto the position axis we get the transverse beam profile in the horizontal and vertical plane. A more detailed description of beam emittance and profile will be given in chapter 2.

The transverse beam profile and emittance can be monitored using different

instruments, such as pepper-pots, wire scanners, SEM grids and beam gas ionisation profile monitors to name a few. A brief overview of some of these instruments will be presented in this section. To meet the requirements of the LIU project, the instrument must fulfill the following requirements:

- Non-invasive measurements of the transverse beam profile.
- Continuous measurements throughout the beam cycle.
- Bunch-by-bunch beam profile measurement at 1 kHz.
- Turn-by-turn beam profile measurement for up to 100 turns.
- Reliable operation to minimise interventions and downtime.

Only one of the instruments mentioned above, the pepper-pot, provides direct measurements of the emittance and the rest measure the beam profile. The pepper-pot is also the only instrument that is completely destructive to the beam as it intercepts the whole beam during a measurement [4]. There are similar methods to directly measure the emittance but they all rely on intercepting the beam by passing it through a slit or some open aperture. These methods therefore do not meet all the requirements outlined above. The emittance can instead be calculated from a measured beam profile, based an a procedure which will be detailed in section 2.2. An overview of instruments that measure the beam profile will be given in the following section.

1.2.1 Wire Scanner

A wire scanner consists of a thin wire that is moved transversely through the beam over a period of time, during which the beam will interact with the wire. The movement can be either linear or rotational and the resolution of the measurement is determined by the wire size and the resolution in the mechanical movement. With a linear scanner the movement of the wire can be very precisely determined using stepper motors and mechanical gearing, but the movement is slow. The usage of these monitors is therefore limited to low intensity beam as the wire will heat up

due to the interaction with the beam during the slow movement. Rotational wire scanners move faster through the beam and reduce the exposure time, which enables their use with high intensity beams, but they require a fast and precise measurement of the rotational angle which can be challenging. At CERN, both linear and fast rotational wire scanners are used [5].

In fig. 1.2, the working principle of a linear wire scanner is illustrated. A motor moves a fork with an attached wire through the beam in a linear motion. The interaction between the wire and the particles in the beam produces a shower of secondary particles as indicated by the red arrows in the figure. Some of these particles will pass through the vacuum chamber and hit a detector consisting of a scintillator and a photomultiplier tube (PMT). A scintillator is a material that emits light via fluorescence as the secondary particles interacts with it and the light is detected by the PMT that converts it into an electrical signal. The number of particles in the beam is proportional to the intensity of the secondary particle shower and therefore proportional to the electrical signal from the PMT. The position of the wire is measured with an encoder that is attached to the motor or the fork. For each wire position the intensity signal from the PMT is recorded and the combination of these two measured values is used to reconstruct the beam profile.

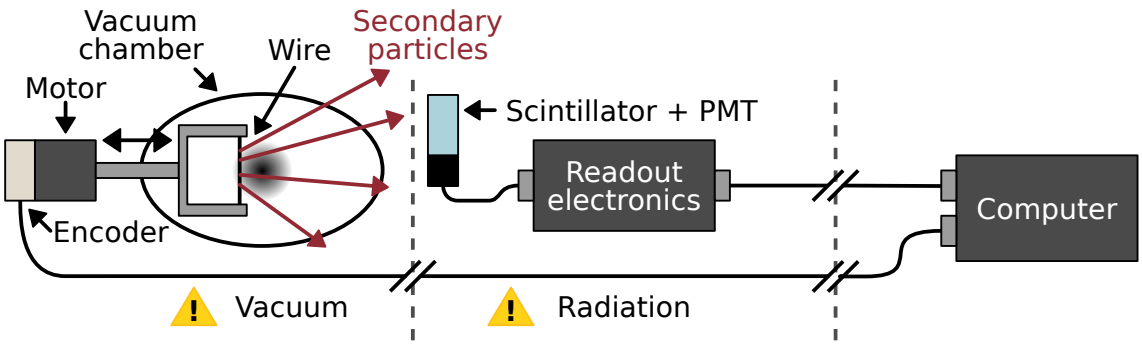

Figure 1.2: Illustration of a linear wire scanner with a motor moving a fork with an attached wire through the beam. The interaction between the wire and the beam particles creates secondary particles that pass through the vacuum chamber and are detected by a scintillator and PMT. The position of the motor is encoded and combined with the intensity of the secondary particle shower to reconstruct the beam profile in a computer.

A beam profile measurement with a linear wire scanner is slow and takes in the order of seconds, while a rotational scanner can measure in a couple of milliseconds. This means that for a wire scanner, each position of the measured beam profile is sampled at different times in the accelerator cycle. For example, in fig. 1.2 where the wire is moving from left to right through the beam, the first sample of the profile on the left will be at turn 1 and the last sample at turn $1 + n$, where n is the total number of turns it takes for the wire to pass through the whole beam. Interactions between the beam particles and the wire causes sublimation of the wire itself and can lead to the wire breaking [6]. Additionally, emittance blow-up can result from the multiple Coulomb scattering of the beam particles with the wire material [7]. The wire scanner is therefore an invasive beam profile monitor. Despite these limitations, it is used as the standard beam profile monitor at the moment in the PS accelerator.

1.2.2 SEM Grid

A secondary electron emission grid (SEM grid) is similar to a wire scanner but instead of a single wire, multiple wires or thin foils are arranged in a grid pattern. The main difference is that the grid is not scanned through the beam, instead the grid is inserted in the beam path before a measurement and retracted afterwards. The principle is illustrated in fig. 1.3 where the motor moves the fork with the wire grid in and out of the beam. Since the fork is not moving during the measurement, there is no need for an encoder. When the beam passes through the grid it will interact with the wires and a secondary emission current can be measured in each individual wire. The number of wires, the spacing between them and their diameter will determine the spatial resolution of the measurement. The pitch of the wires in the grid is usually in the mm range with a wire diameter in the order of 100 µm. Due to the low wire density, most of the particles in the beam will pass through the grid and multiple turns can be measured in a circular accelerator. The main limitation comes from the interaction between the beam and the wires, similar to a

wire scanner. The static position of the grid with respect to the beam causes the wires to heat up and potentially break, resulting in instrument damage that requires breaking the vacuum to fix. In the PS, the SEM grid profile monitors are therefore limited to 30 turns after which the beam is dumped [8].

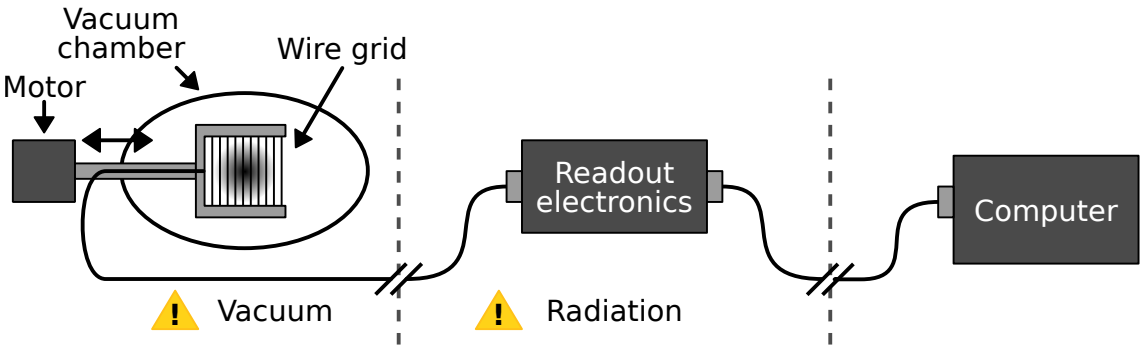

Figure 1.3: Illustration of a SEM grid with a motor moving a fork, with a grid of wires attached to it, through the beam. The interaction between the wires and the beam particles causes secondary electron emission in the wires that is read out as a currents. The wire currents and their fixed positions are combined in a computer to reconstruct the beam profile.

The SEM grid is therefore also an invasive beam profile monitor but it has the advantage over the wire scanner that the full beam profile is sampled in one turn. Two grids that are mounted perpendicular to each other can also be combined in one instrument, which then allows for measurements of the horizontal and vertical beam profile in one passage of the beam. Due to their relatively simple construction and readout electronics, consisting of amplifiers and a multi-channel analog-to-digital converter, they are used extensively for beam profile measurements in low energy accelerators and transfer lines. The newest linear accelerator at CERN, Linac 4, have 30 SEM-grids installed to measure the transverse beam profile at different stages of the accelerator [9].

1.2.3 Beam Gas Ionisation Profile Monitor

A beam gas Ionisation Profile Monitor (IPM) does not interact with the beam directly, as is the case for the wires in the wire scanner and the SEM grid. Instead,

it takes advantage of the ionisation process between the highly energetic particles in the beam (in the GeV range for the PS) and residual gas in the beam vacuum. From this process, electron-ion pairs are created proportional to the density of the beam particles. By measuring the density of the electron-ion ionisation products, the beam profile can be reconstructed. A more detailed description of the working principle of an IPM will be given in section 2.3.

The reference to the first IPM dates back to 1967 and the underlying principle of operation has stayed much the same since then [10]. In 1968 the first IPM in the PS was installed and the monitor was said "to be a most convenient operational method for observing, in an immediately available analog form, the beam position and size during the whole of the machine cycle" [11]. They also concluded that to measure profiles of single bunches, the sampling time should be in the 1 ns range, which could not be achieved with the ionisation electron detector used at the time. This monitor seems to have been removed from the PS at some point and there is no reference indicating that an IPM has been installed since.

Figure 1.4 shows a simplified illustration of an IPM. Inside the vacuum, electron-ion pairs are created close to the beam particles in the centre between two metal plates, which form an electric field cage. Due to the electric field between the plates and the different charge of the electrons and ions, they will start to accelerate towards opposite sides of the field cage, illustrated here as the different colored circles. A detector records the position and the intensity of the ionisation particles that are accelerated onto it and the information is used to reconstruct the beam profile.

The lack of moving parts in the IPM and the non-invasive measurement increase the reliability and enable continuous measurements throughout the beam cycle. It therefore fulfills three out of the five requirements outlined above. For the remaining two requirements, the sampling rates, it depends on how many electron-ion pairs are created for each passing of the beam and whether the detector is able to record enough ionisation electrons in a sufficiently short time window to reconstruct bunch-

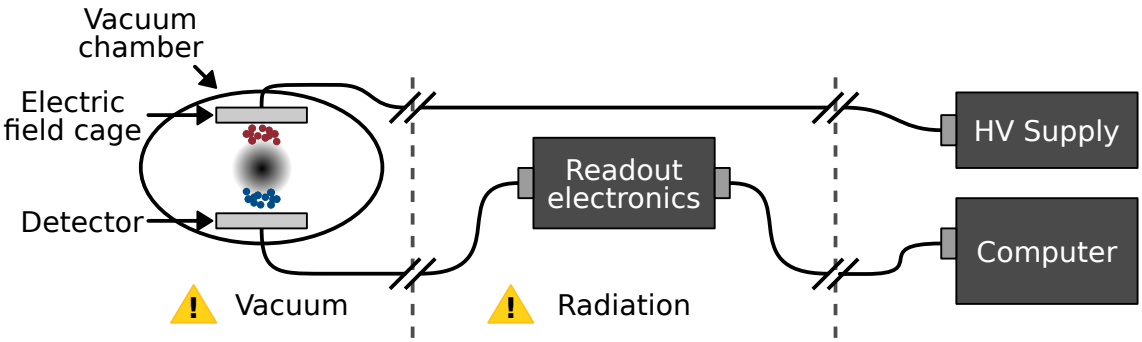

Figure 1.4: Illustration of a beam gas ionisation profile monitor, where the beam particles ionise residual gas inside the vacuum chamber and create electron-ion pairs. An electric field cage separates the electrons (blue) and ions (red) and accelerate the electrons towards a detector where their intensity and position is recorded.

by-bunch beam profiles at the required beam profile sampling rates. In section 2.3.3, different detectors that can be used for an IPM are detailed and discussed, one of which is called a hybrid silicon pixel detector that has the potential to fulfill the sampling rate requirement. As part of the LIU design report, a feasibility study was launched to determine if such pixel detectors could be implemented in an IPM and what the expected performance would be. The study concluded that it was feasible and that the use of hybrid silicon pixel detectors would allow fast bunch-by-bunch measurement of the beam profile in the PS [12].

1.3 History of Silicon Pixel Detectors

In the mid 1980s silicon detectors were widely used in high energy particle physics to construct high spatial resolution detectors [13]. These where one-dimensional silicon strip detectors, constructed as an array of diodes on a silicon substrate where each element was as small as 20 µm in one dimension and several centimeters in the other. A charged particle passing through the detector would create electron-hole pairs in the silicon, which could then be read out as current signals through the diode elements. Each diode was connected to an external low noise amplifier, which also had to be low power due to the thousands of channels needed for a complete

detector. The current signal in combination with the position of the diode elements provided information to reconstruct the passage of the particle through the detector.

The idea of a two-dimensional silicon detector for particle physics experiments was presented in 1987 in [14] as the "micropattern detector", which would expand on the one-dimensional silicon microstrip detectors. The novel idea mentioned in the paper was to incorporate some of the readout and signal processing for each detector element into the detector itself. At the time, charge-coupled devices (CCDs) were commonly used for digital imaging and provided a two-dimensional matrix of detector elements that were read out by shifting the charge from one element to the next to a common one-dimensional readout at the edge of the matrix. An alternative approach discussed in [14] is to locally store the information of an event near the detector elements and read out all events using a common signal at a later time. This would require the integration of analog and digital blocks into an Application Specific Integrated Circuit (ASIC). The paper also discusses whether the diode elements should be integrated in the same silicon substrate as the readout (monolithic pixel detector) or if the readout ASIC and the sensing elements should be separated and bonded together (hybrid pixel detector). Considering the advantages and disadvantages of both, there was no clear preference for one or the other.

Later, in 1994, results were presented for a first prototype of a hybrid silicon pixel detector, consisting of 16 columns at 500 µm pitch and 63 rows at 75 µm pitch [15]. The readout chip was called "D-Omega-ion" and each pixel was composed of a charge sensitive amplifier, a comparator and digital circuitry. When the output of the amplifier crossed a programmable threshold the comparator triggered and generated a binary event that could be read out through a chain of D flip-flops, which connected all the pixels in one column together. The whole pixel matrix could be readout in about 60 µs. The sensor that was bonded to the readout chip was a 300 µm thick silicon sensor with the same pixel pitch.

This basic concept of analog amplification of the input signal followed by a

comparison against a set threshold and then generating a digital event still holds true for hybrid pixel detectors that are designed today. These basic building blocks are illustrated in fig. 1.5 for a generic detector. In (a), the pixel matrix can be seen with a metallic frame around it for bias connection, which will be discussed more in section 2.4.1, and the readout bonding pads where power and data signals are connected. In (b), a cross section with two pixels in the sensor and one pixel in the readout ASIC is shown. A pixel in the sensor is connected via a solder bump to the input pad of the analog front-end in the readout chip where the signal is compared against a threshold to create a digital event.

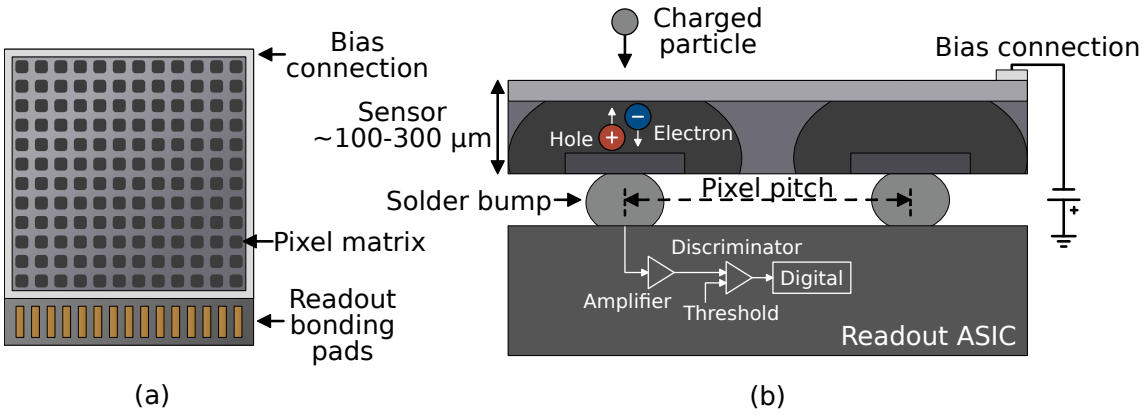

Figure 1.5: Illustration of a hybrid pixel detector with a sensor seen from above (a) and in a cross section view with two pixels visible (b).

In 1997, the Medipix collaboration was created with CERN as one of the collaborators [16]. Over the years, several readout ASICs for hybrid pixel detectors have been developed by the collaboration with more and more features built into the pixels and with faster readout in each iteration. The first Medipix1 ASIC was submitted for manufacturing shortly after the collaboration was formed and contained a matrix of 64 x 64 pixels with a pitch of 170 µm. This was followed a few years later by the Medipix2, which had an increased matrix size of 256 x 256 pixels and a smaller pitch of 55 µm [17]. These Medipix chips were targeted for medical imaging applications involving counting particles and were designed with a configurable analog front-end in each pixel that could be tuned for the specific applications.

In 2005, the design of a new family of ASICs called Timepix was started, where the focus was on extracting as much information as possible from each event in the pixel matrix. The first chip, Timepix, was based on Medipix2 with changes to the pixel front-end which added additional timing information to events based on a common clock in the chip. This new pixel architecture meant that the pixels could operate in one of three modes: (a) counting the number of particle hits, (b) counting the number of clock cycles from when a pixel is activated to the end of the acquisition (Time of Arrival) or (c) counting the number of clock cycles the signal is above the threshold (Time over Threshold). At the same time, the next Medipix3 chip was developed with further improvements to the analog front-end to decrease noise and special features to combine charge from multiple pixels.

In 2013, a new ASIC from the Timepix family called Timepix3 was released, which was the first chip to provide a data-driven readout, where pixel events are pushed out of the chip shortly after they are created, instead of the frame based readout that had been used for all chips up until then [18]. This new readout scheme enabled up to 40 Mhits/s/(cm^2) to be processed. Additionally, each event could include a record of the Time of Arrival (ToA) with a resolution of 1.5625 ns and Time over Threshold (ToT) with a resolution of 25 ns. In this thesis project, the Timepix3 has been used in the novel application of a transverse beam profile monitor and the application will be detailed in section 2.3.3. More specific details about the Timepix3 will be discussed in section 2.4.2.

In the Timepix family it looks as though one chip, the Timepix2, is missing in the timeline. To fill this gap, a Timepix2 chip is currently under development that aims to provide an iterative improvement to the original Timepix by keeping the old frame based readout but incorporating improvements to noise and performance introduced in Timepix3 [19]. The next generation of the Timepix family, the Timepix4, is currently being prototyped and will be discussed later in section 6.2.2.

2 | Theoretical Basis of Beam Diagnostics

This chapter will introduce the beam diagnostics theory that will be referenced and used throughout the thesis. It will start with an overview of beam dynamics concept that are related to beam profile monitoring in section 2.1.

The relationship between transverse beam profile and emittance will be discussed in section 2.2. In some situations the beam profile can be used to diagnose the beam while in other cases the emittance is needed and care must be taken to minimise systematic errors that can be introduced while calculating the emittance from the profile.

Section 2.3 will introduce the working principle of beam gas ionisation profile monitors and the detectors used with them. Finally, in section 2.4 hybrid pixel detectors are detailed with a focus on their application in IPMs.

2.1 Beam Dynamics

This section provides an overview of the theory of beam dynamics, a general broad topic, from a beam diagnostics point of view. We will start the discussion in section 2.1.1 with a description of some of the electromagnetic components in an accelerator and describe how they can be modeled. After this, we will introduce a single particle into the machine and develop a linear model of its transverse motion in section 2.1.2. The last step in section 2.1.3 will expand the discussion to consider an ensemble of particles and how they behave in the machine.

2.1.1 Particle Accelerator Basics

We will focus the following discussion on circular accelerators and in particular, synchrotrons. A circular coordinate system, as seen in fig. 2.1, is commonly used to describe the location of elements in the accelerator and the position of the particles. The dashed line in the figure defines the design orbit and it is the path an ideal reference particle will follow. In reality, the design orbit is not a perfect circle as depicted in the figure, but instead it consists of several straight sections and arcs, which we will approximate as a circle with radius ρ. A local coordinate system (x, y, s), who's origin follows the design orbit in the s-direction, is used to describe positions of particles in the machine.

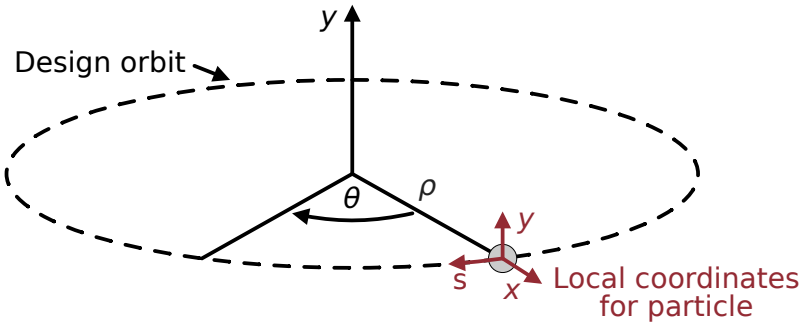

Figure 2.1: Circular coordinate system with the ideal design orbit on the dashed line. The Cartesian coordinates x, y and s define a particles deviation from the design orbit.

To keep the particles in the orbit around the machine, a force perpendicular to their direction of travel towards the centre is needed. With electric $\vec{E}$ and magnetic $\vec{B}$ fields a particle with $\vec{v}$ will experience a force know as the Lorentz force. If the net force is constant we can relate it to the momentum change $d\vec{p}$ of the particle over a time interval dt as:

$$\frac{d\vec{p}}{dt} = q \cdot (\vec{E} + \vec{v} \times \vec{B}), \tag{2.1}$$

where $\vec{p} = \gamma m_0 \vec{v}$ is the relativistic momentum. In most machines the velocity

of the particles is close to the speed of light, $v \approx c$. The main dipole magnets used in the CERN PS have a maximum strength of about 1.5 T. If we were to replace the magnetic field from these magnets with an electric field that exerts an equivalent force it would have to have a strength of 450 MV/m, which is not feasible because current technology limits the maximum electric field to 1 MV/m. Magnets are therefore used all over the accelerator to apply guiding fields for the particles in the transverse directions x and y.

Magnets in Accelerators

There are many different types of magnets used in accelerators with the two most common ones being dipole and quadrupole magnets. Dipole magnets are used to guide the particles around the design orbit with a radius ρ using a magnetic field in the y-direction (B_y). The centripetal force on the particle from the field is in equilibrium with the inertial centrifugal force such that

$$B_y \rho = \frac{p_s}{q}, \tag{2.2}$$

where p_s is the relativistic momentum of the particle in the direction of the orbit. The product of the magnetic field and the radius ($B_y\rho$) is called magnetic rigidity.

Quadrupole magnets have a field such that a particle that passes through it will have a different force applied to it depending on where in the field it is. The force from an ideal quadrupole magnet can be summarised as:

$$F_x = -q \cdot K \cdot x, \tag{2.3}$$

$$F_y = q \cdot K \cdot y, \tag{2.4}$$

$$F_s = 0, \tag{2.5}$$

where K is the magnetic field gradient. At the centre $(x, y, s) = (0, 0, 0)$ there is no

force, but the further away from the centre the stronger the force is. The opposite sign of F_x and F_y leads to focusing in one direction and defocusing in the other depending on the orientation of the magnet. In fig. 2.2 (a) a focusing quadrupole is shown that focus in x and defocus in y, while in (b) it is the opposite. Focusing in both transverse directions can be achieved by combining focusing and defocusing quadrupoles after each other. The combination of a focusing (F) quadrupole, a drift space (O), a defocusing quadrupole (D) and another drift space creates a so called FODO cell. FODO cells are used to create regular structures in an accelerator to keep particles focused all the way around the ring.

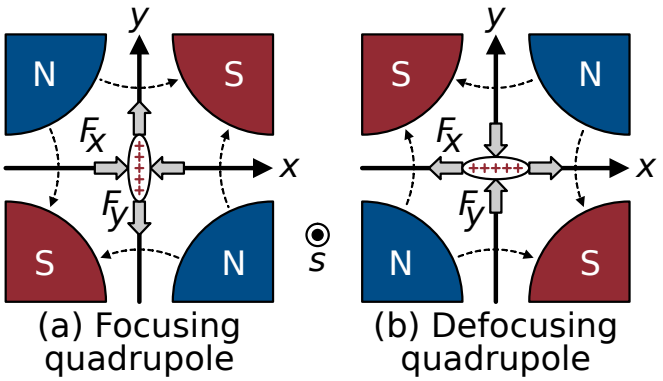

Figure 2.2: Quadrupole magnet in focusing (a) and defocusing (b) configuration.

A circular accelerator can be constructed with dipoles and FODO cells which then defines the optics of the machine. As discussed earlier, beam diagnostic instruments are also needed in an accelerator and have to be included in the optics model, either as simple drift spaces or more detailed models if the instrument for example includes a magnet. There are also higher order magnets in accelerators that are used to correct for non-linear effects which will not be discussed in this thesis.

2.1.2 Transverse Linear Motion

In this section we will introduce a description of the transverse motion of individual particles that pass through the electromagnetic fields of an accelerator. We will limit the discussion to the linear case and neglect effects of particles interacting

with each other in a beam. The framework used here is based on Hamiltonian mechanics to describe the dynamics of the system and the goal is to find a function (the Hamiltonian) that can be used to describe the motion of the particles and to solve it. The details of deriving the Hamiltonian for different accelerator elements and solving them can be found in [20]. For the discussion here, we note that a model of a charged particle's motion through a drift space, a dipole or a quadrupole can be simplified using linear approximations.

Single Particle System

Each element in the machine can be represented as a transfer map, which relates the values of the variables in the motion equation at the entry of the element to the values at the exit. The variables we are interested in for the transverse motion are the horizontal coordinate and momentum (x, p_x) and the vertical coordinate and momentum (y, p_y). The motion in a whole accelerator can then be described by combining the transfer map of every element and solving the equations. Such rigorous analysis can be simplified by describing the beam dynamics at a higher level under the assumption of a periodic beam line. A machine that is built up of FODO cells in a periodic pattern leads to transverse oscillations of particles due to the transverse focusing and defocusing properties of the quadrupoles in a cell. If we track a particle that enters an arbitrary chosen FODO cell in the machine, with some initial values $(x_0, p_{x0}, y_0, p_{y0})$, and observe the new values as it exits each cell we will see that it follows an ellipse in the horizontal and vertical phase space as illustrated in fig. 2.3 [20].

We can describe the change in phase space from the starting point $\vec{u_0}$ before the FODO cell to the point $\vec{u_1}$ after, with the linear transfer map

$$\vec{u_1} = R\vec{u_0}, \tag{2.6}$$

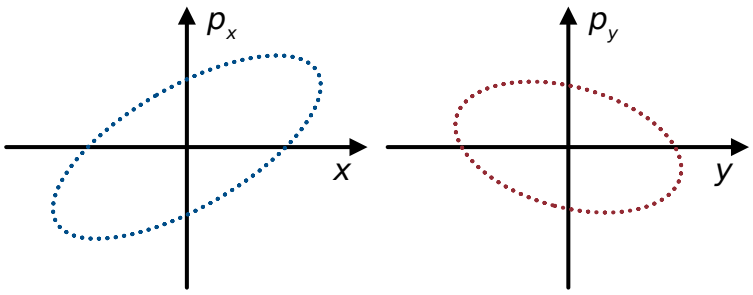

Figure 2.3: Illustration of the phase space for the horizontal coordinate and momentum (x, p_x) (left) and the vertical coordinate and momentum (y, p_y) (right) of a single particle, where each point is the value at the exit of a FODO cell in a periodic lattice.

where R is the transfer matrix for the cell. If we only consider FODO cells in our machine, i.e. no dipoles, we can simplify the transfer matrix by assuming no coupling between between horizontal, vertical and longitudinal beam dynamics as

$$R = \begin{pmatrix} R_x & 0 & 0 \\ 0 & R_y & 0 \\ 0 & 0 & R_z \end{pmatrix}, \tag{2.7}$$

where the block diagonals are 2 x 2 matrices. Furthermore, R_x can be written as

$$R_x = \begin{pmatrix} \cos\mu_x + \alpha_x \sin\mu_x & \beta_x \sin\mu_x \\ -\gamma_x \sin\mu_x & \cos\mu_x - \alpha_x \sin\mu_x \end{pmatrix}, \tag{2.8}$$

where α_x, β_x and γ_x are the Courant–Snyder parameters [21] (also referred to as the Twiss parameters) and μ_x is a parameter called the phase advance. We can now define an action variable J_x that is invariant along the beam line as

$$J_x = \frac{1}{2}(\gamma_x x^2 + 2\alpha_x x p_x + \beta_x p_x^2). \tag{2.9}$$

For fixed J_x, eq. (2.9) defines the ellipse in phase space we saw in fig. 2.3 and the shape is constrained by the Twiss parameters. By solving eq. (2.9) for $x = p_x = 0$ and

finding the minimum and maximum values we can identify how the Twiss parameters affect the shape of the ellipse. Additionally, we can introduce the angle ϕ that describes where on the ellipse a particle is. In fig. 2.4 an example is shown for an ellipse in the horizontal phase space, illustrating how the Twiss parameters define the ellipse. Using the geometric relationships in fig. 2.4, we can define the position

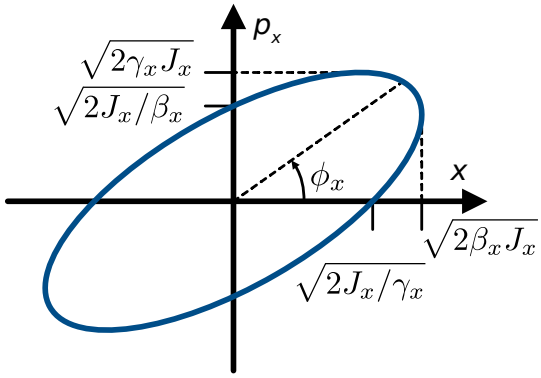

Figure 2.4: Illustration of how the Twiss parameters define the ellipse in phase space.

x and momentum p_x as functions of the action J_x and the angle ϕ_x as

$$x = \sqrt{2\beta_x J_x}\cos\phi_x, \tag{2.10}$$

$$p_x = \sqrt{\frac{2J_x}{\beta_x}}(\sin\phi_x + \alpha_x\cos\phi_x). \tag{2.11}$$

This shows that the transverse position x performs oscillations around the reference trajectory ($x = 0$) with an amplitude that depends on β_x and the angle ϕ_x. β_x is known as the betatron function and has a different value for each position s in the machine. An important parameter in a circular machine is the number of transverse oscillations per revolution, which depends on the values of β_x along the machine, known as the tune

$$Q_x = \frac{1}{2\pi}\oint\frac{ds}{\beta_x}. \tag{2.12}$$

Both a horizontal (Q_x) and vertical (Q_y) tune can be calculated and are common parameters to monitor in a machine using beam position monitors.

2.1.3 Distribution of Particles and Emittance

The discussion so far has focused on how a single particle is affected by the optical elements in the accelerator. We will now consider a group of particles, a bunch, which is what machines commonly operate with. If we make the assumption that the angle variables of the particles in the bunch are uncorrelated with the action variables and that the angle variables are uniformly distributed from 0 to 2π, we can also describe ensembles of particles using the Twiss parameters. In the following discussion the notion $\langle\cdot\rangle$ will be used to indicate the statistical average over all particles in the bunch [20]. At a fixed point in the beam line we can take the statistical average of x^2, xp_x and p_x^2 for all particles using eqs. (2.10) and (2.11), which results in

$$\langle x^2\rangle = 2\beta_x\langle J_x\cos^2\phi_x\rangle = \beta_x\cdot 2\pi\langle J_x\rangle, \tag{2.13}$$

$$\langle xp_x\rangle = -2\langle J_x(\cos\phi_x\sin\phi_x + \alpha_x\cos^2\phi_x)\rangle = -\alpha_x\cdot 2\pi\langle J_x\rangle, \tag{2.14}$$

$$\langle p_x^2\rangle = \frac{2\pi}{\beta_x}\langle J_x(\sin^2\phi_x + 2\alpha_x\sin\phi_x\cos\phi_x + \alpha_x^2\cos^2\phi_x)\rangle = \frac{1+\alpha_x^2}{\beta_x}\cdot 2\pi\langle J_x\rangle. \tag{2.15}$$

We can see that $2\pi\langle J_x\rangle$ is common for all three equations above and this expression is the area of the phase space ellipse. We will therefore define the horizontal emittance of the bunch as

$$\epsilon_x = 2\pi\langle J_x\rangle. \tag{2.16}$$

By rearranging eqs. (2.13) to (2.15), the emittance can be expressed as

$$\epsilon_x = \sqrt{\langle x^2\rangle\langle p_x^2\rangle - \langle xp_x\rangle^2}. \tag{2.17}$$

If the values of $\langle x^2\rangle$, $\langle xp_x\rangle$ and $\langle p_x^2\rangle$ are known at a specific point in the beam line, the emittance can be calculated using eq. (2.17) and the Twiss parameters can be extracted from eqs. (2.13) to (2.15). For an ensemble of particles in a bunch, the emittance is a measure of the area in phase space that the particles occupy. There

are several different definitions of emittance in literature, but for the discussion here we will use the definition in eq. (2.16). The action variable J_x is invariant along the beam line, which means that the emittance is preserved, under the assumption of linear transport of particles that we have been using.

Acceleration of Particles

Accelerating particles means that their energy is increased, which is done using radio frequency (RF) cavities that create oscillating electromagnetic fields. This is a complex topic described in depth in [20]. The discussion in this thesis will only briefly touch on the general concepts of acceleration and how the beam dynamics is affected. From the Lorentz force in eq. (2.1), we can see that an electric field in the same direction as the velocity of the particle is needed to accelerate it. Such an electric field leads to an energy transfer from the RF cavity to the particle.

By introducing RF cavities into the lattice of the machine we break the conserved and static conditions used in the discussion up until now. In practice, this means that the energy increase from acceleration leads to a change in the reference momentum, which results in a decrease in the area of the phase space. This effect is called adiabatic damping. We therefore expect to see a change in the emittance when the particles are accelerated. The change in the emittance from the initial (ϵ_{x0}) to the final (ϵ_{x1}) can be expressed as

$$\beta_{r1}\gamma_{r1}\epsilon_{x1} = \beta_{r0}\gamma_{r0}\epsilon_{x0}, \tag{2.18}$$

where β_{r0}, γ_{r0}, β_{r1} and γ_{r1} are the relativistic parameters. During acceleration, the quantity $\beta_{\mathrm{r0}}\gamma_{\mathrm{r0}}\epsilon_{x0}$, called the normalised emittance, is preserved and not the so called geometrical emittance ϵ_{x0}. We will denote the normalised emittance as $\epsilon_{x,N}$.

Coupling Between Transverse and Longitudinal Motion

Another important set of optical elements we have neglected in the discussion thus far are dipole magnets, which have the important function of guiding the particles around a circular machine. When we introduce dipoles in our machine we also introduce coupling between the transverse planes and the longitudinal because the bending angle depends on the energy of the particle. This causes a change in the trajectory from the reference for any particle with a relativistic momentum deviation δ_p from the reference momentum defined as

$$\delta_p = \frac{P}{P_0} - 1, \tag{2.19}$$

where P is the total relativistic momentum of the particle and P_0 is the reference momentum. This dependency between change in trajectory and momentum can be described by the dispersion function D, which depends on the position s in the machine, similar to the betatron function. The dipoles bend the particle trajectories in the horizontal plane, which will lead to a dispersion function D_x, while in the vertical plane there is no bending and the vertical dispersion D_y is therefore often neglected. In a real accelerator, vertical dispersion may arise due to alignment errors of the magnets that causes coupling between the horizontal and vertical plane, or by special elements such as solenoids and skew quadrupoles.

Equation (2.10) that described the position x of a particle, did not take into account this change in trajectory, but instead assumed the particle would follow the reference trajectory. We therefore need to add the contribution of dispersion to the trajectory as

$$x = \sqrt{2\beta_x J_x}\cos\phi_x + D_x\delta_p. \tag{2.20}$$

With the same assumptions used in eq. (2.13) and additionally assuming that the

momentum deviation δ_p is uncorrelated with the action variables we can calculate

$$\langle x^2 \rangle = \beta_x \epsilon_x + D_x^2 \langle \delta_p^2 \rangle, \tag{2.21}$$

where ϵ_x is the geometrical emittance defined in eq. (2.16). Additionally, if the distribution of particles in the beam is Gaussian in the horizontal and longitudinal plane, with $\langle x \rangle = \langle \delta_p \rangle = 0$, we can write

$$\sigma_x^2 = \beta_x \epsilon_x + D_x^2 \sigma_{\delta_p}^2, \tag{2.22}$$

where σ_x is the horizontal RMS beam size and σ_{δ_p} is the longitudinal RMS momentum spread. These two RMS parameters can be measured in an accelerator and it is therefore common to rearrange eq. (2.22) to calculate the geometrical emittance of the beam as

$$\epsilon_x = \frac{1}{\beta_x} \left(\sigma_x^2 - D_x^2 \sigma_{\delta_p}^2 \right). \tag{2.23}$$

It is important to note that this emittance calculation is only valid under the assumptions above [20]. The normalised emittance follows as

$$\epsilon_{x,N} = \frac{\gamma_r \beta_r}{\beta_x} \left(\sigma_x^2 - D_x^2 \sigma_{\delta_p}^2 \right). \tag{2.24}$$

2.2 From Transverse Beam Profile to Emittance

From the discussion above we can conclude that eq. (2.21) describes the size of the particle distribution (the beam size). If we have a look at eq. (2.21) again we can identify that the beam size is composed of two components as

$$\langle x^2 \rangle = \underbrace{\beta_x \epsilon_x}_{\text{Betatronic}} + \underbrace{D_x^2 \langle \delta_p^2 \rangle}_{\text{Dispersive}}. \tag{2.25}$$

If the dispersive component is small, which is usually true in the vertical plane (y), the beam size only depends on the emittance and the betatron function. Otherwise the dispersive component will have an effect on the beam size, for example in a synchrotron with bending in the horizontal plane. A common method to reduce this effect for beam profile monitors, which are the devices in an accelerator that measure the beam size, is to place these instruments at locations in the ring where the dispersion function D_x is at a minimum. Another option would be to reduce the momentum spread δ_p of the beam, but this might not be possible or in some cases not wanted.

Under the assumptions of Gaussian distributions we can use eq. (2.22), where the RMS beam size σ_x and the longitudinal rms momentum spread σ_{δ_p} are used instead of the statistical averages $\langle\cdot\rangle$. If this does not hold true, the beam size is a convolution of a betatronic distribution and a dispersive distribution and would require a full deconvolution method to separate them. An example of such deconvolution is shown in [22], where it was known prior to the measurement that the longitudinal distribution was not Gaussian. This thesis project will not use deconvolution, instead a method will be used to quantify if a measured transverse beam distribution is Gaussian or not.

The goal of a beam profile monitor is to measure emittance by calculating it from the measured beam size at the location of the instrument. This can be done using the Gaussian assumptions and eq. (2.24) for the normalized emittance, which is

$$\epsilon_{x,N} = \frac{\gamma_r \beta_r}{\beta_x} \left(\sigma_x^2 - D_x^2 \sigma_{\delta_p}^2 \right). \qquad \text{(2.24 revisited)}$$

We therefore need to measure or calculate the parameters in the equation. First of all, we will measure the beam size σ_x using the beam profile monitor. Secondly, the optics function β_x and D_x can either be calculated from a model of the lattice or by measuring them. It is important to note that the optics values are only valid for

a specific location in the accelerator and we should therefore calculate or measure them at the location of the profile monitor that we are using. Lastly, the momentum spread δ_p must be measured and can vary from one type of beam to another as well within a beam cycle. We will discuss these three steps below, which are needed to calculate the emittance.

Before proceeding, we will first look at some examples of these values seen in table 2.1 for a beam profile monitor at a fixed location in the CERN PS. The values are shown here in commonly used units, and emittance is usually specified in units of [µm] or [mm mrad]. If we insert the numbers from the table for the different beam types into eq. (2.24) we can calculate the normalized emittance at the location of the instrument as seen in rightmost column. We can also do the calculation in reverse and get the expected beam size if the normalized emittance is known.

Table 2.1: Example of optics values and measured beam size and momentum spread for beams in the CERN PS at extraction, from [23, 24].

Beam type	β_r	γ_r	β_x [m]	D_x [m]	σ_x [mm]	δ_p	$\epsilon_{x,N}$ [µm]
LHC	1.0	28.15	6.2	3.0	2.8	0.9×10^{-3}	2.5
TOF	1.0	21.68	6.6	2.4	6.5	2.6×10^{-3}	10
SFTPRO	1.0	14.95	11.6	2.2	3.8	1.2×10^{-3}	10

Another important detail is that the values we use for the emittance calculation will depend on when in the beam cycle we measure. We saw in eq. (2.24) that the normalised emittance is preserved during acceleration of particles in the machine. During a beam cycle the optics and the momentum spread may change, which means that if we want to monitor the evolution of the emittance during a beam cycle we need the evolution of these parameters also. In [25], emittance measurements are presented for a full beam cycle based on beam profile measurements with an ionisation profile monitor (IPM). This was achieved by modeling the optics at 1 ms intervals and then calculating the emittance for each step. For measurements of the evolution of the emittance, such a rigorous approach is necessary.

The above discussion highlights the connection and dependencies for measuring the emittance from beam profile measurements. It is challenging to measure the emittance in a non-invasive way as it requires accurate measurements and calculations from both beam instrumentation and beam physics. Close collaboration and transparency is therefore crucial to achieve accurate emittance measurement by making sure all the parties involved have an understanding of and trust in the measured quantities.

2.2.1 Beam Size Measurement

The beam size can be measured using for example a wire scanner, a SEM grid or an IPM as shown earlier in section 1.2. As discussed, it is the transverse position x of the particles in the beam that we want to measure for a horizontal profile and the position y for the vertical. If we know the beam profile we can calculate the beam size $\langle x^2 \rangle$ in several ways, two of which are commonly used. Either we fit the profile with a Gaussian model and extract the size as the standard deviation of the Gaussian distribution, after testing if the fit is good, or we calculate the RMS beam size directly from the profile. The advantage of the fitting method is that robust procedures are available for fitting a model to the beam profile data, which can also include corrections to account for noise and distortion. On the other hand, this means that the result depends on the chosen model and if it does not describe the beam particle distribution well, the calculated beam size may not be correct. The RMS beam size is calculated as the standard deviation of all the samples in the profile and does not make any assumption of the shape of the beam profile distribution and therefore do not require a model. However, noise present in the measured beam profile will influence the RMS beam size and must be minimised.

The different beam profile instruments have their own way of measuring and reconstructing the beam profile. In addition, the method for calculating the beam size from the beam profile can vary from one analysis to another based on the

preferences of the individual doing the analysis. This means that there are several ways of achieving a beam size measurement and in an ideal case they would all lead to the same result. In practice there are often discrepancies and it is therefore important to analyse and understand how each step affects the final beam size result. Chapter 4 will detail these steps for the beam gas ionisation profile monitor using hybrid pixel detectors that have been developed and studied in this thesis project.

2.2.2 Optics Calculation and Measurement

As we saw previously in section 2.1.2, the components in an accelerator can be modeled using a transfer matrix where each element in the machine is included. Such models are typically studied in detail before a machine is even constructed, as the layout and configuration of the elements in the machine have to be determined before construction starts. This means that the optics parameters β_x and D_x for all the positions in the machine can be extracted from the model. The main use of the optics model after a machine is built is by the beam physicists to study specific beam dynamics in the machine where non-linear effects are often included. In the context of this thesis we will not go further into the topic of optics modeling.

A model might not include details of a machine that affects the optics, for example due to limited computational power or simplifications of a modeled element. As an alternative, we can instead measure the optics values in a machine. We discussed earlier that the betatron function β_x defines the amplitude of the transverse oscillation at a given location in the machine. If we measure the amplitude of the particles in the beam at different locations in the machine, we can reconstruct the betatron function. The beam position monitors are able to measure this amplitude and are commonly spread out all over the machine. We can therefore use them to measure the transverse oscillation and such a measurement is detailed in [26].

Similarly, the dispersion function D_x can be measured by taking advantage of the change in trajectory due to a momentum offset in the beam. This results in a change

in the mean position of the particles in the beam, which can be measured with a beam profile monitor, with the added advantage that the value of the dispersion function is measured at the instrument itself. In [27], such a measurement is presented using a wire scanner where a linear correlation between beam position and change in momentum is shown. This of course requires a dedicated measurement campaign for each beam profile monitor that will be used to calculate the emittance and changes to the optics in the machine requires a new set of measurements to be done.

We have focused on the horizontal transverse plane, but the same arguments hold for the vertical plane where the optics values can either be calculated or measured. A caveat for the vertical plane is that the dispersion function is usually orders of magnitude lower compared to the horizontal due to the lack of dipole magnets in the vertical plane. As discussed earlier, there could still be dispersion in the vertical plane due to coupling and misalignment of magnets, but the lower amplitude means that it could be impractical to measure the vertical dispersion. This is another reason why it is common to neglect the vertical dispersion when calculating the vertical emittance.

2.2.3 Momentum Spread Calculation and Measurement

Similar to the optics values, the momentum spread can be calculated from a model of the optics that includes the RF cavities or it can be measured. The measurement method is analogous to that of a transverse beam profile monitor as the goal is to quantify the longitudinal positions of the particles in the beam. By using tomography, the longitudinal phase space can be reconstructed from the longitudinal shape of the beam, as detailed in [28]. The basic idea is that snapshots of the 1D longitudinal beam profile are captured over several turns in the machine. For each turn, the longitudinal information in the 2D phase space has rotated and therefore results in a different 1D projection of the longitudinal beam profile. By combining the information from the 1D snapshots, the 2D phase space can be reconstructed.

2.3 Beam Gas Ionisation Profile Monitor

A beam gas ionisation profile monitor (commonly referred to as IPM or BGI at CERN) is an instrument that enables non-invasive measurements of the beam profile continuously throughout a beam cycle. IPMs have been deployed in laboratories since the late 1960s, as discussed in section 1.2.3. Though the design and detector technology has changed over time, the general working principle has stayed the same. IPMs can operate in one of two modes: electron collection or ion collection. In fig. 2.5, a schematic overview of an IPM operating in electron collection mode is shown. The electrons and ions, created from the ionisation process at the centre of

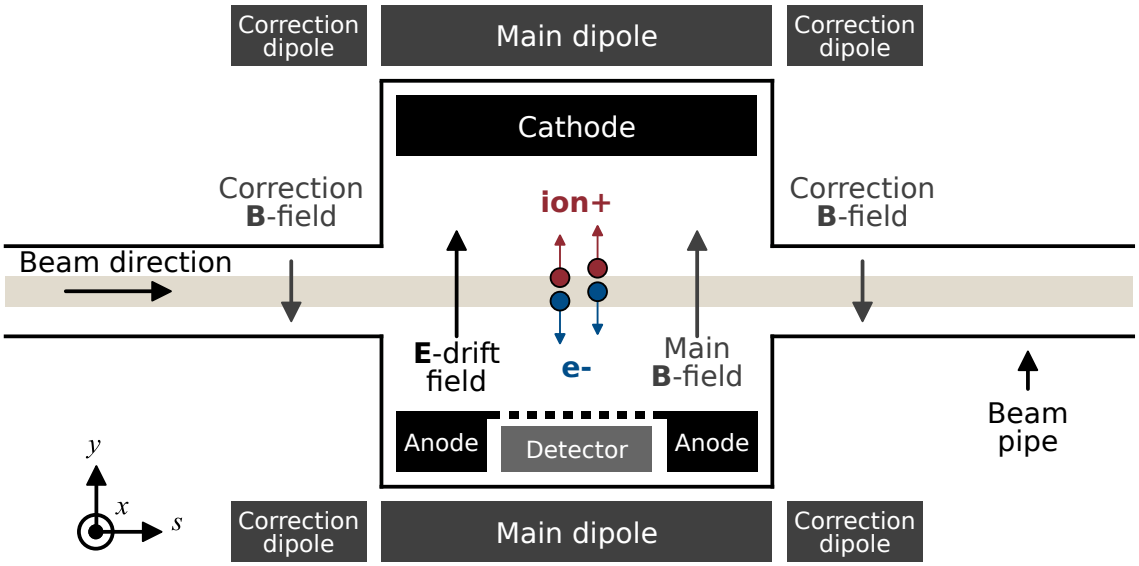

Figure 2.5: Schematic overview of a beam gas ionisation profile monitor for electron detection.

the instrument, are accelerated by the electric drift field with electrons accelerated down towards the detector in this configuration. A dipole magnet (main dipole) is located outside the vacuum chamber of the instrument and creates a magnetic field parallel to the electric drift field. The magnetic field confines the electrons in a helical motion towards the detector and reduces the transverse spatial drift caused by the electric field from the charged particles in the bunch. The magnet is not used for ion collection since the motion of the heavier ions is much less affected by the

electric field of the bunch. If there was only one magnetic field at the centre of the instrument the beam orbit could be perturbed, depending on the strength of the magnet, which is why two additional correction dipoles are used on either side with a magnetic field in the opposite direction and half the strength of the main dipole.

The following sections will discuss the gas ionisation process inside an IPM in section 2.3.1 and the electric and magnetic fields in section 2.3.2. It will end in section 2.3.3 with an overview of commonly used detectors in IPMs, highlighting their advantages and disadvantages.

2.3.1 Gas Ionisation

There are several parameters that can affect how many electron-ion pairs are created in the gas ionisation process in an IPM. One is the length of the detector in the direction of the beam, s in fig. 2.5, as the ionisation occurs all along the beam pipe where the beam particles are interacting with the residual gas. The extreme case would be a detector that covers the whole length of an accelerator, but then there would be no space for other elements which are needed to steer and accelerate the beam. Additionally, IPMs depend on an electric field (and magnetic for electron mode) over the detector where the cost and complexity to create it increases with the detector length.

Another parameter is the background gas composition; how many gas molecules there are and of what type. The IPM requires a minimum amount of gas molecules in the beam pipe, otherwise there would not be any gas ionisation, but if there are too many molecules they can start to affect the beam lifetime and potentially saturate the detector. A way to control the amount and type of gas is to precisely inject gas into the centre of the IPM where the interaction occurs using a gas injection system. More often, the residual gas pressure that is normally present in the beam vacuum pipe is used, which could be on the order of 1×10^{-8} mbar to 1×10^{-10} mbar. The temperature also has an effect as the velocity of the molecules changes, where a slow

moving molecule in low temperature is more likely to be hit by a particle in the beam and create an electron-ion pair.

Lastly, the composition of the particle beam and the particle type have an influence on the ionisation. Different particles have a stronger or weaker interaction with the gas molecules, as defined by the ionisation cross section. With more particles in the beam (higher intensity) there is also a higher probability for ionisation to occur.

The ionisation electron yield per bunch and per turn of the accelerator can be estimated using a formula derived from eq. (5) in [29] as

$$n_{\mathrm{e,b,t}} = L_{\mathrm{s}} \cdot \sigma_{\mathrm{ion}} \cdot N_{\mathrm{b}} \cdot p \cdot \frac{1}{k \cdot T}, \tag{2.26}$$

where L_{s} is the detector length in the beam direction, σ_{ion} is the ionisation cross section of the residual gas, N_{b} is the number of particles per bunch, p is the static pressure, k is the Boltzmann constant and T is the ambient temperature. These are the parameters mentioned above. In section 3.5.2 the expected electron ionisation yield for the IPM implemented in this thesis project will be detailed.

2.3.2 Electric and Magnetic Fields

We saw in fig. 2.5 that the electric field accelerates the electrons (or ions, depending on the direction and collection mode of the IPM) to the detector and that a parallel magnetic field can be added to confine that trajectory. Ideally, we want to transport the ionisation electrons onto the detection plane without distortion to the transverse position of the electron. Additionally, the initial kinetic energies of the ionisation products are low and might be below the minimum threshold for the detector. The electric field therefore serves two purposes, the first is to transport the electrons from the beam centre to the detector and the second is to accelerate the electrons to higher kinetic energy. The magnetic field does not increase the energy of the ionisation electrons, its main purpose is to maintain their transverse position.

Considering that the fields are responsible for translating the ionisation products from the point of creation to the detectors, it is important to ensure the uniformity of the fields over the full detection volume. Several simulation tools have been developed to aid in the design of the electric field cage in an IPM, where field maps can be imported and the transport of the ionisation products simulated [30]. Based on the results from the simulation, changes to the electric field cage design might be required, such as adding additional side electrodes or increasing the electrode size.

In table 2.2, the electric and magnetic field strengths used in IPMs for different accelerator machines are shown. Most of the IPMs can operate in both electron and ion collection mode by changing the polarity of the electric field cage, but the presence of a magnet indicates that electron mode was foreseen.

Table 2.2: Examples of electric and magnetic field strengths used in IPMs at different laboratories.

Laboratory	Machine	Electric Field [kV/m]	Magnetic Field [mT]
CERN	SPS [31]	100	60
Fermilab	Tevatron [32]	115	200
Fermilab	Main Injector [33]	10	100
GSI	ESR [34]	60	80
J-PARC	RCS [35]	86	46

One of the main benefits of using electron mode is that it reduces the drift time, which could be a limitation when a beam with a high repetition rate is used and in bunch-by-bunch measurements. If the ionisation products have not been detected before the beam comes around again and creates more, there will be a pileup and information will be lost. On the other hand, if a magnet is required for electron mode, the cost and the size of the IPM will increase.

2.3.3 Detector and Readout

The detector in an IPM must be able to detect the ionisation products (with kinetic energies of around 10 keV) inside the vacuum of the instrument, as seen in fig. 2.5.

Operating inside the vacuum requires the use of materials that are compatible to minimise outgassing and also requires stringent cleaning procedures of the components. Additionally, high radiation levels inside and close to the instrument further restricts the detector implementation.

The energy of the ionisation products may be too low to detect and therefore requires that the detector is constructed as several stages, where the signal is amplified or converted in each stage to enable detection. This could introduce point spread, which means that the physical location of the ionisation product that we want to measure is distorted from the first stage of the detector to the last. A staged detector could also limit the spatial or temporal resolution of the detection, for example if the amplifier stage has a limited feature size and needs time to recover after a signal has propagated through it. The detector must be read out from inside the vacuum over a distance of up to 200 m in some cases, to a counting room outside the accelerator where the signals are processed and the beam profile reconstructed.

In the following section we will first discuss commonly used detectors in IPMs that are constructed in the staged approach and after that a novel implementation using hybrid pixel detectors will be detailed.

IPM Detector Using Microchannel Plate

A microchannel plate (MCP) is a device that is constructed from a glass substrate with thousands of 10 µm diameter channels cut through it. The MCP requires a potential between the two side which creates an electric field in the channels. When a charged particle enters a channel on the input side of the MCP it will start an avalanche of electrons inside the channel which are expelled from the other side, effectively multiplying the input signal. The gain varies with the applied electric field and the properties of the MCP that is used, with gains in the order of 1×10^4 commonly seen [36]. Several MCPs can also be stacked together to further increase the gain.

After the ionisation products have been amplified by the MCP, the output electrons can be detected either optically or electrically. In fig. 2.6 these two methods are shown. For the optical case (a), a phosphor screen is placed in close proximity to the output side of the MCP, which will output photons when the electrons interact with the screen. This light is then guided through a prism or mirror out of the vacuum to a camera which then images the light. The electrical method (b) instead places strips of metal (anode strips), that are closely spaced together, under the MCP and the electrons will then create currents in the strips. The current is read out from each strip individually via a vacuum feedthrough connector and is amplified close to the instrument and digitised further away from the radiation. These two cases show the general principle of the detectors but slight variations of these can of course be seen in IPMs used in different laboratories. Implementations of optical IPMs can be found in [31, 34, 37, 38] and anode strip IPMs in [32, 39, 40].

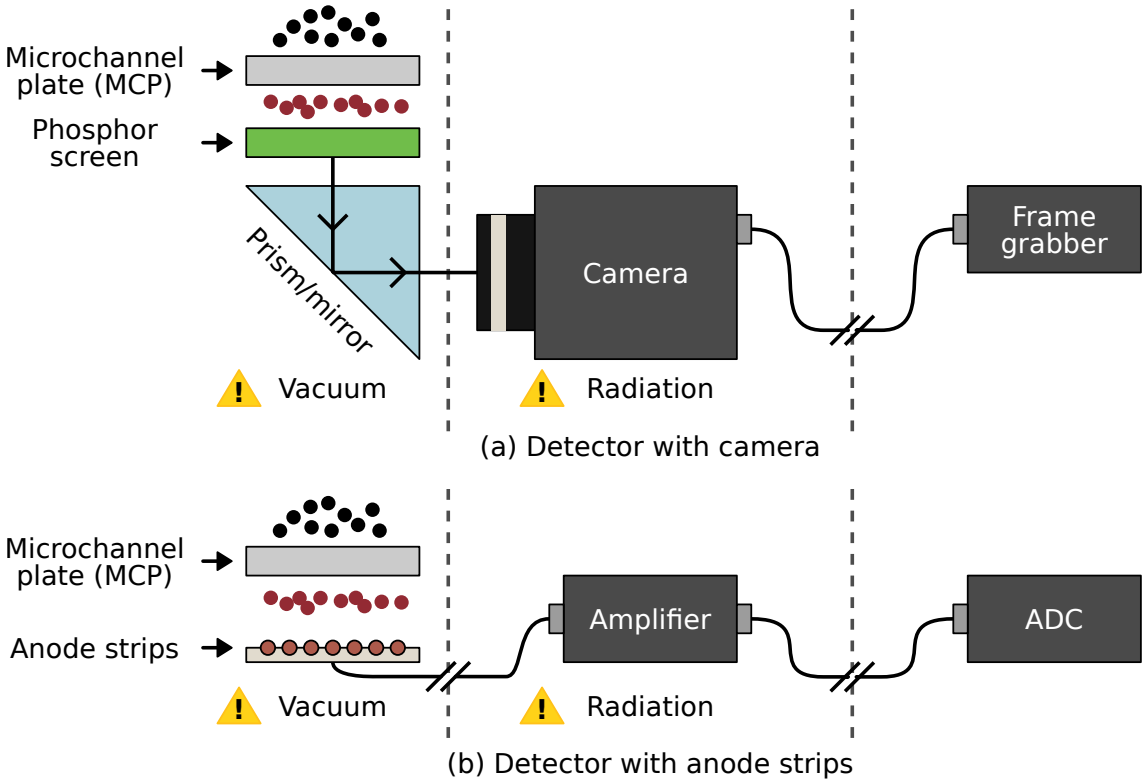

Figure 2.6: IPM detectors using MCPs with a camera (a) or anode strips (b).

The main advantage of the optical readout is spatial resolution, which is limited

by the optical path and the pixel size of the camera (in the 10 µm range), whereas the electrical is limited by the width and spacing of the strips (in the 200 µm range). A challenge with the optical readout is that light intensity is inversely proportional to the square of the distance, which means that to maximise the signal the optical path must be reduced. Ideally, the camera should be mounted inside the vacuum directly after the phosphor screen but this requires a vacuum compatible camera that can also withstand radiation. With the prism inside the vacuum, the camera can be placed outside but the radiation might still be high enough to cause damage to the camera. Additionally, the frame rate of the camera limits the sampling rate of the beam profile and a radiation tolerant high speed camera would therefore be required.

The sampling rate of the electrical strip readout is limited by the analog to digital converter (ADC), the amplifier and the amount of current that is created in the strips. Modern ADCs are capable of more than 1 GS/s per channel, but the large number of channels needed for an IPM will limit the ADC choice to one that fits within the budget of the instrument. For example, to cover a range of 20 mm using strips with a pitch of 200 µm, 100 channels are needed. All these channels have to be routed via vacuum feedthroughs to individual cables and then connected to the ADC. Depending on where the ADC is placed, it might also need to be radiation tolerant which drastically reduces the part availability and performance. Since the electrical signal is analog from the strip to the ADC it is important to maintain the signal integrity by shielding the cables and separating them from noise sources in the accelerator, which adds additional material weight and cost. The biggest advantage with the strip readout is in its inherent simplicity. Voltage and current signals are digitized in many different instruments and equipment in the accelerator, such as beam position and intensity monitors to name a few, and there is therefore great expertise in handling such signals at laboratories. This is further emphasised in [10] where the first IPM was using an anode strip detector and the design was aimed

towards simplicity and compatibility with the control system.

A common limitation for the optical and the electrical readout is the MCP itself, which has a limited lifetime. The position of the beam in an accelerator is often stable, which means the MCP will have a flux of input particles at the same location for most of the time. The channels in the MCP located under the centre of the beam will therefore degrade faster compared to the ones on the sides that see a lower flux of particles, which over time creates a non-uniform response of the MCP. Recent development aim to address this issue with an additional electric field that blocks particles from hitting the MCP and is activated when the IPM is not measuring the beam [33]. Another approach is to remove the MCP altogether, which is only possible if signal amplification is not needed, for example in [41] where an IPM is operating in a 1×10^{-2} mbar vacuum that leads to a high enough ionisation yield that a MCP is not needed.

IPM With Hybrid Pixel Detectors

In section 1.2.3 the feasibility of using hybrid pixel detectors in IPMs was mentioned and how it could fulfill the LIU requirements for an IPM. Another advantage of using pixel detectors in IPMs is that they enable detection of single ionisation electrons. For example, the Timepix3 readout ASIC has a minimum threshold of about 500 electrons, which means that it can detect a particle that deposits more than 1.8 keV of energy in a silicon sensors connected to the Timepix3. Since the ionisation electrons in an IPM have a kinetic energy of about 10 keV when they reach the detector, the Timepix3 can directly detect them and there is no need for a MCP.

The working principle is illustrated in fig. 2.7, where the hybrid pixel detector with a silicon sensor is located inside the vacuum under the beam. Since the analog to digital conversion takes place within each pixel of the detector, all the signals coming out of the detector are digital and therefore less susceptible to noise compared

to the strip readout for example. Later in section 3.1.3 the implementation of the detector assembly inside the vacuum used in this thesis project will be detailed. The first part of the readout electronics needs to be located outside the vacuum within a couple of metres of the instrument due to limited strength in the signals that connect between the readout and the pixel detector. The details of the readout that was designed for this project will be discussed in section 3.2.

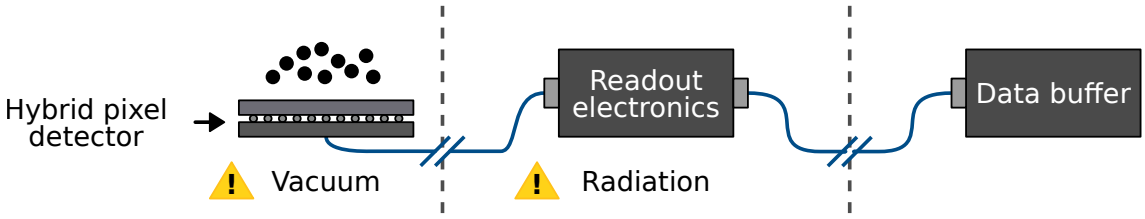

Figure 2.7: IPM with hybrid pixel detector.

One additional advantage of removing the MCP is that the height of the detector can be reduced, which in turn means that the overall height of the instrument vacuum chamber can be reduced. The reduction in aperture allows for a more compact magnet and potentially makes it cheaper as the good field volume is reduced, which relaxes the magnet specification.

The spatial resolution is determined by the pitch of the pixels in the detector and by the number of pixels, which for the Timepix3 detector is 256 x 256 pixels with a pitch of 55 µm. Additionally, a pixel based detector in an IPM leads to a smaller point spread compared to an MCP based detector due to the direct detection.

The biggest disadvantage with using hybrid pixel detectors in IPMs is the lack of experience due to its novelty in this context. This argument is true for both the design and operation of the detector which require specific knowledge of the pixel detector that is used. Prior to this work, we are not aware of pixel detectors being used inside the primary vacuum of an accelerator beam pipe. Pixel detectors are commonly used as part of experiments, outside the primary vacuum, in high energy particle physics, which usually operate independently from the accelerator physics

and instrumentation. This thesis project can therefore be seen as bridging the gap between the field of pixel detectors and the field of beam diagnostics.

2.4 Hybrid Pixel Detectors

This section will discuss the details of using hybrid pixel detectors to detect ionisation electrons in IPMs. The properties of silicon sensors will be discussed in section 2.4.1 and the Timepix3 readout ASIC will be detailed in section 2.4.2.

2.4.1 Silicon Sensor

The sensor in a hybrid pixel detector is where the interaction with the charged particles take place. It is bonded to the readout ASIC with microscopic solder balls. This separation of the sensor from the ASIC, compared to a monolithic pixel detector, allows for different materials and dimensions to be used with the same readout ASIC. Silicon is the most commonly used material and is well characterised and accessible due to its use in the semiconductor industry over the last 60 years. Other materials such as GaAs, CdTe or diamond have also been used in hybrid pixel detectors in specific applications [42, 43], but we will focus on silicon here.

Semiconductor

The working principle of a silicon sensor relies on the semiconductor properties of silicon as a material and is described in detail in [44]. In its intrinsic form, the concentration of free charge carriers (electron and holes) is determined by the silicon itself and the electron density is equal to the hole density. A semiconductor at a temperature above absolute zero have thermally generated free electrons and holes, where the concentration is equal between the two but increases with temperature. This property of a semiconductor is what makes it different from an insulator and a conductor. As shown in fig. 2.8, for an insulator there is a gap between the

conduction band at the top and the valence band at the bottom called the band gap. Electrical conductivity is the defined by excitation of electrons from the valence band to the conduction band and an insulator therefore has the least conductivity due to the wide band gap. An intrinsic semiconductor has a smaller band gap than an insulator and electrons can therefore move to the conduction band easier due to thermal or photon excitation. This means that a semiconductor is partially (semi) conducting, hence the name. On the other end, a conductor (metal) has the valence and conduction bands overlapping without any gap which allows electrons to flow between the bands.

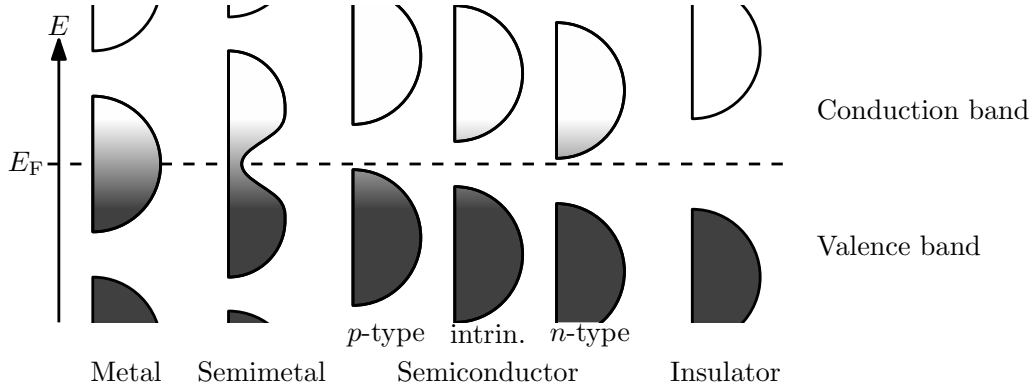

Figure 2.8: Band filling diagram of different material categories, where E_F is the Fermi level. Black indicates that all electronic states are filled and white that none are.

The properties of an intrinsic semiconductor can be modified by introducing impurities into the material, known as doping, which turns it into an extrinsic semiconductor. A doping material with one valence electron more than the intrinsic semiconductor can donate an electron that can move to the conduction band and is therefore called a donor material. The semiconductor doped with the donor is called n-type. On the other hand, if the dopant has one valence electron less it can instead easily trap an electron and produces a free hole. Such a dopant is called an acceptor and the doped semiconductor is called p-type. An example of dopant used with silicon is phosphorus for n-type and boron for p-type. The band-filling of a n-type and a p-type semiconductor is shown in fig. 2.8, where the n-type has the

conduction band closer to the Fermi level, making it easier for an electron to move to the conduction band, and the p-type has the valence band closer which allows it to accept an electron easier.

p-n junction

By combining a n-type semiconductor with a p-type we can create a p-n junction as shown in fig. 2.9. In the interface between the two regions is the so called p-n junction, where the free electrons in the n-type diffuse into the p-type and recombine with the holes. This creates a region at the junction that is depleted from free charge carriers and is therefore called the depletion region. The donor atoms within the depletion region on the n-type side are now positively charged due to the loss of their donor electrons and conversely, the atoms in the p-type side are negatively charged after accepting electrons. The charged side causes an electric field that cancel out the diffusion of free carriers in thermal equilibrium.

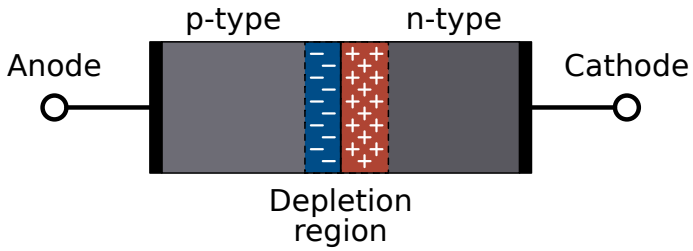

Figure 2.9: p-n junction formed in the interface between a n-type and p-type semiconductor. The blue region is negatively charged and the red region is positively charged. The grey regions are charge neutral.

By applying an external voltage across the p-n junction we can create an electric field in the same direction as the internal field from the junction which will further extend the depletion region. A positive voltage applied to the cathode side in fig. 2.9 relative to the anode side will create this condition and is called reverse biasing of the p-n junction. The width of the depletion region increases with the applied voltage up to a point where it breaks down and current begins to flow. We will not consider operating at the breakdown point in this thesis project. In [44] a detailed derivation

of the depletion region width is shown, which can be approximated as:

$$W \approx \sqrt{\frac{2\epsilon_0 \epsilon_{\mathrm{Si}}}{eN_{\mathrm{D}}} V}, \tag{2.27}$$

where ϵ_0 is the permittivity of free space, ϵ_{Si} is the relative permittivity of silicon, e is the elementary charge, N_{D} is the concentration of donor atoms and V is the externally applied voltage across the junction. A couple of examples are shown in table 2.3 to give an idea of the order of magnitude for the different parameters involved. Voltages below 100 V results in a depletion region width of hundreds of micron. The concentration of donor atoms (N_{D}) requires insight into the semiconductor manufacturing process that is used or it has to be approximated.

Table 2.3: Examples of depletion region widths.

	Example A	Example B
ϵ_0 [C m/V]	8.85×10^{-12}	
ϵ_{Si}	11.68	
N_{D} [m^{-3}]	1×10^{18}	
V [V]	30	90
W [µm]	200	340

Charged particle interaction with silicon

If a free charge (electron or hole) enters the depletion region of a reverse biased p-n-junction, the electric field and the lack of charge carriers in this region will accelerate the charge towards either side of the junction, depending on its polarity. This is the fundamental principle that is used by semiconductor pixel detectors where the charge is collected and quantized. The mean energy required to create an electron-hole pair in silicon is about 3.6 eV [13]. If we have a charged particle that deposits an energy of, for example, 10 keV in a pixel detector sensor, up to 2800 electron-hole pairs can be created in the silicon.

A sensor is constructed as a matrix of individual p-n-junctions, as was shown

earlier in fig. 1.5 (a), that defines the pixels. The junctions are reverse-biased by applying a potential to one side of the sensor, to create the electric field, while the other side is connected to the readout. This concept is illustrated in fig. 2.10 where two possible doping options are shown, n-type in p-bulk (called n-on-p) in (a) and p-type in n-bulk (called p-on-n) in (b). The figure shows a cross section with two pixels and at the top surface an ohmic contact is created with p+ or n+ implantation as well as a conductive metal layer where the bias voltage is applied. This metal layer could cause problems if the incoming charged particles are absorbed in it and therefore do not reach the depletion region in the silicon. A workaround for this is to selectively apply the metal layer to the top surface of the sensor such that it only covers a fraction. At the bottom of the sensor, the n+ or p+ contacts are connected to the readout ASIC through solder bumps, as was shown earlier in fig. 1.5.

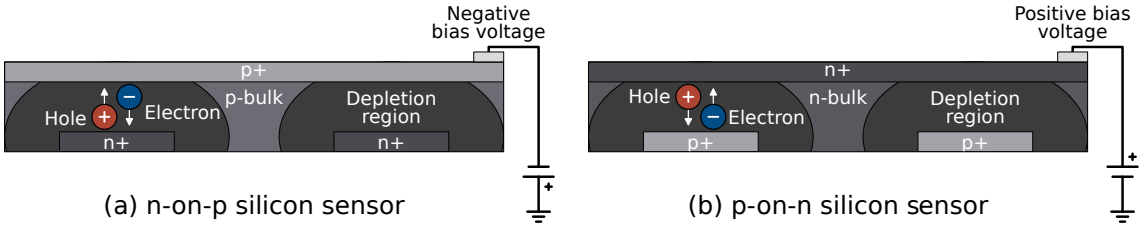

Figure 2.10: Cross section view of silicon sensors of n-on-p type (a) and p-on-n type (b).

Even with the metal layer removed, the charged particles we want to detect must have enough energy to reach the depletion region in the silicon. The continuous slowing down approximation (CSDA) range of, for example, a 10 keV electron in silicon is $3.46 \times 10^{-4}\,\mathrm{g/cm^2}$ [45]. By dividing the CSDA with the density of silicon ($2.33\,\mathrm{g/cm^3}$) we get an average path length of 1.48 µm for such electrons through silicon. This means that the depletion region must be closer than this to the surface of the sensor, where the charged particles enter the silicon, for electron-hole pairs to be created and collected.

If electron-hole pairs are created in the depletion region, the charge (either the electron or hole) can be collected in a pixel by the readout ASIC as a current signal.

On the other hand, if they are created in the undepleted region they will instead recombine and the charge cannot be detected. It is therefore essential that a bias voltage is applied that maximises the depletion region. The electric field from the bias also affects the time it takes for the charge carrier to be transported from where it was created to the bottom collection electrode (drift time). Additional factors affecting the time are the temperature, the charge carrier type and the transport distance. In [46], experimental results for the mobility of electrons and holes in silicon are discussed and summarised, with values of $\mu_n = 1430\,\mathrm{cm}^2/(\mathrm{V\,s})$ for electrons and $\mu_p = 480\,\mathrm{cm}^2/(\mathrm{V\,s})$ for holes. The important conclusion from these values is that electrons are three times more mobile than holes and therefore have a shorter drift time. In table 2.4, estimated drift times for two sensors with different thicknesses (d) and bias voltages (V_{bias}) are shown, under the assumption that the charge carrier have to drift through the whole sensor. This indicates that a thin n-on-p sensor collecting electrons, should be used to minimise the drift time. For the prototype instrument in this thesis project both n-on-p and p-on-n sensors have been used and in section 3.1.3 the different detector assemblies will be detailed.

Table 2.4: Examples of drift times for electrons (t_n) and holes (t_p) through a d thick silicon sensor with a bias voltage of V_{bias} applied across.

d	V_{bias}	t_n	t_p
100 µm	30 V	2.3 ns	6.9 ns
300 µm	90 V	7.0 ns	20.8 ns

2.4.2 Timepix3 Readout ASIC

The Timepix3 pixel detector readout ASIC, introduced earlier in section 1.3, had just been released at the start of the project and was identified in [12] as a potential candidate detector for the new IPM instrument in the PS. It consists of a pixel matrix of 256 x 256 pixels with a pitch of 55 µm and it can process events at a rate of 43 Mevents/s/(cm^2), which equates to 85 Mevents/s for the full detector. This

maximum processing rate can be sustained under the condition that all the events are distributed evenly across the detector. The 1.5625 ns time resolution is lower than the shortest bunch spacing in the PS accelerator (25 ns) and would therefore allow for bunch-by-bunch measurements. A photo of a Timepix3 ASIC with a silicon sensor mounted on top of it can be seen in fig. 2.11. On the left side of the chip in

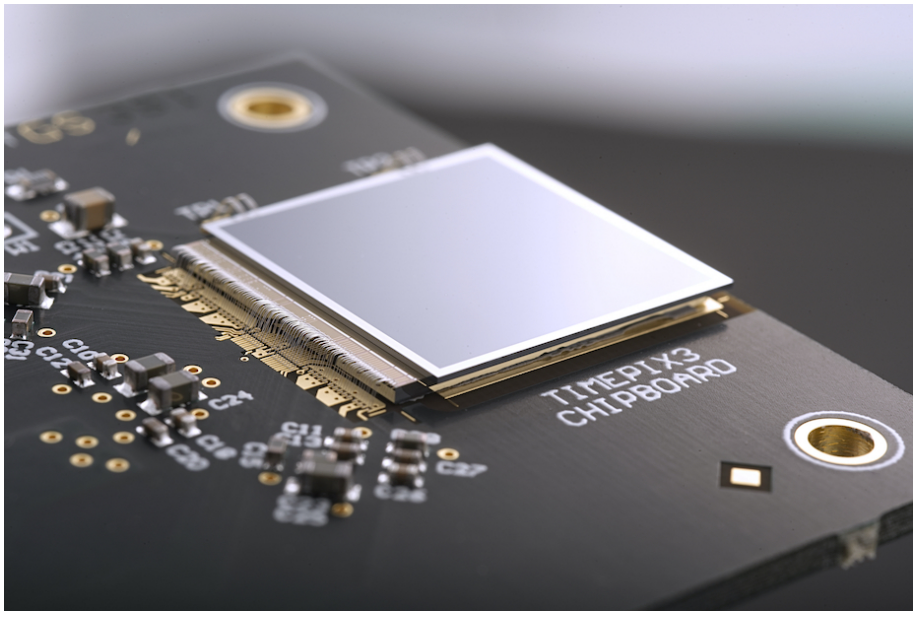

Figure 2.11: Photo of a Timepix3 hybrid pixel detector with a silicon sensor, from [47].

the photo, are the thin wire bonds that connect power and data signals to the chip. The visible metal border around the sensor is where the bias voltage is applied.

Data is read out from the Timepix3 through eight serial links that can be configured to each have a data rate from 40 Mbit/s to 640 Mbit/s, giving a total maximum data rate of 5.12 Gbit/s. Each event is composed of 60 bits. If we divide the maximum data rate by the number of bits per event we get 85 Mevents/s, which is how the maximum event processing is calculated. The Timepix3 also has several control input signals that are used to operate the detector. One of these is the shutter, which works similar to a camera shutter in the sense that it controls the time during which the detector is acquiring events. Additionally, there are control signals for reset and power management and data input where commands can be sent to

configure the Timepix3. The implementation of the Timepix3 readout and control will be detailed in section 3.2. The details of the analog front-end in each pixel of the Timepix3 will be detailed below as well as a description of the event timing and the information that is stored in each event.

Analog Front-end

The analog front-end in each pixel of the Timepix3 is designed to process the charge collected from the sensor (the design is detailed in [48]). In fig. 2.12 a schematic is shown for the front-end with the input pad that is connected to the sensor on the left and the trigger for the creation of a digital event on the right. There are several settings in the front-end that can be configured individually for each pixel in the matrix. If we follow the signal from the input pad on the left to the output on the right we first encounter a capacitor and a switch for the testbit. This is used to inject a test charge into the analog front-end that will pass through the same path as charge from the sensor and is useful for testing the pixel matrix without requiring a source of ionising particles. A similar switch is available in the digital block on the right in fig. 2.12, which bypasses all the analog circuitry and generates a digital event instead.

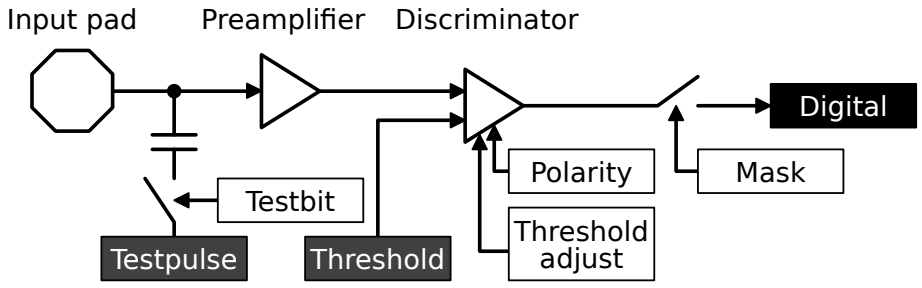

Figure 2.12: Schematic of the analog front-end for one pixel in the Timepix3 ASIC.

The charge, from the sensor or the test capacitor, then enters a charge sensitive preamplifier with low noise characteristics. Besides amplifying the signal it also takes care of the current to voltage conversion. The output of the amplifier will go high when collecting electrons and low when collecting holes from the sensor, which is

determined by the type of sensor that is used and how it is biased as shown earlier in section 2.4.1. The voltage signal is then fed to a discriminator where it is compared to a threshold that is common to all pixels in the matrix. Each discriminator will have its own unique response, due to temperature and manufacturing process variations and therefore requires a setting to adjust the threshold to equalise the response of all discriminators in the matrix. The equalisation method will be described in more detail later in section 3.3.2. There is also a setting in the discriminator that determines the polarity that should be detected, which configures the discriminator to trigger when the input signal goes above the threshold or when it goes below the threshold.

After the discriminator the signal is digital in the sense that it is a pulse with either a high or a low voltage. Before entering the digital part of the pixel circuitry a switch can be set to disconnect the output of the discriminator and therefore mask out the pixel. This provides a method to block a pixel from generating data and can be used on pixels that are broken and always causes the discriminator to trigger or when only a specific set of pixels should be enabled. An important detail is that the analog front-end in the pixel is still powered on when it has been masked out and will therefore respond to input signals even though it will not generate any digital events. We will come back to this inherent feature later when we discuss the equalisation method in section 3.3.2 where the total and peak power usage is important.

Event Timing and Stored Information

The timing of the analog and digital signals in the pixel is important as it defines what information is stored with each event and what it represents. Figure 2.13 shows an example event. In this case the amplifier output has a positive slope that indicates electrons are collected. If the polarity of the discriminator is configured correctly for electron collection, it will trigger as soon as the signal goes above the

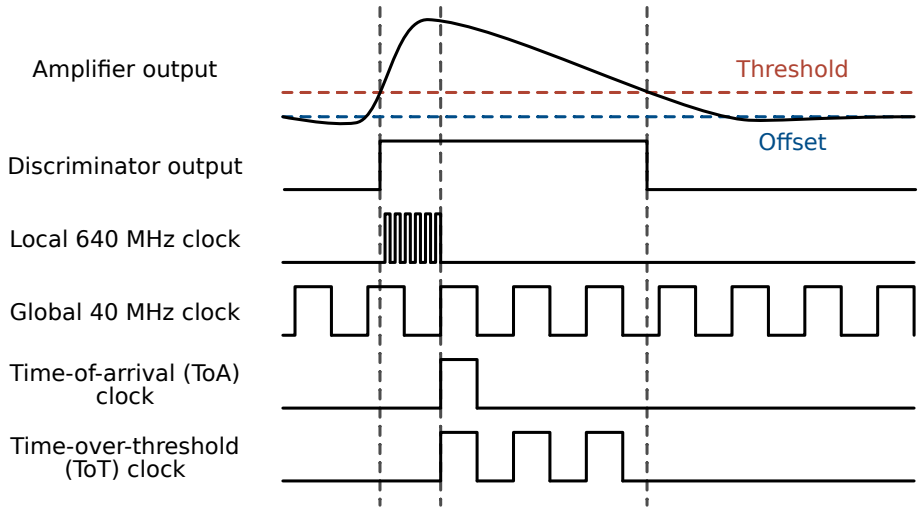

Figure 2.13: Timing diagram for the processing of an event in the Timepix3.

threshold. The amplifier output signal has an offset that is controlled globally for all pixels and defines the voltage level the output will be at when there is no charge collected. When collecting electrons, the threshold must be set above this offset and the discriminator must be configured to detect the rising slope of the signal. For hole collection, the threshold must be set below the offset and the discriminator must be configured to detect a falling slope of the signal. If the pixel is not configured in the correct way it may have an erratic response to stimuli. It is therefore important that the configuration is correct for the application it is used in, otherwise random events will be generated.

As seen in fig. 2.13, the output from the discriminator is digital and asynchronous to the clocks that are used in the Timepix3. When the output of the discriminator goes high, a local 640 MHz clock will start and increment a counter for each pulse until the first rising edge of the global 40 MHz clock. The value of this counter will be stored with the event as the fast time-of-arrival (fToA) information and has a resolution of 1.5625 ns. At the first rising edge of the 40 MHz clock, a pulse is generated that stores the value of a global time-of-arrival (ToA) counter with a 25 ns time resolution. The combination of fToA and ToA defines the time at which the event arrived.

The last information that is stored with the event indicates the time-over-threshold (ToT), which as the name suggest defines how long the signal was above the threshold. During the period when the discriminator is high, the ToT clock is incrementing a counter on each rising edge of the 40 MHz clock. When the discriminator goes low the current value of the ToT counter is stored with the event.

The discriminator signal will go low as soon as the amplifier signal goes below the threshold. When this happens, the location of the pixel, the fToA value, the ToA value and the ToT value are combined into a 48 bit packet (fig. 2.14), which is transferred from the pixel to the output of the Timepix3 through the data bus. By default, all packets are 8b/10b encoded and the output from the chip is therefore 60 bit words. The example shown here in figs. 2.13 and 2.14 is one out of six acquisition modes available in the Timepix3 and is the mode that has been used during this thesis project. There are also settings in the Timepix3 that can change the amplifier output shape for a faster or slower response, which in turn affects the power consumption, which will be detailed in section 3.3.3.

Figure 2.14: Packet with 48 bits created by the Timepix3 for each event.

An important characteristic of the Timepix3 is that to generate an event, the signal has to first go above and then below the threshold, as demonstrated in this example, and there is therefore a hard requirement that the input signal is not constant. This could be problematic if the repetition rate of the input signal is high enough that the output of the amplifier does not have enough time to discharge. There is also a limit in how fast a pixel can recover due to this discharging and the generation of a digital event, which depends on how the data is read out from the chip. In data driven mode, where events are sent out as soon as they are processed, a dead time of 475 ns plus the length of the pulse is expected after each event [18]. For example, an event with a ToT value of 30 has a signal pulse of 750 ns, assuming

a 40 MHz clock, which gives a total dead time of 1.225 µs. Additionally, a pixel with an event has to be read out before it can process another input signal and it has to be done before the 14-bit ToA counter overflows after 409.6 µs. The response of the Timepix3 pixel detector in the context of this thesis project will be discussed later in section 3.3.

3 | Implementation of a Beam Gas Ionisation Profile Monitor With Hybrid Pixel Detectors

This chapter will introduce the details of a beam gas ionisation profile monitor using the Timepix3 hybrid pixel detectors, which is called the Proton Synchrotron Beam Gas Ionisation profile monitor (PS-BGI).

We will cover the mechanical design of the instrument and the detector assembly in section 3.1, which will highlight the novel features of the design. After this, section 3.2 will present the electronics that are needed to transfer data from the Timepix3 pixel detectors, which are mounted inside the vacuum of the accelerator, to the computer located a couple of hundred metres away in an accessible location. This section will also detail the choice of components and the system architecture that enable the use of pixel detectors in a radioactive environment and in vacuum.

Section 3.3 will detail how the pixel detectors respond to stimuli and how they can be configured to detect ionisation electrons. Finally, in section 3.4 the hypothesis that ionisation electrons, which originate from the interaction between the accelerator beam particles and the residual gas, can be detected with the instrument will be tested.

3.1 Mechanical Design

The mechanical design of the PS-BGI profile monitor follows in many ways the traditional IPM designs described in section 2.3, with the entire instrument mounted on support arms that are attached to a vacuum flange. Such a design allows the instrument to be removed independently from its vacuum chamber (which makes up

a section of the beam pipe of the accelerator) for ease of maintenance. To create an Ultra High Vacuum (UHV) seal between the flange and the vacuum chamber, a rectangular ConFlat-type flange sealing system was successfully developed and tested at CERN for this specific instrument, which reduced the complexity of manufacturing and increased reliability [49]. As can be seen in fig. 3.1, there is a cathode and an anode to create an electric field perpendicular to the beam that is passing between the two electrodes, which accelerates the ionisation electrons down to a cutout in the anode where the detector is. A 0.2 T magnet is mounted outside the vacuum chamber and provides a magnetic field parallel to the electric field and therefore confines the electrons in a helical motion towards the detector, as discussed in section 2.3.

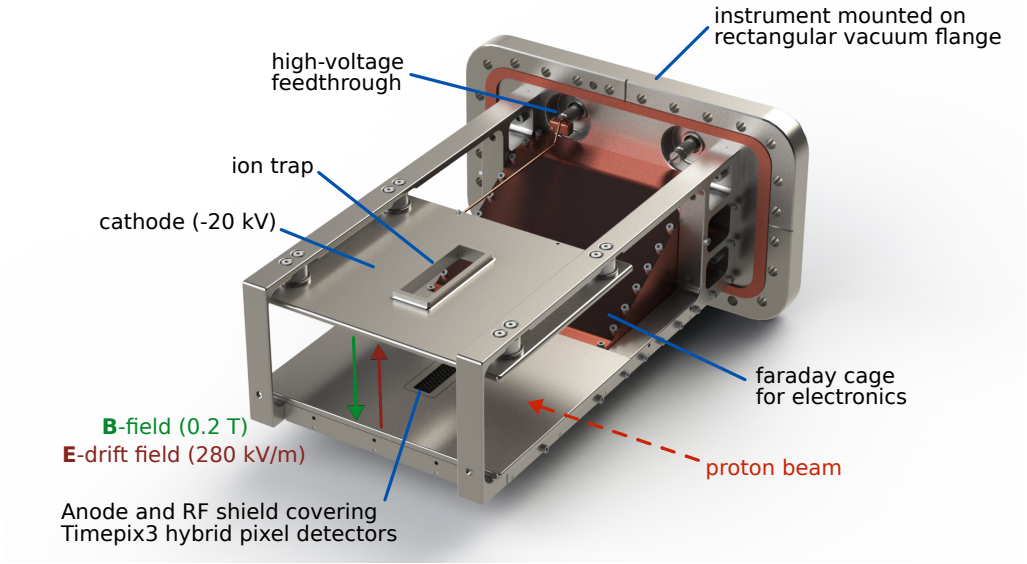

Figure 3.1: 3D rendering of the PS-BGI instrument.

The electric and magnetic fields of the horizontal PS-BGI instrument are shown in table 3.1 together with the previously shown IPMs from table 2.2. The PS-BGI electric and magnetic fields are among the strongest for the IPMs that are shown and they will be discussed in the following sections.

Table 3.1: Examples of electric and magnetic field strengths used in IPMs at different laboratories, including the PS-BGI.

Laboratory	Machine	Electric Field [kV/m]	Magnetic Field [mT]
CERN	**PS**	**280**	**200**
CERN	SPS [31]	100	60
Fermilab	Tevatron [32]	115	200
Fermilab	Main Injector [33]	10	100
GSI	ESR [34]	60	80
J-PARC	RCS [35]	86	46

3.1.1 Electric Field Cage Design

For the PS-BGI there are two aspects of the field cage design that differ from other IPMs and are presented in [50]. The first is the lack of field shaping side electrodes, which are normally used to improve the uniformity of the electric field and therefore reduce the distortion of the measured beam profile. A simulation of the electric field cage design was done in [51] to determine the effect without side electrodes which are commonly used in IPMs. The result can be seen in fig. 3.2 where the standard deviation of the accelerated electrons from a point source gets worse the further away from the center of the instrument. The detector in the PS-BGI instrument is $\pm 28\,\mathrm{mm}$ wide around the center, which according to fig. 3.2 results in a negligible difference between 0 and 3 side electrodes. Side electrodes were therefore not used in the implementation of the electric field cage.

The top cathode in the field cage is electrically isolated using ceramic spacers from the rest of the instrument (at ground potential) and it is connected to a high voltage power supply via a feedthrough connector in the flange. During operation, the cathode is kept at a potential of $-20\,\mathrm{kV}$ that creates an electric drift field between the cathode and anode with a strength of $280\,\mathrm{kV/m}$, which accelerates ionisation electrons down onto the Timepix3 detectors.

Secondly, to minimize the detection of secondary electrons, which are created due to the interaction between the ions from the ionisation process and the cathode,

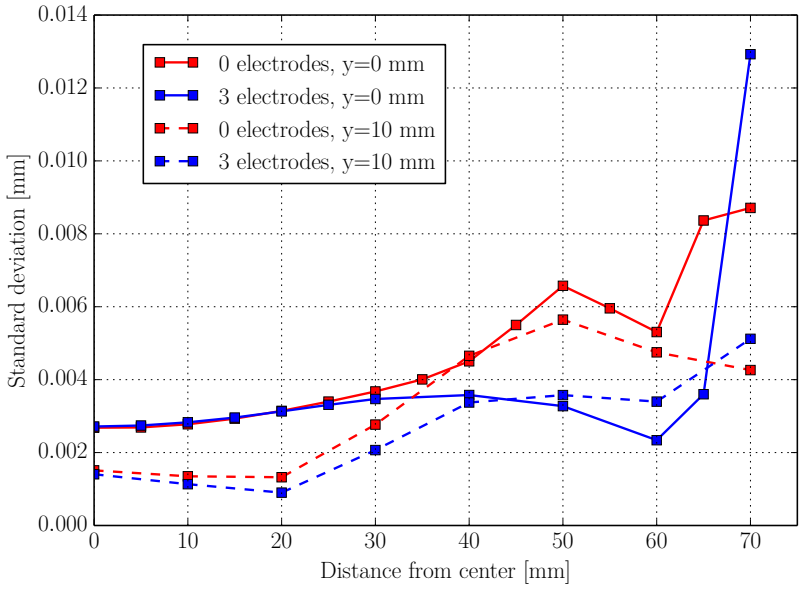

Figure 3.2: Simulation of the electric field with 0 and 3 side electrodes. Ionisation electrons are produced in a point source at the middle between the cathode and anode ($y = 0\,\text{mm}$) or 10 mm above. Figure from [51].

a novel idea illustrated in fig. 3.3 was implemented. The positively charged ions are accelerated towards the cathode that has a rectangular opening above the detectors that allows the ions to pass through. When they have passed through, their attraction to the cathode will cause them to bend in an arc that will move them away from the opening and instead hit the top side of the cathode, as shown by the red lines for the two ions in fig. 3.3. The electric field is three times stronger between

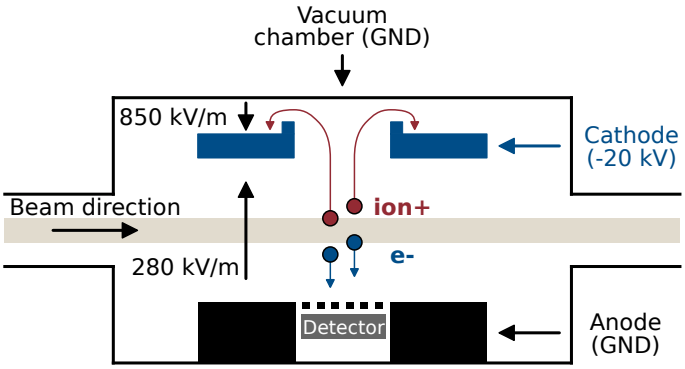

Figure 3.3: Illustration of the ion trap cutout in the cathode that traps the ions and minimises the creation of secondary electrons.

the top of the cathode (at $-20\,\mathrm{kV}$) and the grounded vacuum chamber compared to the region between the cathode and the anode where the ionisation occurs. The ions are trapped between the cathode and the vacuum chamber and any secondary electrons created here will not reach the detectors, which inspired the name "ion trap" [52].

3.1.2 Magnetic Field

IPMs that are detecting electrons are typically constructed with an external magnetic field that is parallel to the electric field, as shown earlier in fig. 2.5. For the PS-BGI a normal conducting triplet dipole magnet with a maximum magnetic field strength of $0.2\,\mathrm{T}$ is used [50], which is one order of magnitude less than the $1.5\,\mathrm{T}$ main bending magnets in the PS. It is important that this field is uniform in the detection region to avoid any distortions to the measured beam profile.

The horizontal PS-BGI magnet has been designed with a good field region of $x = \pm 25\,\mathrm{mm}$, $y = \pm 50\,\mathrm{mm}$ and $s = \pm 50\,\mathrm{mm}$ with respect to the central orbit of the beam, which almost covers the complete pixel detector area of $x = \pm 28\,\mathrm{mm}$ and $s = \pm 7\,\mathrm{mm}$. The detector is $\pm 3\,\mathrm{mm}$ outside the horizontal (x) good field region which is because the magnet was specified before the detector was chosen. Another important specification for the magnet is the integrated field quality which is defined as

$$\frac{\int B_y\, ds}{\int B_y(0,0,s)\, ds}. \tag{3.1}$$

For the horizontal PS-BGI magnet, the design specification required it to be less than $\pm 1 \times 10^{-3}$. The magnetic field of the produced magnet was measured at CERN and the integrated field had a maximum error of 0.63×10^{-3} which is below the specification [53]. In addition, the magnetic field at center of the instrument was measured to $0.198\,\mathrm{T}$.

The triplet configuration of the PS-BGI magnet compensates for the deflection

of the beam from the magnetic field, with one main dipole field at the centre of the instrument and two compensation dipole fields upstream and downstream, which have half the field strength in the opposite direction of the main field. This means that the magnet is self-canceling for the integrated magnetic field along the beam direction (s) and the design specification requires the integrated field to be as low as possible. Furthermore, the main and compensation dipoles are controlled from separate power supplies which allows tuning of the compensation during operation of the accelerator.

3.1.3 Detector Assembly

The ionisation electron detector assembly consists of four Timepix3 hybrid pixel detectors located under the anode plate in fig. 3.1. Each Timepix3 is individually bonded to 100 µm thick edgeless silicon sensors, each with 256 pixel columns and 256 pixel rows with a pixel pitch of 55 µm. The edgeless sensors minimise the gaps between the detectors and therefore improves the detection efficiency compared to using a standard sensor which has a guard ring around the periphery to reduce signal distortion for pixels near the edge [54]. Two detector assemblies have been produced, Mk I in 2017 [50] used p-on-n silicon sensors and Mk II in 2018 [55] used n-on-p silicon sensors. The main reason for switching to n-on-p was for availability from the manufacturer and a reduction in price due to n-on-p being more common.

RF-shield

The aperture above the detectors should have the same size as the detector area (14 mm by 56 mm) to maximise the detection efficiency of the assembly. This would allow ionisation electrons to reach the detectors without interacting with any other materials, but as indicated in fig. 3.1 this is not the case and a RF shield partially covers the detectors. The shield is there to protect the pixel detectors from electromagnetic interference (EMI) generated by the beam, which passes 3 cm above the

detector assembly. All the electronics inside the UHV is covered by a Faraday cage (highlighted in fig. 3.1), which is at the same potential as the anode. The shield above the detectors is also part of the Faraday cage and is made from 5 mm thick stainless steel with hexagonal apertures as illustrated in fig. 3.4. The walls of the shield have a width of about 150 µm, which is equivalent to 2 to 3 pixels. Due to the low energy of the ionisation electrons, they are blocked by the walls of the shield, which therefore leads to a reduction in detection efficiency. The shield is in some places deformed, which creates a non uniform acceptance over the detector area. It is therefore crucial to correct for this during the beam profile reconstruction to avoid systematic distortions of the measured beam profile. This procedure will be covered later in section 4.2.3.

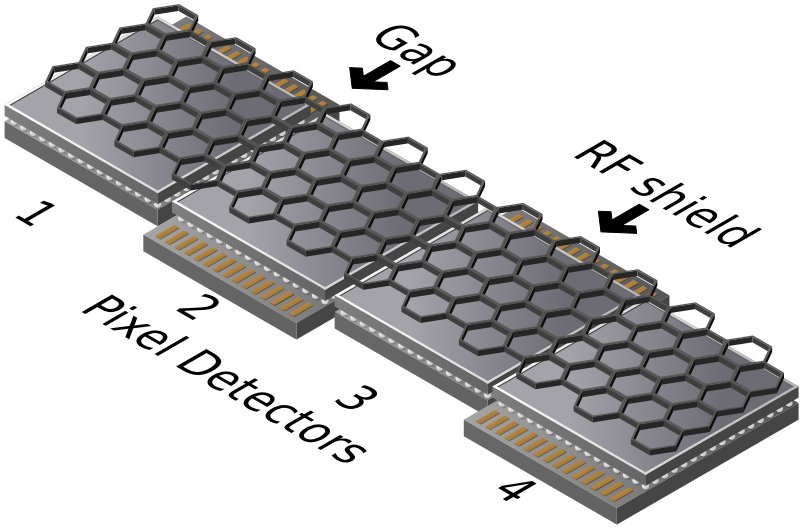

Figure 3.4: Illustration of the detector assembly with four Timepix3 pixel detectors and a RF shield above with hexagonal apertures.

Gaps Between Detectors

Another important aspect highlighted in fig. 3.4 is that there are small gaps between the four detectors, with a width in the range of 50 µm to 200 µm. This is due to variations in the sensor widths and limited precision when placing the detectors next to each other on the detector assembly. The gaps are measured with a microscope after

the assembly process and the result is used during the beam profile reconstruction in chapter 4 to correct for the true position of each detector.

Material and Component Selection

The choice of materials and components is critical in the detector assembly to reliably operate in the high radiation environment of up to 1000 Gy/year at the assigned location in the PS and to ensure UHV compatibility. In fig. 3.5 the Faraday cage and the electrode have been removed to show the detector assembly, with the four Timepix3 pixel detectors visible as the black rectangles in the centre. The detectors are mounted on an aluminium oxide substrate (ceramic), which has two metal layers to carry the power and data signals to flexible cables that are connected to the instrument vacuum flange [50]. Ceramic was chosen because of its UHV compatibility and radiation tolerance. In 2017, the first detector assembly (Mk I) was made using a low outgassing epoxy to attach the Timepix3 to the ceramic substrate. Before the assembly was installed in the accelerator, one of the detectors, the second from the left in fig. 3.5, did not respond to any commands or send out data and was therefore not operational. The remaining detectors worked without any problems and provided data during the year that was used to verify the principle of operating pixel detectors in a beam gas ionisation profile monitor inside the UHV of the accelerator [56].

The second assembly (Mk II) was prepared for installation in 2018, where the epoxy was replaced with an adhesive film, called Staystik 672. The film has been qualified for operation in UHV and radiation test shows it is an appropriate material to use for this application. After installation, all four detectors in the Mk II assembly were fully operational. Unfortunately, the first detector from the left in fig. 3.5 started to show erratic behaviour after a couple of months and stopped working shortly after. The location of this broken detector in Mk II was less critical compared to Mk I, as the three remaining detectors were next to each other.

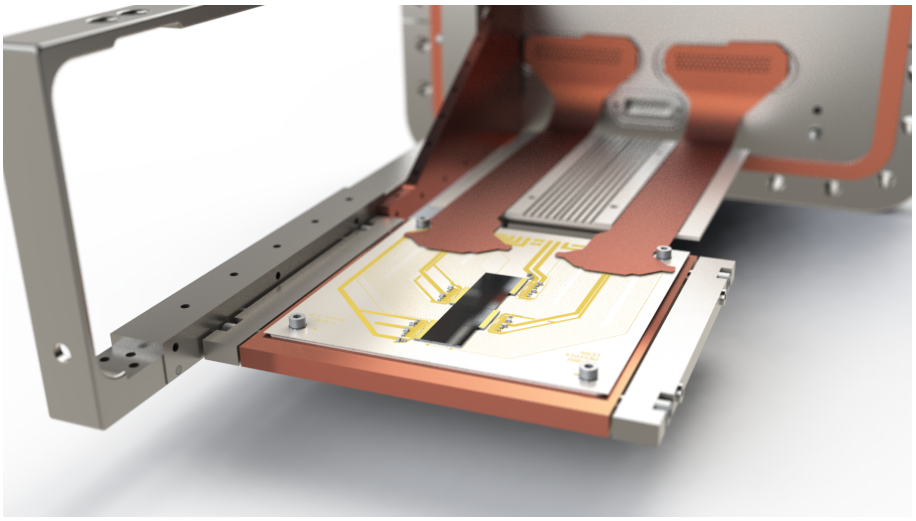

Figure 3.5: 3D rendering of the isolated detector assembly in the PS-BGI instrument.

Further analysis of the Mk II assembly has not been conducted due to resources being allocated to the development of a Mk III assembly. The common component in both Mk I and Mk II was the ceramic circuit board which was pushing the boundaries of the manufacturing process to the size of the board and the narrow and dense tracks for the differential signals. It is unlikely that the Timepix3 detectors were broken since they were tested and approved after the silicon sensor assembly by the manufacturer. Loose or broken connections between the detectors and the small connectors on the ceramic board is likely the cause of the failures. The new Mk III version aims to mitigate the yield issues of the large detector assembly, used in Mk I and II, by moving to a modular design with four independent Timepix3 modules that are attached to a mechanical support. More information about the Mk III detector assembly is not yet available as it is currently under development. In this thesis, only data recorded with the Mk II detector assembly in 2018 will be presented (in chapters 4 and 5).

Cooling

The detector assembly seen in fig. 3.5 is screwed down to a copper plate that is welded onto a pipe, where cold water at a temperature of around 6 °C circulates

during operation in the accelerator. The four detectors combined dissipate about 10 W of power that needs to be transferred by means of conduction from the pixel detectors to the cooled copper plate, via the adhesive film and the ceramic substrate. During operation of the Mk II detector assembly, it became apparent that the conduction between these layers was limited and the temperature of the detectors had to be monitored to avoid heating them up beyond the maximum operating temperature. The Timepix3 has an internal temperature sensor which was read out every second and displayed in the operational software. If any of the four detectors was approaching 90 °C during an acquisition they were all manually put into a low power state and the acquisition was canceled. There was no logging capability in the readout at that time and there are therefore no recorded temperature measurement. Various methods to limit the power consumption and to stabilise the temperature will be detailed later in section 3.3.3.

3.1.4 Orientation of the Detectors

The detector assembly is located under the bottom anode plate inside the PS-BGI instrument and it is protected by the stainless steel RF shield. The four Timepix3 detectors are mounted side by side and due to constraints in the routing of the signals on the ceramic carrier board, the first and the third detectors are rotated by 180°. A schematic view of the detectors is shown in fig. 3.6 with the shield above. The rows of the detectors are oriented parallel to the beam direction s and the columns are in the perpendicular transverse direction x. We therefore expect the ionisation electron rate to be constant for all rows, but to vary between columns with the beam profile density. Later in section 3.4 we will test this hypothesis.

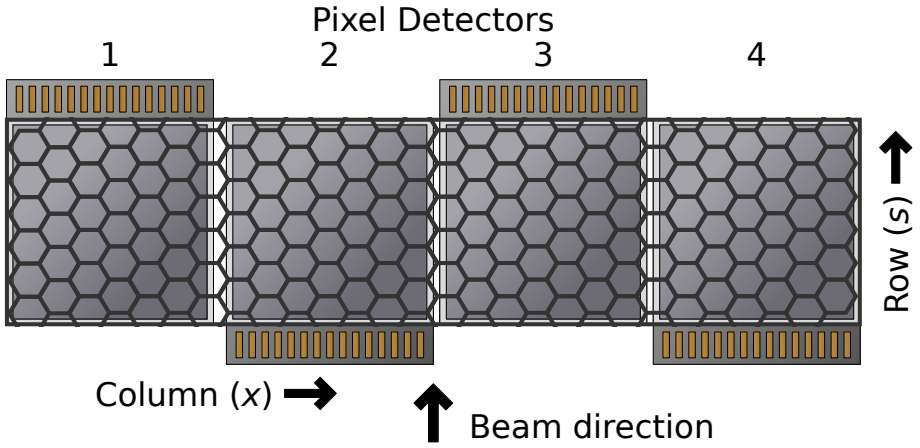

Figure 3.6: Orientation of the four Timepix3 pixel detectors with respect to the beam direction.

3.2 Readout Electronics and Software

The readout electronics is responsible for transferring data generated by the Timepix3 pixel detectors in the accelerator to a data acquisition (DAQ) computer. In the DAQ computer, the beam profile can then be reconstructed, a procedure which will be detailed in the next chapter (chapter 4). Physical hardware, field-programmable gate array (FPGA) code (called firmware), middleware and analysis software had to be designed and implemented from the ground up, due to the novelty of using pixel detectors in an IPM.

The PS-BGI hardware can be divided into three separate stages based on location and functionality, as shown in fig. 3.7. On the right in the figure is the "Detector" assembly, discussed earlier in section 3.1.3, and it is located inside the UHV of the accelerator. From the detector assembly, a set of flexible cables bring the electrical signals to vacuum feedthrough connectors that interface with the second stage, which is the "Readout Front-end" located just outside the vacuum chamber and will be discussed in section 3.2.1. The front-end has one FPGA that control the Timepix3 detectors and a second FPGA that forwards the data from the detectors to optical links that are connected to the third stage.

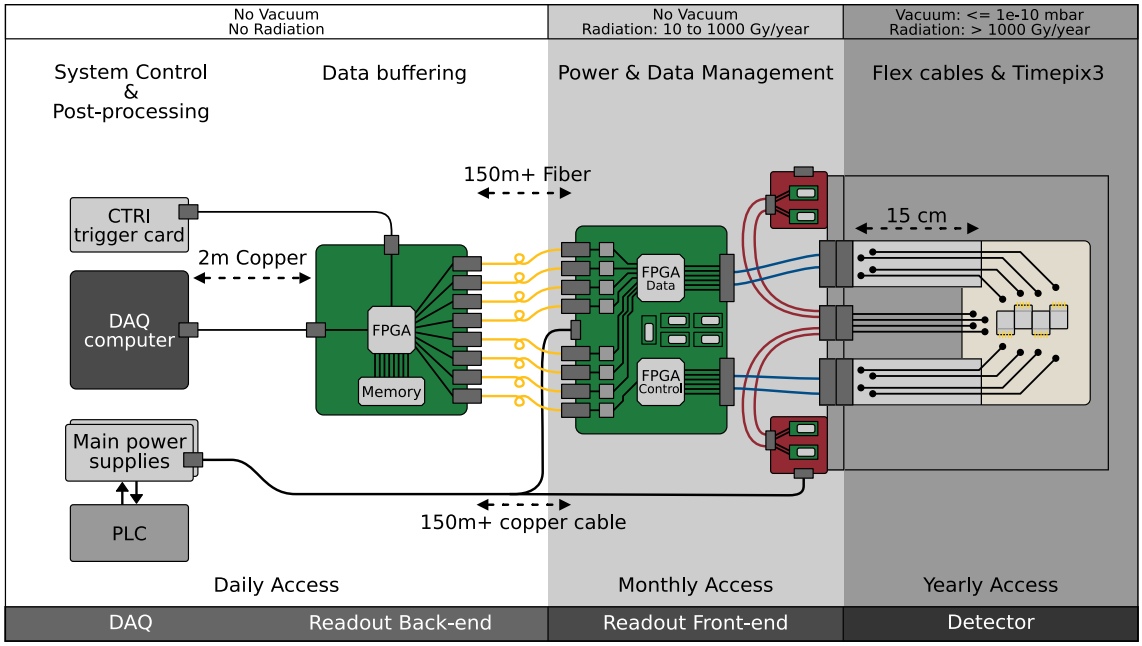

Figure 3.7: Overview of the readout electronics for the PS-BGI instrument which is composed of three stages with different environmental constraints and access.

The third stage contains the "Readout Back-end" and the "DAQ". An FPGA in this stage is connected to the optical links and buffers the data it receives from the front-end before it is read out by the DAQ computer over a 1G Ethernet connection. All the control of the instrument is done from the DAQ computer via a graphical user interface (GUI). To synchronise the measurements with the accelerator, a trigger card in the DAQ computer outputs signals to the back-end FPGA which are used to control the shutters in the Timepix3 detectors. The trigger card is connected to the central timing system of the PS and the outputs can be configured to trigger on specific beam cycles with adjustable delays.

Power for the front-end electronics is distributed on 150 m long copper cables from power supplies located in the back-end where there is no radiation. A programmable logic controller (PLC) manages the critical control of the power supplies and the cooling of the instrument and operates independently from the detector readout and control.

The different stages of the readout will be discussed in detail, starting with

the front-end hardware in section 3.2.1 and FPGA firmware in section 3.2.2. The back-end hardware will be discussed in section 3.2.3 and the FPGA firmware in section 3.2.4. Finally, in section 3.2.5 the GUI and the software running on the DAQ computer will be discussed.

3.2.1 Front-end Hardware

As shown in fig. 3.7, in the second stage of the readout (the "Front-end"), the hardware is located in air but it is still exposed to a high level of radiation, with an expected total ionising dose (TID) from 10 Gy/year to 1000 Gy/year. We therefore use components here that are radiation-hardened by design or that have been proven to be radiation-tolerant in similar environments. These components in the front-end will be detailed below.

Radiation and Magnetic Field Tolerant Power Supplies

Both the Timepix3 detectors and the front-end FPGAs require locally regulated power supplies for their operation, summarized in table 3.2. At the front-end of the PS-BGI instrument the TID is estimated up to 1000 Gy/year and another constraint for the power supplies is the instrument dipole magnet, which has a maximum field strength of 0.2 T. The power supplies must be able to withstand this radiation and magnetic field, which eliminates commercial-off-the-shelf (COTS) power supplies unless a batch of components are tested and qualified for operation.

Table 3.2: Voltage and power requirements for the front-end power supplies.

	One Timepix3 detector	FPGAs (control and data)
Voltage	1.5 V (analog), 1.5 V (digital)	1.0 V, 2x 1.5 V, 1.8 V, 2.5 V
Total power	2 W	20 W

At CERN, a DC/DC converter ASIC called FEAST has been developed and optimized for radiation tolerance and high efficiency [57]. The ASIC is assembled

together with an air core inductor, capacitors, biasing circuitry and a RF-shield into a module called FEASTMP [58], which has the specifications summarised in table 3.3. The FEASTMP module fulfills the radiation and magnetic field require-

Table 3.3: FEASTMP converter module specifications.

	Min	Max
Input voltage	5 V	12 V
Output voltage	0.9 V	5 V
Output power	-	10 W
TID	-	>2 MGy
Magnetic field	-	>4 T

ments at the front-end and it can also provide the necessary output voltages for the Timepix3 detectors and front-end FPGAs.

Each of the four Timepix3 detectors should have separate power supplies for the analog domain and the digital domain to achieve the lowest noise and to minimize coupling of digital noise to the analog circuitry inside the detector. The materials used in the FEASTMP module are unfortunately not vacuum compatible, especially the use of thermal paste inside the module, and can therefore not be mounted inside the vacuum of the instrument close to the Timepix3 detectors, which would be the ideal location. Instead, the modules are mounted outside on the instrument flange as can seen in fig. 3.8 and the power is brought to the detectors via a vacuum feedthrough connector. The detectors share the power supplies in groups of two, but keeping the analog and digital supplies separated, so that a total of four FEASTMP modules are mounted on the flange instead of eight. This is done to increase the reliability of the system as a reduction in the number of power supplies will increase the mean time between failures (MTBF), which is important as the access to the front-end is limited and dangerous due to radiation during operational periods.

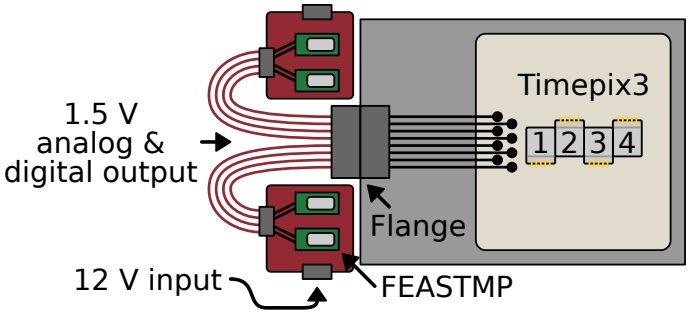

Figure 3.8: Connection of FEASTMP power supplies to the Timepix3 detectors.

Front-end FPGA Board

The front-end FPGA board is responsible for reading the data that is sent by the Timepix3 detectors and to send commands to the detectors to configure them and to start and stop a measurement. No data processing is done on this board, instead it acts as a router where the Timepix3 data is packaged and sent to the optical links that connect to the back-end. This real-time streaming architecture therefore requires a bandwidth that is greater than the combined maximum bandwidth of four Timepix3 detectors, which is 20.48 Gbit/s.

In the first version of the front-end, a generic readout board developed at CERN called GEFE was used [59]. This board was designed to be radiation tolerant and contained a flash based FPGA from Microsemi (part number A3PE3000-FGG896) that was tested in radiation up to a TID of 750 Gy. It also had two FEASTMP modules for power. For data transmission, an optical link based on a GBTx chip with a VTRx optical transceiver module was used, which are radiation hardened components from the GigaBit Transceiver (GBT) project [60]. The available bandwidth of the optical link was 3.2 Gbit/s in both directions, which was not enough to run the four Timepix3 detectors at their full data bandwidth and they were therefore limited to 640 Mbit/s per detector. At the start of the project, this was the only available hardware that met the radiation requirement and it was therefore used for the operation of the Mk I instrument in 2017. The limited bandwidth of this readout caused issues in the beam measurements, due to data piling up in the

detectors because it was not read out fast enough by the FPGA on the GEFE board and the limited optical link bandwidth.

The development of a new readout board, version 2, was therefore initiated. A lot of time had already been spent developing and debugging the firmware on the flash based FPGA on the GEFE board and the experience from operating this board showed that it was a suitable FPGA for the front-end. The new board was therefore split into two parts, with one FPGA handling the control of the Timepix3 detectors, reusing the firmware from the GEFE board, and a second FPGA taking care of the data from the detectors. For the control part, a smaller version of the Microsemi ProASIC3 FPGA (part number A3PE1500-FG484) was chosen. This FPGA was not sent for radiation tests due to a tight schedule and limited resources in the project. For the data FPGA, a Kintex-7 from Xilinx (part number XC7K70T-1FBG676C) was used. Part of the decision for the Kintex-7 was the adjustable delays of the input and output pins that can be used to correct for any length mismatch between the 32 data links from the Timepix3 detectors and the common clock. The ProASIC3 also has adjustable delays but they are static and can only be changed with a full reconfiguration of the FPGA. The Kintex-7 enabled real-time adjustment of the delays and could therefore be used to ensure that the Timepix3 data links are stable when they are captured by the clock in the FPGA.

There are few published radiation tests of the Kintex-7 targeting the TID effects. One experiment, where they tried to find the point at which the device failed, indicated that the devices were still functional after 3000 Gy [61]. The Kintex-7 is a SRAM based FPGA and is therefore more susceptible to single-event upsets (SEUs) compared to a flash based FPGA, due to the large configuration memory that could alter the behaviour of the devices if an upset occurs there. Mitigation techniques can be deployed in the user logic of the FPGA, for example triple modular redundancy (TMR), but before that undertaking it is important to first evaluate how critical a failure is to the system as a whole. The PS-BGI instrument will be used to

monitor the beam profile on demand by the operators and it is not connected to a beam interlock and should therefore not be considered a critical equipment for the operation of the accelerator. If an upset occurs in the Kintex-7 that requires the system to restart, it could be done within a few seconds. On the other hand, if the rate of such upsets is so high that it interferes with the operation of the instrument, mitigation techniques must be deployed. In [62], a case study is presented for the use of Kintex-7 in space missions, where upsets are categorised based on their severity and observability and practical solutions are suggested in each case. A similar analysis could be carried out for the PS-BGI instrument in case the upset rate is interfering with the operation of the instrument, but during the operation of the instrument with the new readout in 2018 from September to December, no upsets were observed.

An illustration of the version 2 of the front-end FPGA board is shown in fig. 3.9. In addition to the two FPGAs there are five FEASTMP power supply modules (voltages shown in table 3.2) and a total of eight GBTx chips connected to VTRx optical modules. A total of 22.4 Gbit/s optical bandwidth for transmission is connected to the data FPGA, which enables the Timepix3 detectors to operate at the maximum bandwidth of 20.48 Gbit/s. The control FPGA does not need to send or receive large amount of data and is therefore only assigned 1.28 Gbit/s for transmitting and 1.92 Gbit/s for receiving. Commands and settings for the modules inside the data FPGA, which will be discussed in section 3.2.2, can be received at a rate of 2.56 Gbit/s. A common 40 MHz clock is used for the whole board, which is also the clock that is sent to the Timepix3 detectors, and it is originating from the bottom GBTx chip as seen in fig. 3.9. The board can also be configured to use an external clock.

The front-end FPGA board is assembled and mounted inside a metal crate with all the connections facing one side as can be seen in fig. 3.10. RJ-45 connectors are used to connect to the Timepix3 detectors, using six shielded CAT6A Ethernet

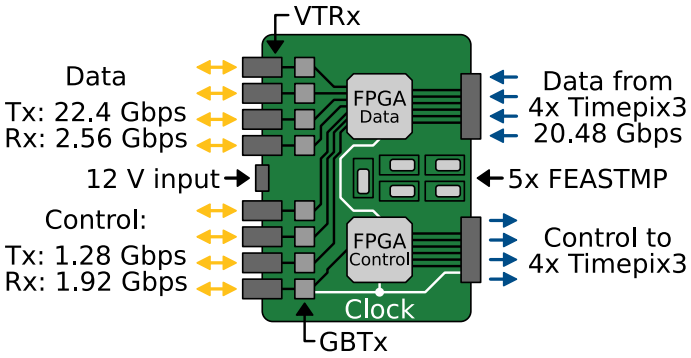

Figure 3.9: Illustration of the front-end FPGA board with FEASTMP modules for power and GBTx + VTRx for optical data transmission.

cables per detector, over a distance of 2 m. Below the Timepix3 connectors are the VTRx modules that are using standard optical LC connectors, shown with the black protective covers mounted in the figure. A total of 16 fibres could be connected, but only 10 are needed due to the asymmetric bandwidth shown above. Additionally, there are four SMA connectors that connect to an analog-to-digital converter (ADC) on the board and are used to sample an analog output signal from each Timepix3 detector. These analog output signals are configurable and are used for temperature measurements and monitoring of the voltages inside the detectors, which we will see later in section 3.3.1.

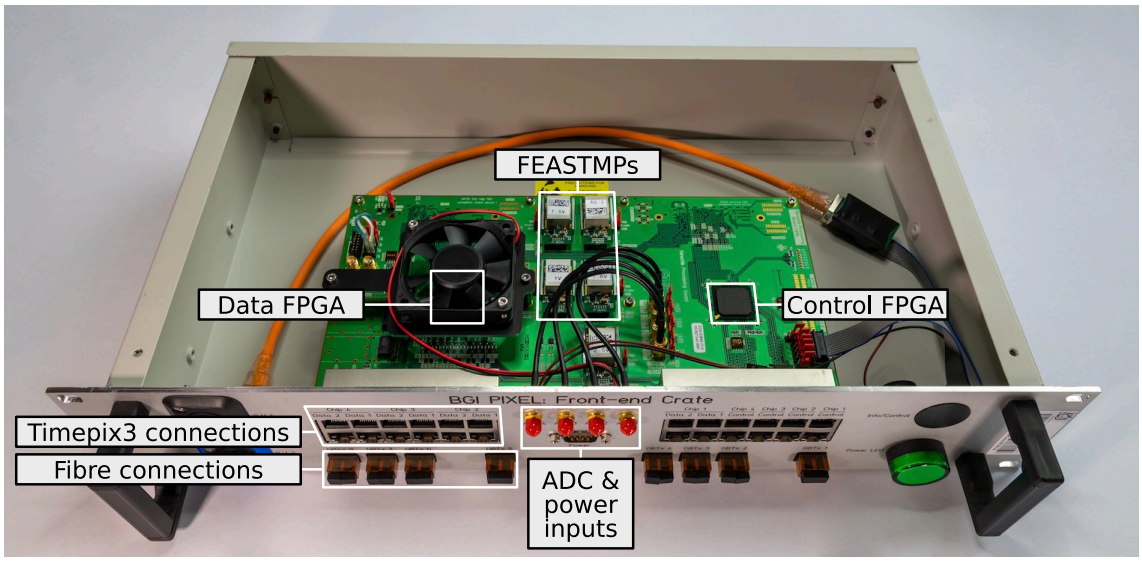

Figure 3.10: Photo of the assembled front-end FPGA crate.

3.2.2 Front-end FPGA Configuration and Firmware

The front-end firmware controls how the two FPGAs are configured and the procedure to load this firmware differs for each due to the different technologies that are used. For the flash based ProASIC3 control FPGA the configuration is stored in non-volatile flash memory cells on the device itself. It is therefore retained after power is removed and as soon as power is applied again and the internal reset has been released, the FPGA is configured and operational. A dedicated programmer has to be connected to the control FPGA to load the memory cells with the configuration and therefore requires physical access to the front-end FPGA board if the configuration has to be changed.

The SRAM based Kintex-7 device stores the configuration in volatile SRAM cells that loose their state when power is removed, which means that this FPGA has to be configured from an external device every time it is powered on. A common method is to connect a flash memory chip, where the configuration is stored, to the Kintex-7, which allows a module inside the FPGA to pull the data from the flash memory and configure all the SRAM cells. This is not a viable option due to the radiation requirements of the readout. Instead, a chip on the front-end FPGA board called GBT-SCA that communicates over one of the GBT links takes care of the configuration [63]. A JTAG master on the GBT-SCA is connected to the JTAG port on the Kintex-7. This enables the configuration to be sent from the back-end over the optical link to the FPGA via the GBT-SCA. It is therefore possible to remotely change the firmware of the data FPGA and no access to the front-end board in the accelerator is needed for this device.

Control FPGA Firmware

The majority of the control FPGA firmware was developed for the first GEFE based version of the readout. This was done as a master thesis project where the focus was

on exploring mitigation techniques for use of FPGAs in radiation and how to detect single event upsets [64]. During the summer of 2016 a GEFE board was installed in the PS accelerator with a firmware that monitored single event upsets. Over the course of 99 days with the board continuously running, 11 faults were detected, but only one of those required the system to be restarted. The firmware employed triple modular redundancy (TMR) techniques on all registers, which means that each register had three identical copies and the outputs passed through a majority voter that made a decision on the combined output. Such a triplication ensures that single bit errors can be detected and corrected, but it also increases the number of required registers by a factor of three. The readout firmware on the first version with the GEFE board implemented this TMR technique on all registers and since the firmware was reused for the second version, the control FPGA firmware has this implementation also.

A block diagram of the control FPGA firmware is shown in fig. 3.11 and it consists of three main parts. The GBTx controller takes care of the communication with the GBTx chip by decoding the received data and packaging data that will be sent to the back-end. The received data consists of data packets that will be forwarded to the specified Timepix3 controller and settings for internal registers in the firmware. A bus connects all the modules in the firmware and the GBTx controller acts as the master on the bus while the Timepix3 controllers are slaves. The four Timepix3 controllers handle the control signals that are connected to the detectors and as can be seen in fig. 3.11, there are no signals going from the detectors to the controllers. This is due to how the Timepix3 was designed, where both control and data packets from the detector are sent through one of the eight data links that are connected to the data FPGA. The Timepix3 controllers also take care of the reset and power enable signals, which are used to put the detectors in different operational modes, and a time synchronisation signal that resets the internal counters of the detectors and defines the timestamp of each detected event. The shutter, which

controls the data acquisition, is routed directly from the GBTx controller to all Timepix3 controllers that ensures time synchronisation between the detectors when the shutter is opened. These signals can be controlled individually per detector or simultaneously for all four to synchronise them.

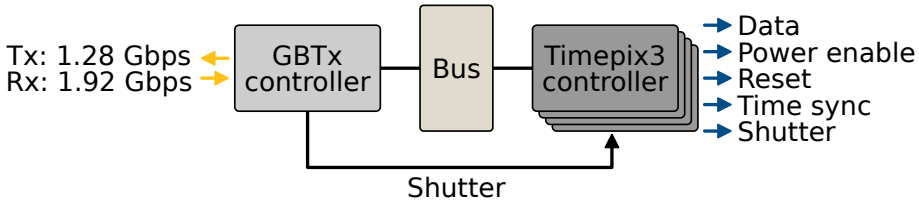

Figure 3.11: A block diagram of the firmware for the control FPGA with a GBTx controller connected to four Timepix3 controllers through a common bus.

Data FPGA Firmware

The data FPGA firmware is also built around a common bus with controllers to interface with the seven GBTx chips and four Timepix3 controllers to synchronise and capture data coming from the detectors. A block diagram of the firmware can be seen in fig. 3.12. One of the GBTx controllers receives 2.56 Gbit/s of data from the back-end and acts as the master on the bus. A dedicated data path connects the four Timepix3 data synchroniser modules to the GBTx controllers and the 20.48 Gbit/s of Timepix3 data is distributed over the 22.4 Gbit/s optical transmission bandwidth. This ensures that data is kept in a continuous stream through the system without any congestion.

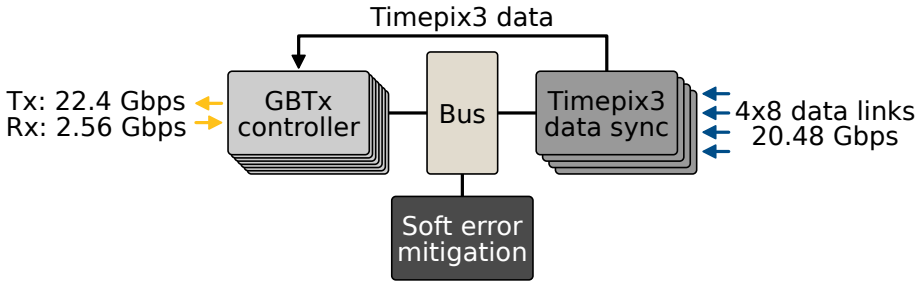

Figure 3.12: A block diagram of the firmware for the data FPGA with GBTx controllers, Timepix3 data synchronisers and a soft error mitigation module.

The Timepix3 data synchronisers contain adjustable delay blocks, which was one of the main reason to use the Kintex-7 as the data FPGA as discussed earlier. Each Timepix3 detector has eight data links connected to the FPGA that can be individually delayed to ensure that the data is stable when it is captured by the clock. The delay can be set manually or automatically through a syncronisation mechanism. When the Timepix3 is idle it is transmitting a unique comma symbol called K.28.5 on the data links, which is a common comma symbol used for the 8b/10b encoding that the Timepix3 is designed with. The synchronisers search for this comma symbol in the automatic mode and will scan through the different delay settings until they find it. When the comma is found, it will lock the delay value and start transferring the captured bits to the GBTx controllers. The delay value that is needed depends on the path through which the Timepix3 data has to propagate from inside the instrument all the way to the front-end FPGA board. This is greater than 2 m, but more importantly each data link can have a different path length compared to the rest. The flex cables inside the instrument and the flange connection board have been designed to minimise any length mismatch, but in the interfacing connectors and in the Ethernet cables there will inevitably be a mismatch.

Another important feature in the front-end data FPGA is the extended timestamp counter. The time-of-arrival (ToA) counter for a Timepix3 event is 14 bits long and will therefore overflow after 409.6 µs. At the start of an acquisition, the time synchronisation signal from the control FPGA clears the ToA counter and at the same time a counter in the data FPGA is also cleared. This counter is 42 bits and it is synchronous with the ToA in the detector. The value of this extended counter is sent over the GBT links to the back-end where it is combined with the ToA count for each event, which increases the overflow time to about 30 h. We will discuss more how the back-end handles the data sent from the front-end in section 3.2.4.

The last module in fig. 3.12 is the soft error mitigation (SEM) and its purpose is

to mitigate and detect SEUs in the configuration memory caused by radiation. This is based on an intellectual property (IP) core designed into the Kintex-7 device that can be configured in different modes. The configuration memory in the Kintex-7 is divided up into frames where each frame consists of 101 32-bit words. For this application, the core has been configured to operate in an automatic repair mode that is continuously scanning the configuration memory and can correct single-bit errors within one frame of the memory. The module will also report when errors are detected, which can be used to monitor the SEU rate during operation. During the four months at the end of 2018, when the system was installed in the accelerator, the SEU rate was not monitored, but it is foreseen to be monitored in the operational system at the end of 2020.

The front-end firmware was named WaterBear after the micro-animal (formally called Tardigrade) that can withstand doses of radiation 1000 times higher than other animals.

3.2.3 Back-end Hardware

In the last stage of the readout system, the back-end, there are no vacuum or radiation requirements and COTS components can therefore be used. A block diagram of the hardware components is shown in fig. 3.13. The front-end FPGA board is

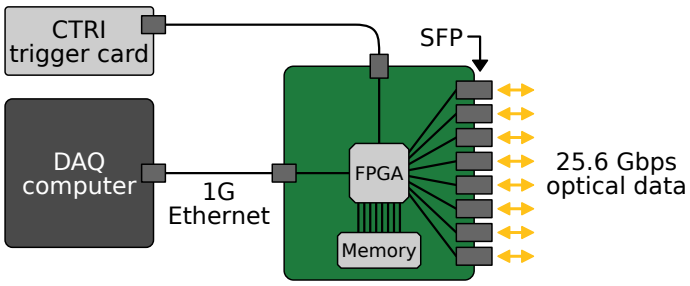

Figure 3.13: Illustration of the back-end hardware with the FPGA board, the computer and a trigger card.

connected over 150 m long optical fibres to the back-end FPGA board that will buffer the data in memory before it is read out over a 1G Ethernet connection by

the computer. There are two trigger signals connected to the back-end FPGA board to synchronize the measurements with the accelerator. One trigger signal indicates when a beam is injected into the accelerator, which is used to start a measurement, and another trigger signal indicates when the beam is extracted, which will stop a measurement.

Back-end FPGA Board

The back-end FPGA board is a commercially available product called VC707 from Xilinx and it contains a Virtex-7 FPGA. The board also has a 1 GB DDR3 memory and two high speed connectors for FPGA Mezzanine Cards (FMCs). In one of the FMC connectors a mezzanine is installed where 10 small form-factor pluggable (SFP) modules are mounted and connected to the optical fibres that go to the front-end. All the components are mounted inside a rack mountable crate as can be seen in fig. 3.14. The 1G Ethernet and the fibres are connected at the front of the crate as well as four trigger inputs using BNC connectors.

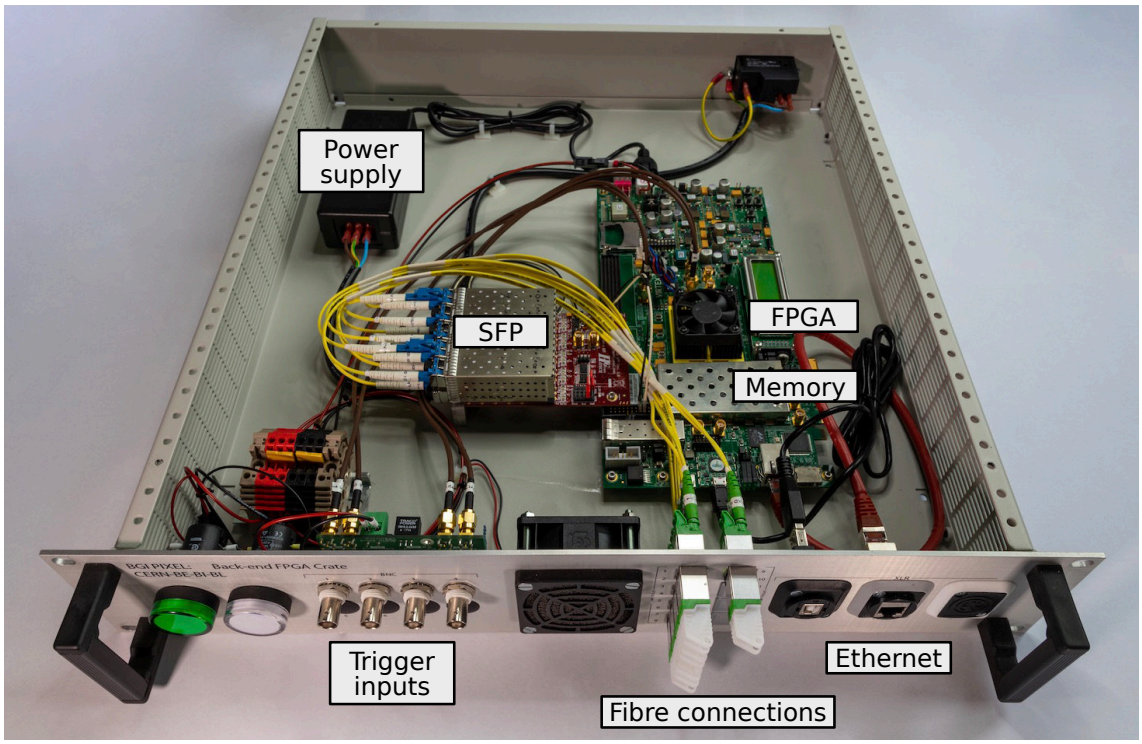

Figure 3.14: Photo of the back-end FPGA crate.

Environmental Control and Monitoring

As was shown in the system overview in fig. 3.7, the main power supplies are located in the back-end and are connected using 150 m long copper cables to the front-end. The lack of radiation in the back-end enables the use of COTS power supplies that are controlled with a programmable logic controller (PLC) independently from the rest of the instrument. In the PLC, a software is running that continuously monitors the states of the power supplies and processes request to change the states from the network over a Ethernet connection. An illustration of the PLC and the four power supplies that are needed for the instrument is shown in fig. 3.15. The power supply for the magnet is part of a separate system with its own dedicated controller.

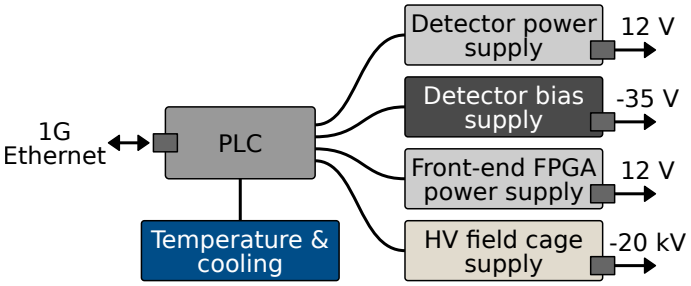

Figure 3.15: Illustration of the back-end PLC that controls and monitors the power supplies, temperature and cooling.

For the detectors there is one 12 V supply that is connected to the FEASTMP modules on the flange, shown earlier in section 3.2.1, and one supply for the sensor bias. The front-end FPGA board is powered from a separate 12 V supply. The PLC also controls the high voltage power supply that is connected to the field cage inside the instrument and it can supply a voltage down to −20 kV.

At the instrument, inside the accelerator, there are cooling pipes that connects to the detector assembly, discussed earlier in section 3.1.3, and the PLC can control the flow of water in these pipes with a solenoid. The temperature of the cooling water is also measured by the PLC and a switch indicates if there is water flowing or not. Several safety procedures are employed in the PLC software that make sure

the instrument is operating in a safe state. For example, if the cooling is disabled or the flow suddenly stops, the PLC will turn off the detector and bias supplies to prevent damage.

3.2.4 Back-end FPGA Firmware

The main task of the back-end FPGA firmware is to assemble Timepix3 events from the data it receives from the front-end and then to buffer these in a memory before they are read out through the 1G Ethernet connection by the computer. A block diagram of the firmware, which was named PolarBear, is shown in fig. 3.16. On the right in the figure, the fibres connect to eight GBT-FPGA cores that will decode the data sent by the GBTx chips in the front-end. This core was developed by the GBT project for this specific purpose of interfacing FPGAs with the GBTx [65]. Each channel in the core handle 80 bits of data at a frequency of 40 MHz in both direction, which is the 3.2 Gbit/s of available bandwidth on a GBT link. The back-end can therefore process eight GBT links at the full bandwidth, and the bits that are not used by the front-end will be ignored.

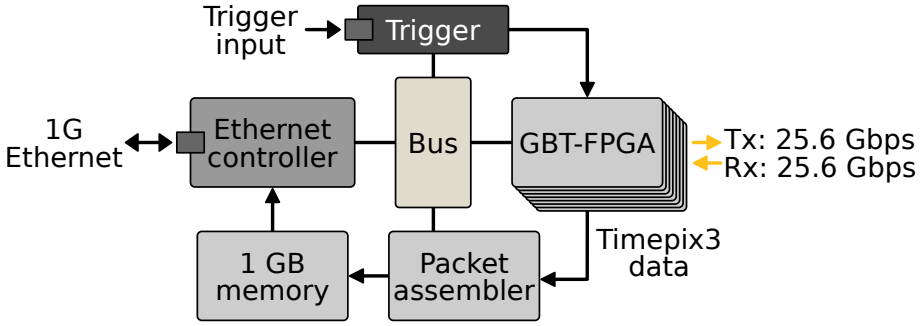

Figure 3.16: A block diagram of the firmware for the back-end FPGA with GBT-FPGA blocks, packet assembler, memory, Ethernet controller and trigger.

After the GBT-FPGA cores have processed the data coming from the front-end, it will be sent to a packet assembler where the encoded 10 bit chunks of Timepix3 data are first decoded to 8 bit (one byte) and then assembled into 48 bit Timepix3 packets. Additionally, the packet will be combined with a detector identifier, a

trigger count and the extended counter from the front-end to form a 84 bit event word. The ordering of the bits in the event word can be seen in fig. 3.17. The 28 most significant bits from the counter are used as the timestamp, which creates a 42 bit total timestamp together with the 14 bit ToA information in the Timepix3 packet.

Figure 3.17: 84 bit event word constructed in the packet assembler with detector identifier, trigger count, timestamp and Timepix3 packet.

After the word has been constructed in the packet assembler it will be sent to the 1 GB memory, which is the external DDR3 memory that is mounted on the back-end FPGA board shown in fig. 3.14. Data has to be transferred to the memory as 128 bit words and the event word is therefore padded with zeros. This means that a total of 67 million events can be stored in the memory. Four Timepix3 detectors can process around 320 Mevents/s (at maximum and under the assumption the events are evenly distributed as discussed in section 2.4.2). This means that in this worst case scenario, the memory will be full after 200 ms. During measurements with the PS-BGI instrument the pixel detectors will not experience such high event rates. The expected event rates will be discussed later in section 3.5.

The memory has a dedicated connection to the Ethernet controller, as shown on the left in fig. 3.16. The controller also acts as the master on the bus and enables a computer, connected over the 1G Ethernet, to communicate with the modules in the firmware and to read data from the memory. Inside the Ethernet controller, a core called IPbus has been implemented that takes care of the formatting of the Ethernet packets and provides a reliable and high performance link with accompanying software running on the computer [66]. During testing, a throughput of over 500 Mbit/s was observed when reading events from the memory on the FPGA board to the computer. This is still lower than the 20.48 Gbit/s the Timepix3 detectors can output and it is therefore the bottleneck for a completely stream based readout

system. The memory buffer on the board alleviates this limitation.

For the operational instrument, which will be installed at the end of 2020, local processing of the Timepix3 events will be implemented in the back-end FPGA and will therefore reduce the necessary bandwidth over the Ethernet. Additionally, automatic control and monitoring of the readout and the detectors will be added to the firmware. For the measurements done with the instrument in 2018, an expert software application was used with which the readout system was controlled and the events stored in the back-end memory buffer were read out.

3.2.5 Control and Data Readout Software

The control and data readout software operating on the DAQ computer is written in C++ with a graphical user interface (GUI) built using the Qt framework. In keeping with the theme of bear names used for the FPGA firmware, the control and readout software GUI was named Panda and an overview of the software can be seen in fig. 3.18. Communication with the back-end FPGA board is done through a dedicated Ethernet connection with the computer using the IPbus protocol [66]. A library called µHAL is provided with the IPbus software package that exposes a C++ API for read and write transactions. It is possible to directly communicate with the IPbus firmware in the back-end FPGA using µHAL but with reduced reliability. An additional software called ControlHub, which runs independently of the user application, acts as a mediator between the µHAL and the FPGA and will automatically try to correct any missing or damaged packets.

Figure 3.18: Block diagram of the Panda GUI software that is communicating with the ControlHub software using the µHAL library.

The Panda GUI provides access to all the available settings in the back-end

FPGA, front-end FPGAs, GBTx chips and Timepix3 detectors. It also monitors the temperature of the detectors, the synchronisation status of the data links inside the front-end FPGA, and the status of the back-end memory buffer. When the GUI is running it is continuously communicating with the back-end FPGA in the background and reading Timepix3 events from the buffer as soon as there is something available. The events are then displayed in a graphical preview window with options to view counts, ToT, ToT distribution or ToA for all the events that have been read out. A dedicated tab in the GUI, shown in fig. 3.19, was implemented for operation of the instrument during measurements, with the essential information provided to the user at a quick glance including the preview of the events. From this view the user can adjust the thresholds of the detectors, start and stop a measurement, adjust the trigger options and monitor the status of the system.

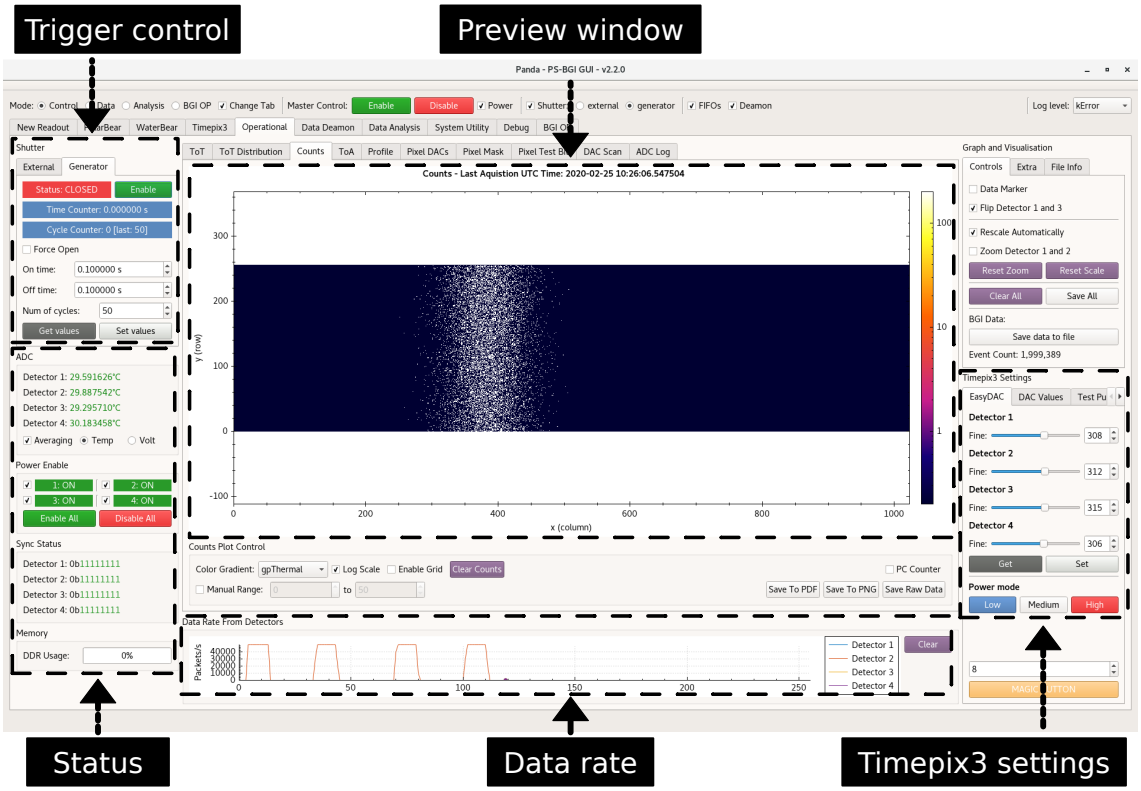

Figure 3.19: Screenshot of the Panda GUI with the operational tab open that provides the user a preview window of the acquired Timepix3 events and control of various parameters.

The events that are read out by the GUI can be saved to file for later analysis. Initially, a library called Protocol buffers was used to structure and serialise the events before writing them to file [67]. It was later replaced with a custom C++ class, which writes the events to a human-readable JSON text file, due to performance limitations in Protocol buffers when processing millions of events. The human-readable JSON format also has the advantage of being self-describing and not requiring specific libraries to be used.

The features in the Panda GUI mentioned thus far are those needed to operate the PS-BGI instrument for measurements. Before a measurement is started, the Timepix3 detectors have to be initialised using a startup sequence that can be triggered from the GUI.

Startup Sequence for the Timepix3 Pixel Detectors

The startup sequence, outlined in fig. 3.20, begins with a reset of the Timepix3 ASIC through a dedicated signal for each detector, which allows the operator to only run the sequence on specific detectors. After the reset, the Timepix3 is in a default state, with for example only one data link enabled, and requires some hardware settings to be configured. This includes configuring the clocks, data links and power settings to allow the front-end readout FPGAs to communicate with the Timepix3 over all data links. Next step is to configure each pixel in the matrix with individual threshold adjustments and mask settings obtained through an equalization method, which will be described in more detail later in section 3.3.2. The last steps prepares the detector for acquisition by setting the correct threshold and other digital-to-analog converter (DAC) values that controls the response to stimuli but also affects the power consumption.

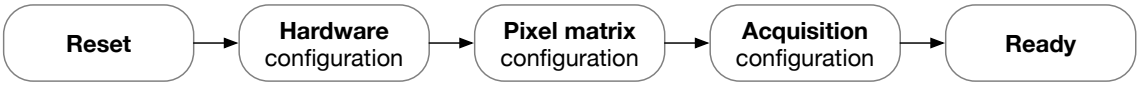

Figure 3.20: Startup sequence for the Timepix3 pixel detectors.

Data Acquisition Sequence for the Timepix3 Pixel Detectors

After the detectors have been configured for operation, the data acquisition involves only a few steps, shown in fig. 3.21. As soon as the injection trigger arrives to the back-end FPGA it will be sent over the fibres to the front-end FPGA that is connected to the shutter signal of the Timepix3. During the data readout the Timepix3 can operate in one of two different acquisition modes, data driven or sequential. In the data driven mode, data will be sent out as soon as the shutter is open and an event is detected, while in the sequential mode all the events are buffered until a specific read command is sent that will instruct the Timepix3 to send all the recorded data. When the extraction trigger is received the shutter will be closed and the data acquisition stops.

Figure 3.21: Data acquisition sequence for the Timepix3 pixel detectors.

3.3 Pixel Detector Response

The vast number of settings in the Timepix3 ASIC makes it a versatile readout chip for hybrid pixel detectors, but it also opens up the possibility to configure the chip in a way that does not lead to the desired response. Additionally, a comprehensive equalisation procedure is needed to normalise the response of each of the pixels in the detector. Basic health checks can be done with the Timepix3 including scanning the DAC settings and to monitor the built-in temperature sensor. Some of the Timepix3 settings also have an impact on the power consumption of the chip and can therefore be optimised to reduce the consumption, but it requires an understanding of the possible side effects. These topics will be discussed in this section.

3.3.1 DAC and Threshold Scan

There are DACs in the Timepix3 that can be controlled through registers. They all have an expected response by design and deviations from expectations could indicate that the chip is not working as intended. During the production test of the Timepix3 ASICs the DAC settings are scanned and only chips that are within a specified tolerance are kept. Measuring the response of the DACs during operation can also give an indication of the chip health and for the PS-BGI instrument such a test can be done when it is installed during operation. On the left in fig. 3.22 an example of a DAC scan is shown for a Timepix3 chip that is within tolerance and periodic scans similar to this throughout the life of the chip can be used to monitor for any deviations.

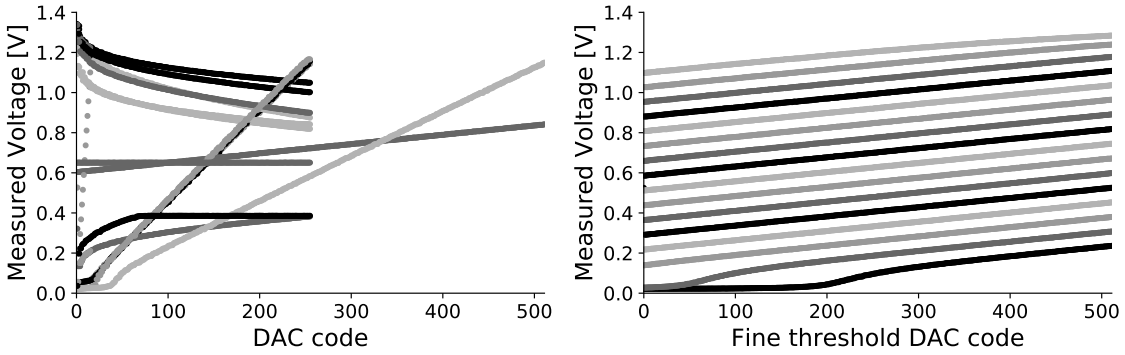

Figure 3.22: Example scan of all possible DAC settings (left) and threshold settings (right).

Two of the DACs in the Timepix3, "Vthreshold_fine" and "Vthreshold_coarse", are combined to generate the threshold voltage that is distributed to all discriminators in the pixel matrix. A similar health check as the DAC scan can be done for all the possible combinations for the threshold settings and an example can be seen on the right in fig. 3.22. Each line is a specific setting for "Vthreshold_coarse", with the lowest value at the bottom. The result from the threshold scan can also be used to determine the combination of "Vthreshold_fine" and "Vthreshold_coarse" that yields the desired threshold voltage.

3.3.2 Equalisation

The global threshold sets the voltage level at which all the pixels in the matrix should trigger if the input signal crosses the threshold. Unfortunately, due to process and design variations the pixels will not have the exact same response and therefore an adjustment is needed to ensure each pixel responds to an input stimulus in a similar manner. This can be achieved using a pixel DAC that is included with every pixel in the Timepix3. The pixel DACs have 16 levels that can be set individually for each of the 65 536 pixels, where an equalisation method is used to choose the optimal values.

The equalisation method takes advantage of the inherent analog noise in the Timepix3 ASIC, which can be considered stable if the temperature of the chip is kept constant. The threshold must be kept above the noise for a pixel to be functional, otherwise it will output a signal due to triggering on the noise. By moving the threshold closer to the noise there is a point at which the pixel will be at the so called noise centre, which is what we will use in the equalisation method as a measure of the pixel response. The noise centre is defined as the threshold value where the pixel has the highest number of counts. This means that the method can also be used to identify noisy pixels, which are pixels that will always have a high count value even when the threshold is far away from the noise. In fig. 3.23 an example is shown for measurements of the noise centre for all the pixels in a Timepix3 pixel detector. On the left in black is the result when the pixel DAC is set to the minimum of 0 for all pixels and on the right in grey is when it is set to the maximum of 15.

We will now use the information in this measurement to calculate an optimal pixel DAC value for each pixel such that the distribution for the noise centre measurement is as narrow as possible, which means that all pixels have a similar response to an input signal for a fixed global threshold. The equalisation procedure can be summarised in the following steps:

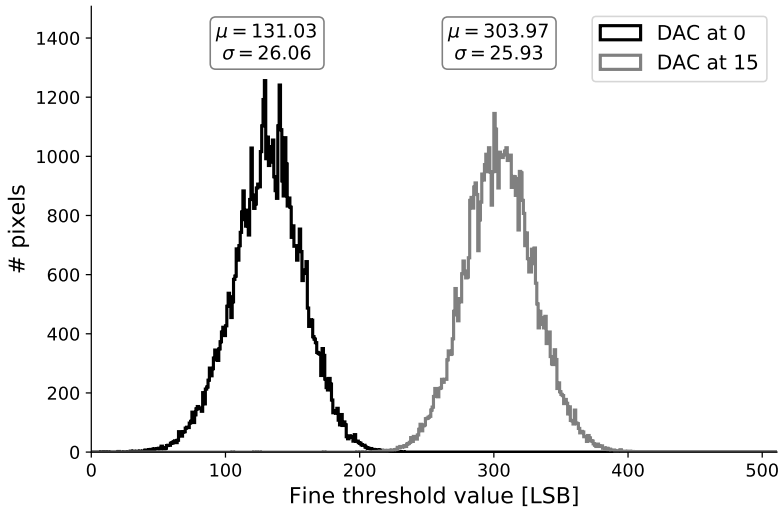

Figure 3.23: Example of measurement showing number of pixels which have a noise centre at a specific threshold value with all pixel DACs set to 0 (black) and to 15 (grey).

- Calculate a target threshold located between the maximum and minimum distributions seen in fig. 3.23.
- For each pixel, find the pixel DAC value which is closest to the target by interpolating between the measured maximum and minimum values.
- Scan again for the noise centre and set a new target threshold as the mean of measured distribution.
- For each pixel, see if there is a DAC value that brings it closer to the new target.
- Iterate the last two steps until there is no observable change.

The result from an equalisation can be seen in fig. 3.24 where the blue distribution in the centre is the distribution of the measured noise centers after the optimised pixel DAC values have been set. The standard deviation of the distribution has decreased, from 26 LSB down to 6.3 LSB, which can be interpreted as more pixels having an equal response to the global threshold. Two iterations were needed to achieve this equalisation result and any further iterations did not yield any improvements.

The Timepix3 has to be kept at a constant temperature and be powered from a stable power source for the duration of the equalisation process to ensure a consistent

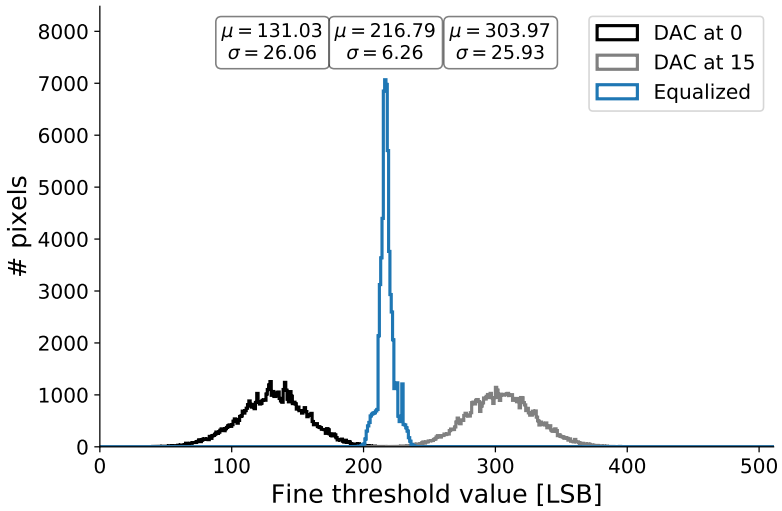

Figure 3.24: Example of noise centre measurement after equalisation (blue).

and reproducible result. This is challenging due to the high data rate that the noise centre scan generates. As soon as the threshold gets close to the noise centre of a pixel it will start to trigger and send data as fast as it can, which draws the maximum amount of current. If we then consider that more than 1000 pixels can have the noise centre at the same threshold, as was seen in fig. 3.23, the total amount of current can cause voltage drops inside the chip that ruins the equalisation result. To get around this problem we can take advantage of the ability to mask a pixel from generating data. In the measurements shown here, the equalisation was done on a quarter of the pixel matrix at a time. This improves the results, but due to a design choice in Timepix3, the masked pixels are still powered up and will therefore consume power even though they are masked. A feature to mask and disable the power to individual pixels would therefore be required to improve the equalisation further.

Power consumption is also of concern during the operation of the detectors in the instrument, where a shower of particles hitting the detectors could trigger enough events to cause voltage drops similar to those seen during the equalisation. There are features built into the Timepix3 that will reduce the dynamic and static power and allows for operating the chip in different power modes. These modes involves

changing DAC settings that affect the response of the pixels and therefore also the result of the equalisation. It is therefore important to run the equalisation with the Timepix3 settings and at the temperature that will be used during operation to ensure the response of the pixels is consistent.

3.3.3 Power Considerations

Cooling is a big concern for operating Timepix3 detectors inside the UHV of an accelerator as the majority of the heat has to be transferred through conduction, which is done by means of a cooled copper plate in the PS-BGI as discussed earlier in section 3.1.3. Reducing the power consumption will also reduce the amount of heat that has to be transferred away from the detectors. Two methods of reducing the power consumption of the Timepix3 have been used in this thesis project. The first method involves running the chip in different power modes, which means changing the DAC settings to reduce the biasing current of transistors in the analog parts of the pixels, while the second method uses a feature in the Timepix3 that distributes the clock phases to reduce peak dynamic power.

Power Modes

A low power mode and a high power mode have been used based on values provided by the chip designers. In the PS-BGI instrument, when the detectors were taking measurements in high power mode, the temperature inside the detectors would quickly rise and after about 1 min the system had to be turned off, which was an indication that the thermal interface between the detector assembly and the cooling plate was not sufficient. The majority of the measurements from 2018 using the Mk II detector assembly were therefore done in low power mode where the temperature could be kept stable. In low power mode the response of the pixels is slower and the minimum threshold is higher, which is caused by reduced noise margins in the analog circuitry [68].

In fig. 3.25 the measured distribution of ToT values in low and high power modes are shown where the peak is located at higher ToT for low power mode. The ToT value comes from a counter in the pixels that indicates how long the signal pulse was above the threshold, which means that in low power mode where the response is slower the signal will stay above the threshold longer. Another side effect of this is that each pixel will take longer to recover and be ready to trigger on a new particle event and therefore causes a reduction in the maximum pixel event rate. We can observe this for the measurements in fig. 3.25 where there is a 40 % data reduction in low power mode compared to high power mode. This can lead to systematic errors in the beam profile measurement if an ionisation electron trigger a pixel before it has finished processing the previous event. We will discuss this in more detail later in section 3.5.3.

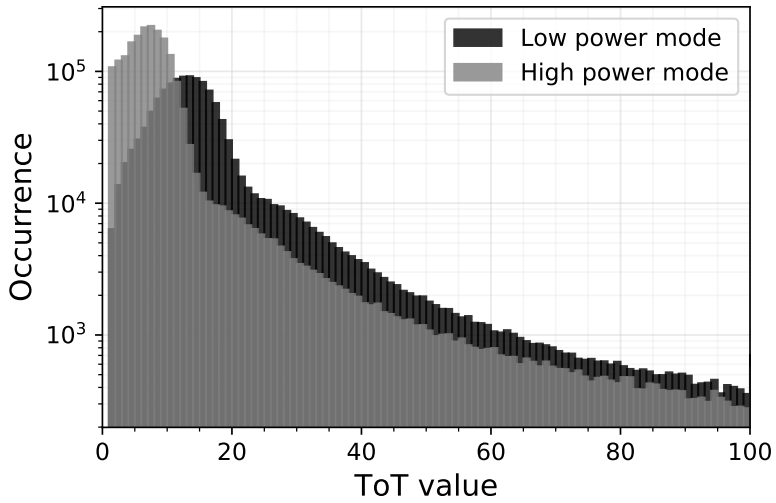

Figure 3.25: Distribution of ToT values for measured data in low and high power mode.

Clock Phases

The clock that is used in all pixels is distributed along 128 pairs of columns, called double columns, and therefore has a significant impact on the peak dynamic power consumption, due to the large fanout in the routing and the need for strong drivers

to reach all the pixels. To reduce the impact this pixel clock has on the power, the phase of the clock going to each double column can be configured inside the Timepix3. As default after reset, a single phase is used, which means all double columns are clocked at the same time. The next option is to have two clock phases for every other double column, such that for example when one double column has a high clock signal it is low in the neighboring double column. Additionally, the chip can use 4, 8 or 16 different clock phases with a reduction in peak power in each step.

The different phases will have an impact on the timestamp for the recorded events due to the shift of the clock. In fig. 3.26, a test pulse has been sent to all pixels in the matrix to trigger them at the same time. On the left in the figure the timestamp is taken directly from the Timepix3 events without correction and 8 regions are visible, which is due to the 16 clock phases distributed over the 128 double columns. We can correct for this clock shift because it is constant, and the result can be seen on the right in the figure. The difference between the highest ToA value and the lowest is now around 3 ns compared to a whole clock cycle (25 ns) without the correction. The remaining difference is due to a limitation in the test pulse signal that has to propagate from the first row to the last row, which takes some time.

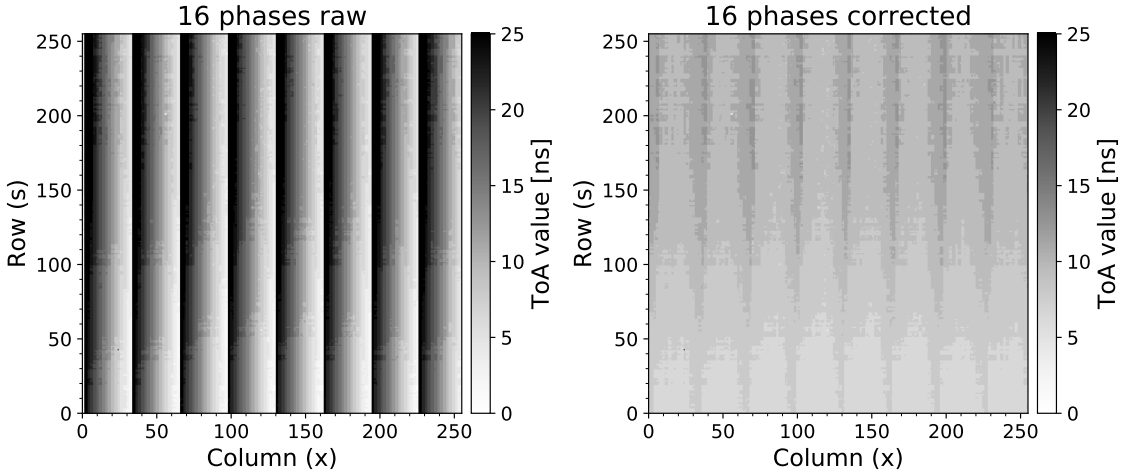

Figure 3.26: Timepix3 clock phase distribution before (left) and after (right) correction.

3.4 Detection of Ionisation Electrons

Now that we have an understanding of the mechanical design of the instrument and the detector assembly, how the readout electronics can transfer data from the pixel detectors to a computer and how the Timepix3 responds, we can investigate if it is possible to detect ionisation electrons generated from beam gas ionisation. We will take advantage of our knowledge of the orientation of the detectors detailed in section 3.1.4 and our understanding of the transport of the ionisation electrons.

3.4.1 Ionisation Electron Transport

The ionisation electrons created at the interaction point between the particles in the accelerator beam and the residual gas are transported to the detectors with the electric and magnetic field, as discussed earlier in section 2.3.2. When the electrons reach the detectors they have gained an energy of 10 keV due to the acceleration, assuming they were created in the middle between the cathode and anode electrodes. Additionally, a dynamic electric field caused by the bunched particles in the beam will influence the ionisation electrons (and ions). With the static electric drift field enabled we expect a uniform directionality of the accelerated ionisation electron towards the detectors, while if it is the disabled the electric field from the bunch will result in a time dependent direction and magnitude of the acceleration.

We can test this hypothesis by taking measurements with and without an electric drift field in the presence of the accelerator beam. The results can be seen in fig. 3.27 where the signal without electric field is seen on the left and with is on the right. For both of these measurements, the magnetic field was present and kept at its nominal strength of 0.2 T and the Timepix3 detectors were operating in high power mode. Without the electric field the detected particles are randomly distributed all over the detector area and with the electric field distinct patterns are visible. The honeycomb shaped RF shield is visible as a shadow where there are low number of

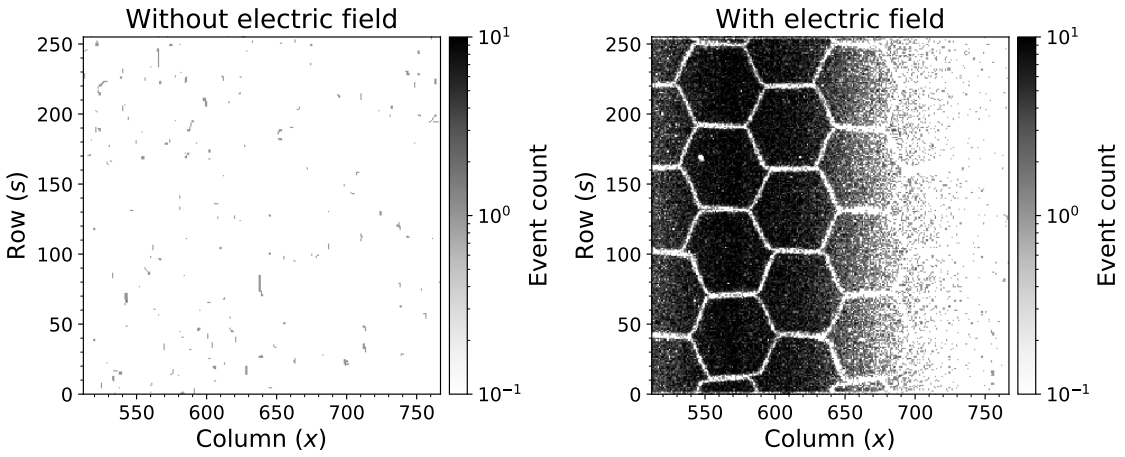

Figure 3.27: Events recorded without (left) and with (right) electric drift field applied.

events detected, which fits with the reasoning that the ionisation electrons do not have enough energy to pass through the walls of the shield. We can also observe that the orientation of the detected signal matches with our understanding of the orientation of the detectors discussed above. In the row direction there is a uniform signal (if we disregard the shield shadows) while in the column direction the signal fades away as we go to higher column numbers, which are presumably beyond the size of the beam. This indicates that there is a distinct difference in the detected signal with and without the electric drift field as hypothesised. Before we can definitively decide if the detected signal with the electric drift field is due to ionisation electrons, we need to examine it further.

3.4.2 Clustering with DBSCAN

Particles that interact with the silicon sensors will deposit different amount of energy which leads to triggering of one or more neighboring pixels, assuming the charge created from the deposited energy is high enough to trigger pixels. Depending upon the trajectory and energy of the particle, we can identify different shaped tracks and blobs from the triggered pixels. We can categorize these pixel events into clusters using an algorithm called DBSCAN [69] which has two important parameters:

- ϵ: the maximum distance between two points in order to form a cluster
- N_{min}: the minimum number of points for a cluster

The input to DBSCAN can, for example, be the 2D pixel image seen on the left in fig. 3.27, from which the algorithm will identify the tracks and blobs and assign a unique identifier to each. We can see in fig. 3.27 that the clusters are formed by pixels that are close together and we can therefore use a value of $\epsilon = 2.5$ pixels. For N_{min} we can use a value of 2, which will enable the identification of clusters with a size ≥ 2. Later in section 4.2.2 we will use this algorithm to separate the ionisation electron events from the background events. To achieve this, we first need to understand the characteristics of the ionisation electrons.

3.4.3 Characteristics of Detected Ionisation Electrons

The geometry of the IPM means that the trajectory of the ionisation electrons is perpendicular to the detector. In section 2.4.1 the average path length of a 10 keV electron in silicon was calculated as 1.48 µm. This short average path length compared to the 55 µm pixel size and the perpendicular direction of the ionisation electrons means that we can expect mostly single pixel events. If an electron hits between two pixels, the charge in the sensor might be distributed between these two and form a two pixel cluster. We will assume for this discussion that a majority of the ionisation electrons will create single pixel events.

If we run the DBSCAN algorithm on the same data sets shown in fig. 3.27, we can calculate the size of the identified clusters (number of pixels). In fig. 3.28, the distribution of cluster size is shown for the cases with and without the electric drift field. For the cluster size of 1, which are the single pixel events, we can see an increase of more than two orders of magnitude when the electric field is active. Based on the assumption that ionisation electrons create single pixel events and that they are accelerated towards the detector when the electric field is active, we can conclude that this increase is due to the presence of ionisation electrons.

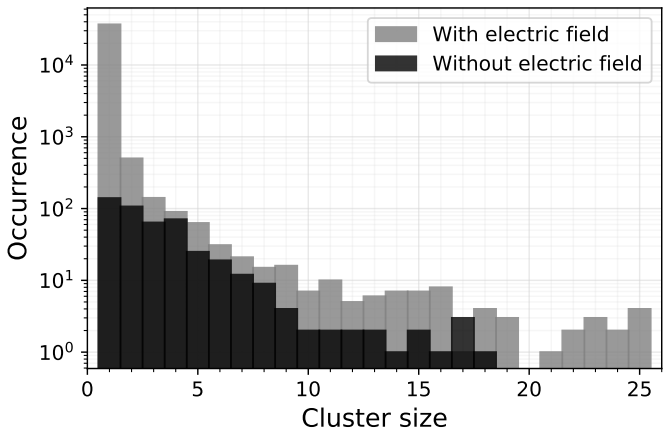

Figure 3.28: Distribution of cluster sizes with and without electric field.

Furthermore, for clusters of size 2 there is a factor 4 increase, which could be due to ionisation electrons that hit on the border between two pixels, and also for larger clusters there is an overall increase when the electric field is applied.

As the major increase is for single pixel events we will focus on these for our continued discussion. The time-over-threshold (ToT) information from the Timepix3 is correlated to the amount of charge collected from the sensor in each pixel, as detailed in section 2.4.2. A minimum ionising particle (MIP) have a most probable energy loss in silicon of 3.72 MeV/cm, which means 37.2 keV in a 100 µm thick sensor [70]. This creates 10 000 electron-hole pairs in the sensor compared to an ionisation electron in the PS-BGI instrument which will deposit at most 10 keV into the sensor and therefore create 2800 electron-hole pairs. A greater number of electron-hole pairs will lead to more charge collected in the pixel, which in turn will result in a higher ToT value. We therefore expect MIPs to have a higher ToT compared to ionisation electrons. A MIP is used in this discussion for ease of comparison, but beam loss particles in an accelerator come from a wide range of sources with a wide energy spectrum. In contrast, the ionisation electrons have a well defined and narrow energy range.

In fig. 3.29, the distribution of ToT values for single pixel events is shown for the case with and without the electric field. Without the electric field we expect to

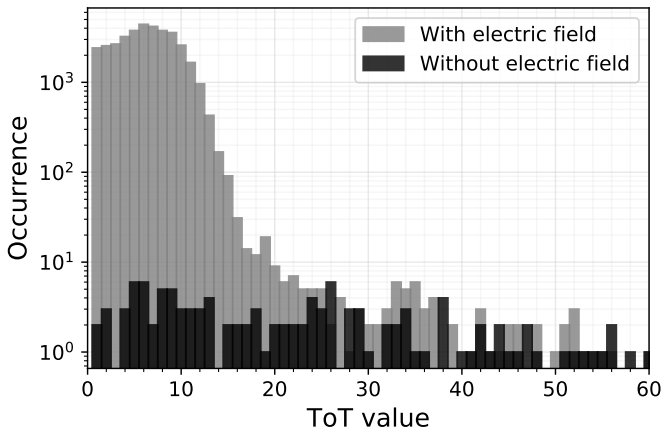

Figure 3.29: Distribution of ToT values of single pixel clusters with and without electric field.

only see beam loss particles that form clusters of more than one pixel, unless the trajectories of the particles are perpendicular to the sensor which could lead to single pixel events. The low number of events we see in fig. 3.29 without the electric field are therefore likely these perpendicular beam loss particles. On the other hand, with the electric field we expect that ionisation electrons reach the detector and create a large number of single pixel events with a relatively low ToT, which is exactly what we see in fig. 3.29. Beam loss should also be present when the electric field is enabled, which we can see in the figure as the higher ToT values.

Based on these observations we can say with high confidence that we are able to detect ionisation electrons that originate from the particle beam's interaction with the residual gas in the vacuum chamber. Later in section 4.2.2 we will use this knowledge to separate the ionisation electron events from the rest, which we will call background events. For some applications of pixel detectors, the energy information of each cluster is the important information to measure. This requires a conversion function to relate the ToT values to the corresponding energy. Such energy calibration procedures are well documented in literature and relies on radioactive sources with peaks in their energy spectra at known and precise values [71]. For this application in a beam profile monitor we are interested in separating the ionisation

electron events from the background events. Determining the exact energy of the events is not of a major interest.

3.5 PS-BGI Response and Measurement Concepts

The PS-BGI instrument, with the Timepix3 hybrid pixel detectors, is able to detect the ionisation electrons that originate from the interaction between the beam in the accelerator and the residual gas in the vacuum chamber. The next step is to understand how the instrument responds to different beams and environmental conditions, which will determine how beam profile and beam size measurement using the instrument can be carried out. Throughout this section we will use two different beams, each using either protons or Pb^{54+} ions, and discuss how the beam parameters affect the ionisation electron yield and rate. First, in section 3.5.1 we will briefly detail the accelerator timing and how it impacts the ionisation electron rate and after in section 3.5.2 the yield will be discussed in detail. Lastly, the limitations in data rate of the Timepix3 detectors and how they impact the instrument will be discussed in section 3.5.3.

3.5.1 Accelerator and Instrument Timing

The PS accelerator is a synchrotron and the particles in the beam travel clockwise around a closed orbit with a circumference of 628 m and a period t_{rev}. The beam consists of one or several bunches ($n_{\mathrm{b}} \geq 1$) that are spaced t_{b} apart, shown in fig. 3.30. These three parameters have an impact on the rate at which ionisation electrons are created and we therefore need to consider these to determine if a beam profile measurement using the PS-BGI instrument is feasible.

The instrument can operate in different modes, which are coupled to these accelerator timing parameters and leads to different requirements of the minimum and maximum ionisation electron rate. The modes can be categorised into four groups

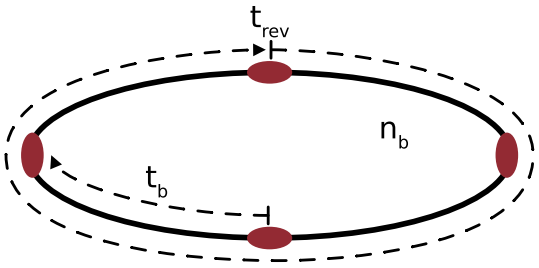

Figure 3.30: Illustration of four bunches ($n_\mathrm{b} = 4$) spaced t_b apart circulating clockwise in the PS with a revolution period of t_rev.

shown in order of decreasing complexity below.

- Single turn, bunch-by-bunch measurement: the beam profile of each bunch is measured for every turn.
- Multi turn, bunch-by-bunch measurement: the beam profile of each bunch is measured and integrated over a number of turns.
- Single turn measurement: the integrated beam profile of all bunches is measured for every turn.
- Multi turn measurement: the integrated beam profile of all bunches over multiple turns is measured.

To synchronise the measurements with the accelerator a trigger card is connected to the readout electronics, discussed earlier in section 3.2. The measurements can therefore be started and stopped synchronous to events in the accelerator, such as injection and extraction. In addition, the Timepix3 hybrid pixel detectors have a timing resolution of 1.5625 ns which is faster than the shortest bunch spacing (t_b) as will be shown below.

The timing parameters vary throughout the acceleration cycle in the PS and the two extremes are at injection and prior to extraction. In table 3.4 the accelerator timings of injection and extraction are shown for the current LHC type beams and for the future HL-LHC beams in the PS, with protons (p) and lead ions (Pb^{54+}). We can see that for most cases, the number of bunches (n_b) increase from injection to extraction, which is done by splitting the original bunches into smaller bunches and

Table 3.4: Accelerator timings for LHC type beams in the PS at injection and extraction, from [3, 72].

Beam (Type):	LHC (p)		HL-LHC (p)		LHC (Pb^{54+})		HL-LHC (Pb^{54+})	
	inj.	ext.	inj.	ext.	inj.	ext.	inj.	ext.
E_k [GeV]	1.4	25	2.0	25	15	1222	15	1222
n_b	6	72	6	72	2	2	2	4
t_b [ns]	284	25	284	25	354	100	354	100
$B_l = 4\sigma_l$ [ns]	180	4	205	4	170	4	200	4
t_{rev} [µs]	2.287	2.096	2.211	2.096	5.636	2.115	5.636	2.115

therefore also reduces the number of particles per bunch. At extraction the bunch spacing (t_b) is shorter which leads to a higher maximum ionisation electron creation rate and could reach the limits of the Timepix3. We will discuss this potential issue later in section 3.5.3. The bunch length (B_l) is commonly defined as four times the standard deviation of the longitudinal distribution ($4\sigma_l$) and therefore represents 95 % of all particles in the bunch. The revolution periods (t_{rev}) in table 3.4 are calculated using the formula

$$t_{rev} = \frac{C}{\beta_r \cdot c}, \tag{3.2}$$

where C is the circumference of the accelerator (628 m for the PS), c is the speed of light in vacuum and β_r is the relativistic beta. This can be calculated from

$$\beta_r = \sqrt{1 - \left(1 + \frac{E_k}{m_0 \cdot c^2}\right)^{-2}}, \tag{3.3}$$

where E_k is the kinetic energy of the particle and m_0 is the rest mass. For protons $m_0 = 938.272\,\mathrm{MeV}/c^2$ and for Pb^{54+} ions $m_0 = 207.2 * 931.494\,\mathrm{MeV}/c^2 = 194\,409.977\,\mathrm{MeV}/c^2$.

3.5.2 Expected Ionisation Electron Yield

In section 2.3.1 a formula for estimating the ionisation electron yield per bunch and per turn of the accelerator was derived as

$$n_{\mathrm{e,b,t}} = L_{\mathrm{s}} \cdot \sigma_{\mathrm{ion}} \cdot N_{\mathrm{b}} \cdot p \cdot \frac{1}{k \cdot T}, \qquad \text{(2.26 revisited)}$$

It is important to quantify the expected number of electrons as it will determine the feasibility of using the PS-BGI instrument under specific conditions. Some of the parameters in the formula are constrained, such as the detector length (L_{s}) of 14 mm for the Timepix3 and the temperature T which is kept at around 21 °C (294 K) in the PS. The number of particles per bunch (N_{b}) can be measured using beam current monitors. The ionisation cross section (σ_{ion}) depends on the composition of the residual gas, the type of particles in the beam and the kinetic energy of those particles. The PS operates mainly with protons with a kinetic energy (E_{k}) from 1.4 GeV at injection up to 25 GeV at extraction. It also operates with different types of stripped ions, such as Pb^{54+}, with a kinetic energy of 0.0722 GeV/u at injection up to 5.9 GeV/u at extraction [72]. The kinetic energies for the ions have been normalised to the atomic mass unit, which in the case of Pb^{54+} with an atomic mass of 207.2 u leads to kinetic energies from 15 GeV to 1222 GeV. An expression for the ionisation cross section can be derived from Bethe theory [73–75] as

$$\sigma_{\mathrm{ion}} = 4\pi Z^2 \left(\frac{\hbar}{m_e c}\right)^2 \beta_r^{-2} \left(M^2 x + C\right), \qquad (3.4)$$

where Z is the charge of the particles in the beam ($Z = 1$ for protons and $Z = 54$ for Pb^{54+} ions), $\hbar$ is the reduced Planck constant ($\hbar = \frac{h}{2\pi}$), m_e is the mass of an electron and c is the speed of light in vacuum. C and M^2 are obtained from measurements in [29] and

$$x = \ln\left(\frac{\beta_r^2}{1-\beta_r^2}\right) - \beta_r^2, \qquad (3.5)$$

where β_r is calculated using eq. (3.3). Ionisation cross sections are commonly expressed in units of Barn [b], with $1\,\mathrm{b} = 10^{-28}\,\mathrm{m}^2$. In table 3.5, the ionisation cross section is calculated for argon (Ar) and hydrogen (H_2) gas using the equations above for the same beam types shown earlier in table 3.4. We can see from table 3.5 that the cross section is only about a factor 4 to 5 different between argon and hydrogen gas, while operating with Pb^{54+} ions instead of protons increases the cross section by several orders of magnitude.

Table 3.5: Expected ionisation cross sections for LHC type beams in the PS.

Beam (Type):	LHC (p)		HL-LHC (p)		LHC (Pb^{54+})		HL-LHC (Pb^{54+})	
	inj.	ext.	inj.	ext.	inj.	ext.	inj.	ext.
E_k [GeV]	1.4	25	2.0	25	15	1222	15	1222
β_r	0.916	0.999	0.948	0.999	0.372	0.991	0.372	0.991
$\sigma_{\mathrm{ion,H_2}}$ [Mb]	0.194	0.226	0.188	0.226	2670	567	2670	567
$\sigma_{\mathrm{ion,Ar}}$ [Mb]	0.924	1.16	0.904	1.16	11 700	2810	11 700	2810

The residual gas composition can vary with time due to, for example, outgassing of components in the vacuum. If a gas injection system is used with the instrument, the gas composition in the ionisation region can be controlled. For the operational PS-BGI instrument, aimed to be completed at the end of 2020, a gas injection system with argon creating a pressure of around $1 \times 10^{-8}\,\mathrm{mbar}$ in the ionisation region is foreseen. Two main cases can therefore be considered, the first with an assumed residual gas mainly composed of hydrogen (H_2) and the second with the gas injection. The static pressure (p) can vary by more than an order of magnitude in the PS from around $1 \times 10^{-9}\,\mathrm{mbar}$ in the beginning of the year to $1 \times 10^{-10}\,\mathrm{mbar}$ at the end. We will calculate the yield for an average pressure of $5 \times 10^{-10}\,\mathrm{mbar}$, which can then be interpolated to other pressures using the linear relationship between electron yield and pressure in eq. (2.26).

In table 3.6 the expected ionisation electron yields ($n_{\mathrm{e,b,t}}$) are calculated for the different cases based on the beams that are produced for the LHC in the PS. The number of expected electrons per turn and per bunch is denoted $n_{\mathrm{e,b,t}}$, and

Table 3.6: Expected ionisation electron yield for different beams and conditions in the PS. Beam parameters are taken from [3, 72].

Beam (Type):	LHC (p)		HL-LHC (p)		LHC (Pb^{54+})		HL-LHC (Pb^{54+})	
	inj.	ext.	inj.	ext.	inj.	ext.	inj.	ext.
E_k [GeV]	1.4	25	2.0	25	15	1222	15	1222
N_b [10^{10} p/b]	168	13.0	325	26.0	0.0550	0.0510	0.109	0.0500
n_b	6	72	6	72	2	2	2	4
Residual gas (H_2, 5×10^{-10} mbar)								
$n_{e,b,t}$	5.62	0.506	10.5	1.01	25.3	4.99	50.2	4.89
$n_{e,t}$	33.7	36.4	63.2	72.9	50.6	9.97	100	19.6
Gas injection (Ar, 1×10^{-8} mbar)								
$n_{e,b,t}$	535	51.9	1010	104	2220	494	4410	485
$n_{e,t}$	3210	3740	6080	7480	4450	989	8810	1940

by multiplying this by the number of bunches in the beam we get the number of electrons per turn, i.e. for each passing of the whole beam through the PS-BGI instrument, as

$$n_{e,t} = n_{e,b,t} \cdot n_b. \tag{3.6}$$

If we are integrating the beam profile measurements over all bunches, $n_{e,t}$ is the number of interest, while $n_{e,b,t}$ determines whether bunch-by-bunch measurements are feasible or not. The electron yield is higher for the gas injection case by about a factor 100, which is expected due to the higher ionisation cross section for argon seen in table 3.5 and the higher pressure. To get more than 50 ionisation electrons from a single passing of a bunch (turn-by-turn measurement), gas injection is necessary. We will test this statement later in section 5.4, which details a turn-by-turn measurement using the PS-BGI during which the pressure was increased during a short period of time by sublimating a vacuum pump to release molecules into the ionisation volume.

3.5.3 Detector Hit Rates

The expected yield indicates how many ionisation electrons are available to be detected by the Timepix3 hybrid pixel detectors in the PS-BGI. As discussed in sec-

tion 3.1.3, the detector assembly consists of four Timepix3 detectors mounted side by side along the transverse direction x, giving a total width of 56 mm. This means that the beam is in the direction of the rows in the detectors, as shown in fig. 3.6. Inside the Timepix3, pixel events are transported down through dedicated buses in the columns to the common readout at the bottom where events from all the 256 columns are collected before being sent out of the detectors. This means that it takes longer to read out 256 pixels in a column than it takes to read out 256 pixels in a row. Unfortunately, the detectors in the PS-BGI had to be oriented as shown in fig. 3.6 for the signals to be routed from the detectors, which reduces the read out rate.

Each Timepix3 is capable of processing up to 85 Mevents/s (see section 2.4.2), where the majority of the ionisation electron hits will generate one event. The important caveat is that this is the maximum processing rate under the assumption that all events are uniformly distributed over the full detector, which is not the case for the PS-BGI as discussed above. Another limitation in the detector is the overflow of the ToA counter after 409.6 µs. If an event has not been read out within this time, the timing information is lost. It is therefore important to determine the expected detector hit rates by combining the timing information in table 3.4 with the expected ionisation electron yield in table 3.6.

The first measure we can calculate is the average number of ionisation electrons that will hit the detectors per second as

$$R_{\mathrm{det}}^{\mathrm{avg}} = \frac{n_{\mathrm{e,t}}}{t_{\mathrm{rev}}}, \tag{3.7}$$

which gives an indication if the detectors can sustain a continuous measurement of the beam. This assumes that the creation of ionisation electrons is uniformly distributed over the circumference of the machine, which is an approximation since the beams are usually bunched. We therefore expect bursts of hits when a bunch

passes through the instrument and no hits between bunches. We can estimate the maximum hit rate during the centre of the bunches by observing that 68 % of the particles in the bunches are located within $\pm\sigma_l$ of the longitudinal mean μ_l. With the bunch length defined as $B_l = 4\sigma_l$, we can calculate a maximum hit rate at the centre of the bunch as

$$R_{\text{det}}^{\text{max}} = 0.68 \cdot \frac{n_{\text{e,b,t}}}{2\sigma_l}. \tag{3.8}$$

The maximum hit rate is only valid during a fraction of the revolution period (t_{rev}) in the machine which we can express as

$$p_{\text{t,max}} = \frac{2\sigma_l n_b}{t_{\text{rev}}}. \tag{3.9}$$

The size of the beam also affects the detector hit rate as a wider beam will spread the electrons over more pixels in the detectors and a smaller beam over fewer pixels. A Timepix3 pixel is 55 µm in the transverse direction x and contains 256 pixels in the direction s of the beam. One of the variables in the equation for the ionisation electron yield (eq. (2.26)) was the length of the detector along s, which means that the yield is distributed over these 256 pixels. With the size of a pixel, the number of pixels in the s direction and the beam size σ_x, we can calculate how many pixels are hit on average by 95 % of the ionisation electrons from the beam as

$$R_{\text{pixel}}^{\text{avg}} = R_{\text{det}}^{\text{avg}} \cdot \frac{55 \times 10^{-6}}{4\sigma_x \cdot 256}. \tag{3.10}$$

Pixels near the centre of the beam profile will also be hit more often than pixels on the edge due to the distribution of the profile. We can therefore calculate a maximum hit rate for the 68 % of the ionisation electrons that are located near the mean of the beam profile distribution as

$$R_{\text{pixel}}^{\text{max}} = R_{\text{det}}^{\text{max}} \cdot 0.68 \cdot \frac{55 \times 10^{-6}}{2\sigma_x \cdot 256}. \tag{3.11}$$

Equations (3.10) and (3.11) require the beam size σ_x, which is known for the currently operational LHC type beams in the PS with protons and Pb^{54+} but not for the HL-LHC type beams since the optics function is not defined at the moment. We can rewrite eq. (2.24) to get an expression for the beam size as

$$\sigma_x = \sqrt{\frac{\beta_x}{\gamma_r \beta_r}\epsilon_{x,N} + D_x^2 \sigma_{\delta_p}^2}. \tag{3.12}$$

In table 3.7 beam parameters are presented for the different cases used before. The optics parameters β_x and D_x are kept the same for LHC and HL-LHC beams due to the missing optics model. Likewise, for the momentum spread (δ_p) values are kept the same where there was no reference to a different value. At the bottom of the table are the calculated beam sizes σ_x using eq. (3.12).

Table 3.7: Beam parameters calculated and taken from [3, 23, 72].

Beam (Type):	LHC (p)		HL-LHC (p)		LHC (Pb^{54+})		HL-LHC (Pb^{54+})	
	inj.	ext.	inj.	ext.	inj.	ext.	inj.	ext.
β_x	11.7	11.8	11.7	11.8	11.6	11.8	11.6	11.8
D_x	2.39	2.33	2.39	2.33	2.38	2.33	2.38	2.33
δ_p [10^{-3}]	0.90	0.90	1.5	1.5	0.40	0.20	0.40	0.20
$\epsilon_{x,N}$ [µm]	2.25	2.36	1.80	1.89	1.30	1.40	1.30	1.40
β_r	0.916	0.999	0.948	0.999	0.372	0.991	0.372	0.991
γ_r	2.49	27.6	3.13	27.6	1.08	7.29	1.08	7.29
σ_x [mm]	3.39	1.00	2.66	0.899	6.14	1.51	6.14	1.51

We can now calculate the expected hit rates for the full detector ($R_{\text{det}}^{\text{avg}}$ and $R_{\text{det}}^{\text{max}}$) and for a pixel ($R_{\text{pixel}}^{\text{avg}}$ and $R_{\text{pixel}}^{\text{max}}$) and the result can be seen in table 3.8 below. The units for the gas injection hit rates are three orders of magnitude greater than the residual gas rates due to the high ionisation electron yield from table 3.6.

Table 3.8: Expected detector hit rates for different beams and conditions in the CERN PS.

Beam (Type):	LHC (p)		HL-LHC (p)		LHC (Pb^{54+})		HL-LHC (Pb^{54+})	
	inj.	ext.	inj.	ext.	inj.	ext.	inj.	ext.
$p_{\text{t,max}}$ [%]	24	6.9	28	6.9	3.0	0.19	3.5	0.38
Residual gas (H_2, 5×10^{-10} mbar)								
$R_{\text{det}}^{\text{avg}}$ [Mhits/s]	14.7	17.4	28.6	34.8	8.98	4.72	17.8	9.25
$R_{\text{det}}^{\text{max}}$ [Mhits/s]	42.4	172	69.9	344	203	1700	341	1660
$R_{\text{pixel}}^{\text{avg}}$ [hits/s]	233	930	577	2080	78.5	167	156	328
$R_{\text{pixel}}^{\text{max}}$ [khits/s]	234	3200	491	7160	616	20 900	1040	20 500
Gas injection (Ar, 1×10^{-8} mbar)								
$R_{\text{det}}^{\text{avg}}$ [Ghits/s]	1.40	1.78	2.75	3.57	0.789	0.468	1.56	0.917
$R_{\text{det}}^{\text{max}}$ [Ghits/s]	4.04	17.7	6.72	35.3	17.8	168	30.0	165
$R_{\text{pixel}}^{\text{avg}}$ [khits/s]	22.2	95.4	55.5	213	6.89	16.6	13.7	32.5
$R_{\text{pixel}}^{\text{max}}$ [Mhits/s]	22.3	329	47.2	735	54.1	2080	91.2	2040

If we assume that each hit of an electron creates one event in the Timepix3 detector, we can compare the values in table 3.8 with the maximum processing of 85 Mevents/s of the Timepix3. From the table we can observe that the average hit rates over the detectors ($R_{\text{det}}^{\text{avg}}$) for residual gas are all below this limit. This implies that a continuous measurement can be sustained under the assumption that the ionisation electrons will hit the detector uniformly over one beam revolution. On the other hand, the maximum hit rates when the centre of the bunch passes through the instrument ($R_{\text{det}}^{\text{max}}$) indicate that the Timepix3 processing limits are reached for most of the beams, only LHC (p) and HL-LHC (p) at injection are below the limit. Although, the maximum is only valid for a fraction of the full beam revolution period as indicated by $p_{\text{t,max}}$, which is below 7 % for the cases where the rate is above the limit.

The pixel hit rates ($R^{\text{avg}}_{\text{pixel}}$ and $R^{\text{max}}_{\text{pixel}}$) give an indication whether a pixel will manage to process an event before the next one arrives. We saw earlier in section 2.4.2 that each pixel has a dead time of 475 ns plus the length of the signal pulse, which is quantified by the ToT value. In fig. 3.29 we see that for ionisation electrons we can expect ToT values of up to 20, which results in a dead time of 1 µs. This means that once a pixel has been hit it will not process another hit until 1 µs has passed and the pixels therefore have a processing rate of 1 Mhits/s. We can see in table 3.8 that the majority of cases are below this limit for $R^{\text{avg}}_{\text{pixel}}$ and $R^{\text{max}}_{\text{pixel}}$, except for some of the $R^{\text{max}}_{\text{pixel}}$, which again, are only valid during a fraction of the revolution period.

The expected hit rates for when gas injection is used are about 100 times greater compared to residual gas, which we already saw for the expected ionisation electron yield in table 3.6. This enables turn-by-turn measurements due to the greater number of ionisation electrons but it also limits the length of the measurements due to the increased hit rate. Even the average hit rates in table 3.8 with gas injection are above the limits of the Timepix3 for continuous operation and such operation can therefore not be achieved. The $R^{\text{avg}}_{\text{pixel}}$ values on the other hand are below the 1 Mhits/s and indicate that the pixels will be able to process the events, but the readout in the Timepix3 is not able to push out the data in time. We can take advantage of a property in the Timepix3 to overcome this limitation. Each pixel in the detector can be considered as a memory cell that holds information from a measured event until it has been read out and if this is done before the ToA counter overflows after 409.6 µs the timing information is intact. We can therefore foresee a mode of operation where the Timepix3 collects events for less than 409.6 µs and then reads them out over a longer period of time. This 409.6 µs is about 100 to 200 revolutions in the machine, based on the revolution periods from table 3.4. We can therefore use the PS-BGI instrument in such a mode to study injection oscillations for example, which will be shown later in section 5.4.

The values presented in this section are estimates of the expected ionisation

electron yield and hit rates. In the formula for the electron yield (eq. (2.26)) we see that the ionisation cross section, the static pressure and the number of particles in the beam all affect the yield equally. We saw in table 3.5 that the ionisation cross section can vary by a factor 5 for different gases and we also discussed the variation in static pressure in the CERN PS that can vary by an order of magnitude. With gas injection these last two parameters can be controlled more precisely and the sensor bias voltage and threshold settings in the Timepix3 can also be changed to lower or increase the detection efficiency, within the limits of the detector. The variation in detector hit rates due to the bunched nature of the beam also presents challenges in predicting the response of the detector. Two different beams with two different particles have been presented in this section without any definition of how the measurement is done, for example if the goal is to measure an integrated beam profile over multiple turns or bunch-by-bunch. It is therefore suggested to use the equations and examples above to calculate the expected ionisation yield and hit rates on a case by case basis, where additional information on the measurement setup can improve the predictions.

4 | Beam Profile Reconstruction

This chapter introduces the methods used for reconstructing the beam profile from beam gas ionisation electrons detected with the Timepix3 hybrid pixel detectors. Many different particles yield readout events in the Timepix3 detectors including those from ionisation electrons as described in section 3.4. It is therefore important to process the events in such a way that ionisation electron events are distinguished from other background events. For the reconstruction, we hypothesise that the transverse distribution of these ionisation electrons can be used to measure the beam profile and from the profile infer the beam size.

The detectors can be simulated to gain insight into how events from the pixel detectors can be analysed to produce a beam profile and size measurement. In section 4.1 such a simulation is described where events are created which mimics the response of the pixel detectors to a beam with a Gaussian distribution. Simulation data will be compared throughout this chapter with measurements acquired for a variety of acceleration cycles in the PS.

Events from detected particles which are not ionisation electrons have to be identified and removed. In section 4.2.2 a method using clustering is described that separates the ionisation electron events from the background. Background events form characteristic clusters in time and space, while the ionisation electron events are spread out as single pixel events. The ToT information in each event is used in conjunction with the clustering to further improve the identification and separation.

A pixel image with 260 000 pixels can now be constructed from the ionisation electron events. Some of these pixels have a different response compared to the rest due to debris on the sensor and the RF shield above the detectors. They therefore need to be identified and removed which will be detailed in section 4.2.3.

After removing the background events and masking pixels we are left with a pixel image which can be used to construct a beam profile (section 4.3). A significant

advantage of using hybrid pixel detectors for an ionisation beam profile monitor is that it enables counting of individual ionisation electrons. This allows for a rigorous mathematical treatment of the data based on Poission statistics which will be shown in section 4.3.1. The output of this is a beam profile measurement as will be seen in section 4.3.2, where profiles from simulation and measurement are presented.

The beam profile is an estimation, based on the measurement, of the transverse beam distribution which can be processed further to produce a beam size measurement. This describes the width of the distribution and two different methods are introduced in this chapter. The first method, which involves fitting a Gaussian (normal) distribution to the beam profile, is described in section 4.3.3 and the second method, based on calculating a RMS value, is described in section 4.3.4.

Lastly, Monte Carlo simulations will be shown in section 4.4 where the limitations on precision (section 4.4.2) and accuracy (section 4.4.3) are studied. A final simulation is presented in section 4.4.4 where the effect of the background events on the beam size is studied.

4.1 Simulation of Detector Response

We present the details of a simulation that allows for full control of the beam profile distribution and how it is sampled in order to explore detector response to a known source. Known properties of the simulation such as beam size and position can be used to compare with equivalent properties of a measurement using the detector.

Figure 4.1 illustrates the basic idea of the simulation. A transverse beam profile is assumed to follow a Gaussian distribution, with a fixed position and size relative to the detector. The creation of an ionisation electron is simulated by randomly sampling from this distribution. Samples that are above a gap between the pixel detectors, or outside the detector region are removed. The samples are then binned into pixels of size 55 µm by 55 µm. As a final step, pixels which should not be included

in the analysis can be masked. This will be detailed later in section 4.2.3. The binned

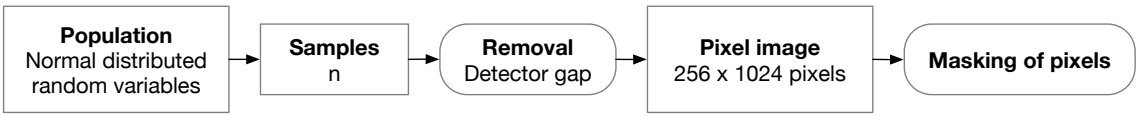

Figure 4.1: Flowchart for simulation of detector response.

samples are equivalent to events recorded by the Timepix3 pixel detectors during a measurement with the instrument. This facilitates the use of the same beam profile reconstruction methods in both cases. A pixel image constructed from the simulated binned samples can be seen on the left in fig. 4.2. In the middle of the figure the masked pixels are marked and on the right they have been removed. The right pixel image is equivalent to a measured pixel image as we will see later in this chapter.

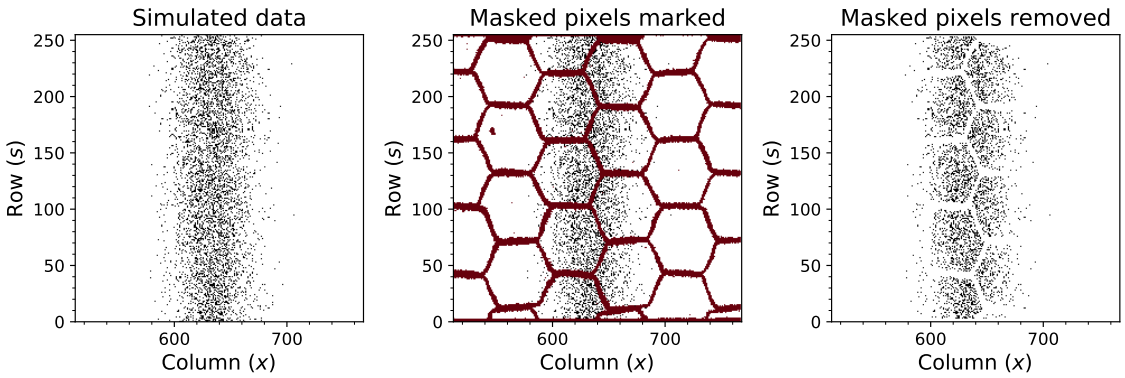

Figure 4.2: Pixel image from simulated data (left) with masked pixels marked (middle) and removed (right).

A further improvement to the simulation would be to add beam loss and background events which is present in measurements. In section 4.2.2 a procedure is shown how to extract the background events from measurements. These events could be added to the simulation at random to ensure a more accurate representation of the detector response. The main use of the simulated data in this thesis has been to evaluate the response of the reconstruction of the beam profile and to understand the underlying statistical limitations as will be seen later in section 4.4.

4.2 Data Treatment

The events recorded by the Timepix3 pixel detector that are processed by the read-out contains spatial and temporal information as discussed in sections 2.4.2 and 3.2. For the beam profile reconstruction the following information is used:

- Detector number (1 to 4)
- Column number (0 to 255)
- Row number (0 to 255)
- Time-of-arrival (ToA) timestamp with a 1.5625 ns resolution
- Time-over-threshold (ToT) with a 25 ns resolution

All the events recorded by the instrument during a measurement are stored in a data set. A measurement can be taken for a full PS cycle, which can be a couple of seconds long, and the first step of the data analysis is to split all the events into multiple integration time windows (section 4.2.1). The events within a time window include both background particles and ionisation electrons and they need to be separated (section 4.2.2). In addition, the presence of the metal RF shield and debris on the sensors result in a non-uniform transmission efficiency of ionisation electrons from the point of creation to the detector plane. Pixels that are located under the shield or debris have to be masked out to avoid distorting the beam profile (section 4.2.3).

4.2.1 Integration Time

The integration time window needs to be optimised depending on the beam parameters (such as intensity) and the period of time within the cycle of interest (e.g. turn-by-turn at injection compared to 100 profile spread over the full cycle). In fig. 4.3 three pixel images with increasing integration time are shown. To get better counting statistics a longer time should be used, but if the beam is not stable during this period the measured beam profile could be disturbed. The smallest integration

time for beam profile measurements is the length of a single bunch of the beam in one turn, which, for example, is 25 ns at the end of the cycle in the PS. Obtaining a profile in this window is limited by the number of ionisation electrons created during the beam gas interaction, as discussed in section 3.5.2. A compromise has to be made between a more precise measurement and a higher beam profile sample rate. In some cases there might be a hard requirement for single turn measurements because the beam is changing from one turn to the next. Such an example will be shown in the next chapter where injection mismatch is studied (section 5.4).

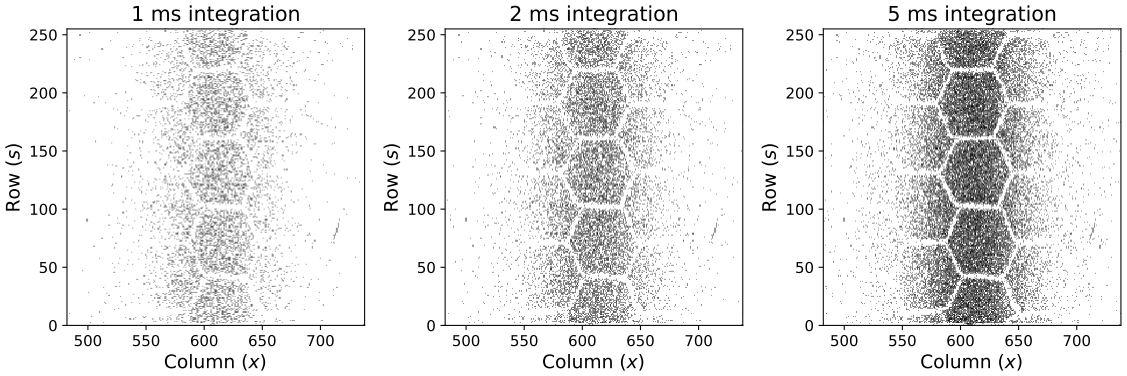

Figure 4.3: Pixel images from measured data with different integration time.

4.2.2 Removal of Background Signal

We can identify the ionisation electrons detected by the Timepix3 hybrid pixel detector based on cluster size and ToT value, as discussed in section 3.4. The ionisation electrons from the beam gas ionisation mostly form single pixel clusters with ToT values between 1 and 30. This combination of low ToT value and small size can be used to separate the ionisation electrons from the rest of the detected particles, which we will call the background particles.

Identification and Removal Using Clustering

The background particles trigger multiple pixels in the form of clusters and the events from these pixels are grouped close together in time, with a maximum time

difference of about 100 ns. This information can be used to first cluster events in time and then cluster them in the 2D space. The combined time and space information can then be used to remove the background particle events from the recorded data. In fig. 4.4 the ToT values of events are shown for a 1 ms measurement with beam. Between column 300 and 500 there are several single pixel events with low ToT values, which is indicative of ionisation electrons.

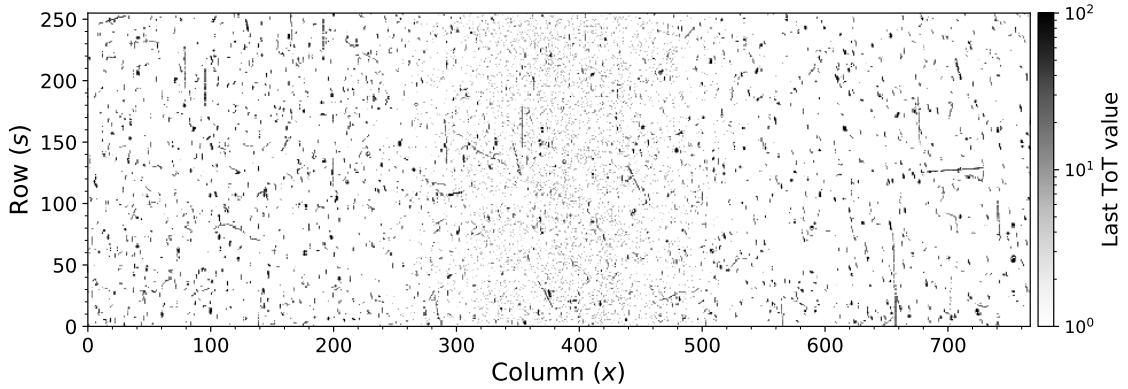

Figure 4.4: Pixel image from 1 ms of measured data with background particles and ionisation electrons.

We will first use the DBSCAN algorithm, introduced in section 3.4.2, to cluster events in time by using the pixel event timestamps, where the DBSCAN parameter ϵ defines the distance in time between two events. If we take a 20 µs time window from the same measured data used for fig. 4.4 there is a distinct time structure of the timestamps of the events as can be seen in fig. 4.5. The DBSCAN algorithm identifies the nine clusters of events in time, where each cluster is marked with a different colour. The distance between adjacent clusters is around 2.2 µs, which is

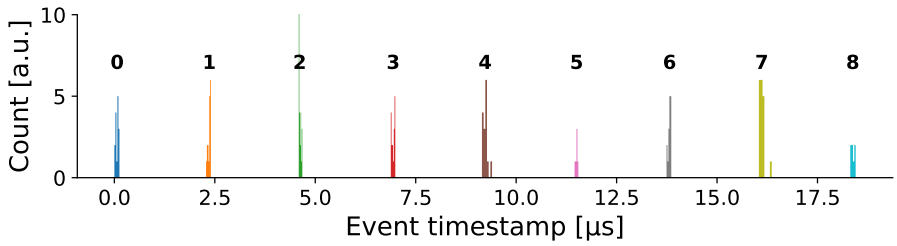

Figure 4.5: Histograms of event timestamp for a 20 µs time window.

consistent with the revolution frequency of the PS accelerator, such that each cluster contains events from a single turn in the accelerator. We will use this method of timestamp clustering in the next chapter to study the time structure of the beam (section 5.3).

If we now run the time clustering for the full 1 ms of the measured data in fig. 4.4, the algorithm identifies 400 time clusters, where the first 9 are the ones shown in fig. 4.5. For each of the time clusters we can run a 2D spatial clustering which, in combination with the time cluster information, give us unique clusters for all the measured data. In total, the algorithm identifies 2000 unique clusters. A majority of the events from low energy ionisation electrons will be identified as "noise" by the DBSCAN algorithm because of their tendency to form single pixel events. These ionisation electron events are shown in fig. 4.6, which can be compared to fig. 4.4. We expect to see a concentration of ionisation electrons somewhere on the detector area corresponding to the position of the beam above, which we do between column 300 and 500 in fig. 4.6. There are additional events outside this region, which will be discussed later in section 4.3 when the beam size is calculated.

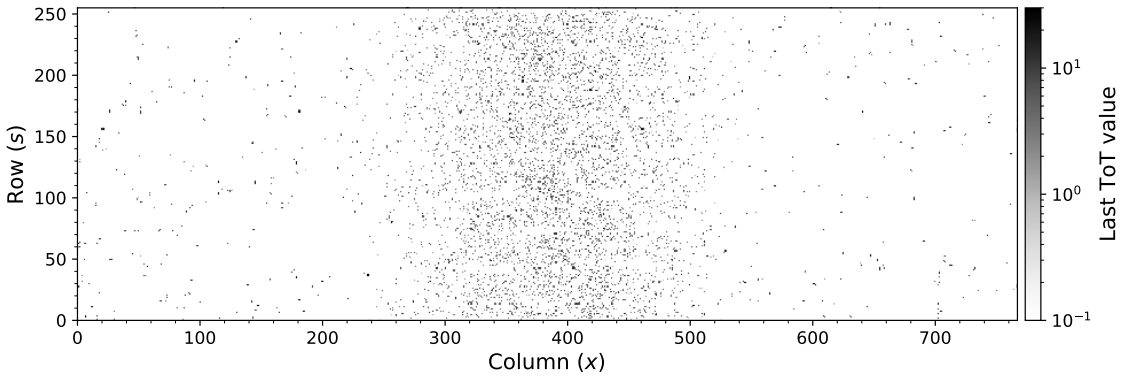

Figure 4.6: Pixel image from 1 ms of measured data with only ionisation electrons.

The background particle events can be separately extracted as shown in fig. 4.7 using the same measured data as before. Due to the high ToT value and shape of these clusters they are inconsistent with ionisation electrons and instead originate mostly from beam loss. As mentioned in section 4.1, these clusters can be used

to improve the realism of the detector simulation. The sum of ionisation electron events from fig. 4.6 and background events from fig. 4.7 is equivalent to fig. 4.4.

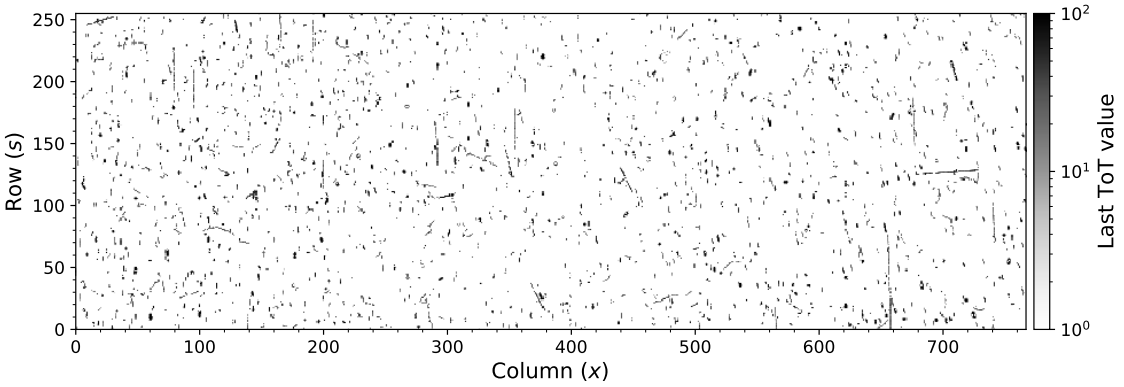

Figure 4.7: Pixel image from 1 ms of measured data with only background particles, mostly coming from beam loss.

Identification and Removal Using ToT Cut

The clustering algorithm above requires computational efforts and time, which for example in a real time application could be critical. An alternative approach instead uses the ToT information alone by removing all high ToT events. Since this cut method only requires information from a single event it has a performance advantage over clustering and could run in real time. In fig. 4.8 the distribution of ToT values is shown for all the events from the measured data used in fig. 4.4 with a distinct peak at a ToT value of 10. If we place a cut on the ToT value at 15 and only keep

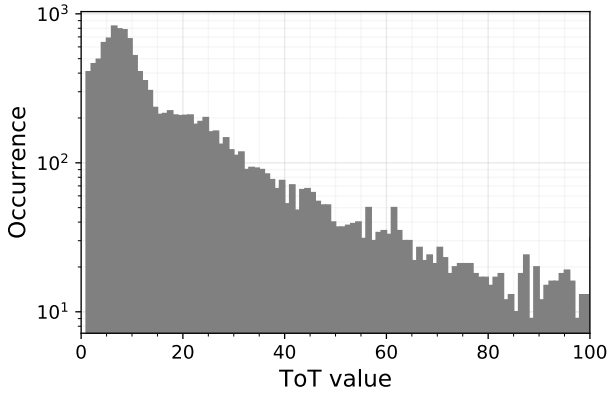

Figure 4.8: Distribution of ToT values in the measured data from fig. 4.4.

events with a value lower than this, most of the ionisation electron events are kept while removing a large portion of the background.

The result is shown in fig. 4.9 where we can see that the ionisation electron signal is more prominent compared to the original data in fig. 4.4 but there is still background signal remaining. There is more background signal compared to the result of the clustering method seen in fig. 4.6, and in total there are 38 % more events. A background particle triggers several pixels and the total deposited energy by the particle in the sensor is distributed among all these pixels. The removal of background signal is therefore more accurate when taking into account multiple events, as is done with the clustering method, compared to only single events with the ToT cut method.

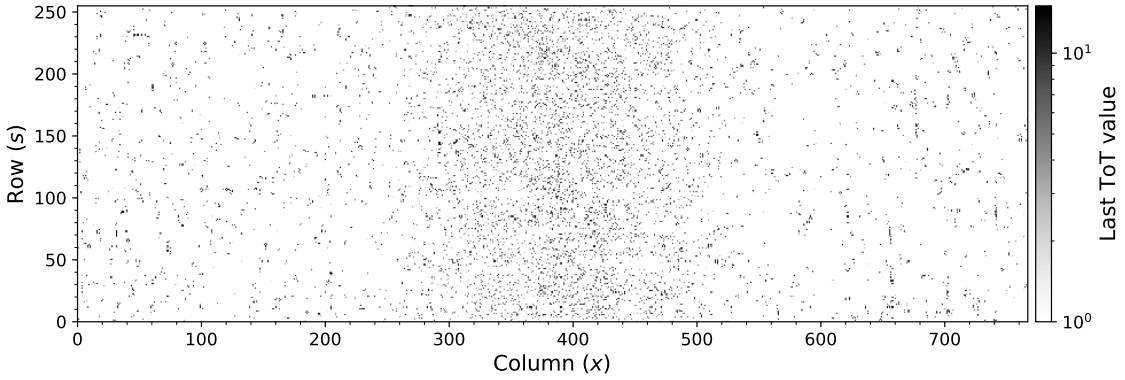

Figure 4.9: Pixel image from fig. 4.4 with only low ToT value events present.

4.2.3 Masking of Pixels

The pixel image constructed from the detector events thus far is not ready to be used for beam profile reconstruction. As evident in fig. 4.3, a shadow of the honeycomb shaped RF shield is visible. This is due to the low energy of the ionisation electrons (10 keV) which will be absorbed within a few µm of the metal shield. The RF shield is about 5 mm thick and is made from stainless steel (see section 3.1.3).

Under the assumption that the Timepix3 detectors have been equalized, detailed in section 3.3.2, all pixels have the same response to a fixed input charge. However,

the transmission efficiency of ionisation electrons from the point of creation to the detector plane is not the same for all pixels due to the presence of the RF shield. This difference in transmission efficiency leads to a lower average event rate for pixels under the shield and due to the irregular shape of the shield, some columns will have more pixels under the shield than others. This leads to a non-uniform efficiency over the detector area and distorts the beam profile measurement in the transverse direction. It is therefore paramount to remove these low transmission efficiency regions of the detector by identifying the affected pixels and masking them out. Additionally, pieces of dust or other debris on the sensors might also lower the efficiency of the pixels under it and these pixels should therefore also be masked. We can categorize pixels that should be masked out into the following three categories:

1. Pixels under the honeycomb shaped RF shield.
2. Pixels under debris on the sensor.
3. Faulty pixels that do not generate any signals or has been disabled.

For simplicity, events coming from a masked pixels will be removed. In fig. 4.10 examples of these types of pixels are highlighted for measured data with a majority of them belonging to the first category.

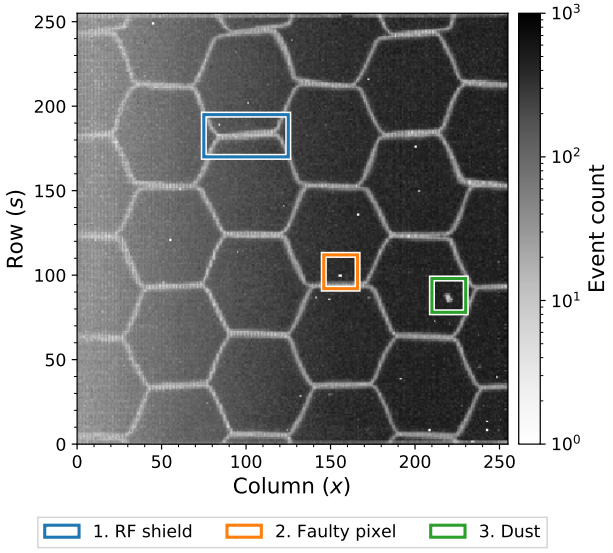

Figure 4.10: Pixel image from measured data with examples of different pixels that should be masked out.

Identification and Separation Using Median

As can be seen in fig. 4.10, the highlighted pixels have a lower event count compared to the neighbouring pixels. We will use this property in order to identify the pixels that will be masked out. The data in fig. 4.10 is from a measurement with beam and we therefore expect that the distribution of ionisation electron count for a single column, aligned to the beam propagation direction, to follow a Poisson distribution. The reasoning behind this statement will be motivated later in section 4.3.1. In fig. 4.11 the distribution (right) of event counts in the highlighted column (left) is shown. It is not a Poisson distribution, but instead there is a population centered around an event count of 250 and another population at lower count values, which are the candidates to be masked out.

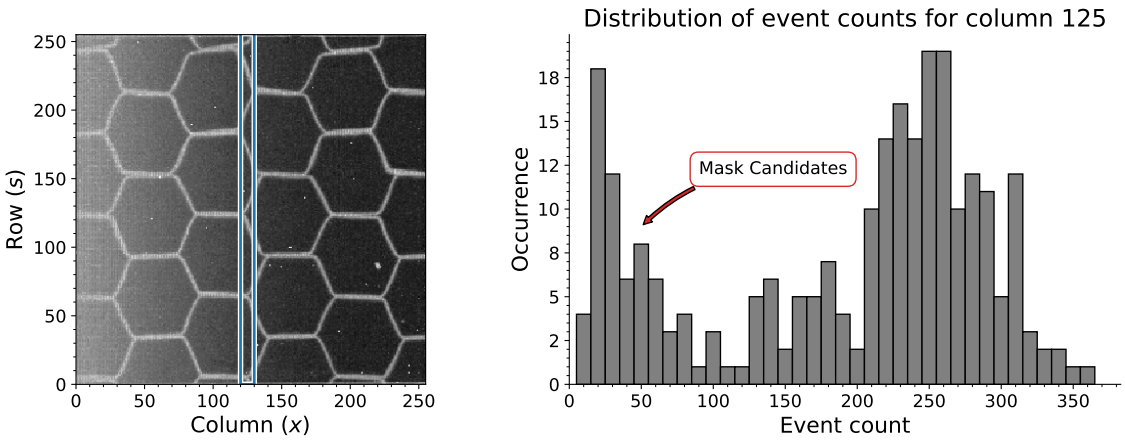

Figure 4.11: The distribution of count values (right) for a highlighted column in the pixel image (left) from measured data.

Ideally we would fit a model to the complete distribution. The contribution to the model from the nominal pixels is a Poisson with an average rate that is proportional to rate at which ionisation electrons are created. Due to the different detection efficiencies of the pixels, the rate for the mask candidate pixels varies and would require a specific model for each pixel and high statistics to resolve. Another approach is to consider the mask candidate pixels as outliers compared to the pixels which have a nominal efficiency. The median $\tilde{x}$ is an appropriate statistical property

to use for this purpose. This gives an estimate of the centre for the nominal pixel population due to them being a majority of all pixels. As an estimate for the width of the population we can use the median absolute deviation (MAD) which is defined as

$$\text{MAD} = \text{median}(|x_i - \tilde{x}|). \tag{4.1}$$

In fig. 4.12 the median together with two limits with MAD included is shown. For this specific distribution a limit on the event count values could be set at $(\tilde{x} - 2 \cdot \text{MAD})$ which will identify the mask candidates while leaving most of the desired signal from the nominal pixels.

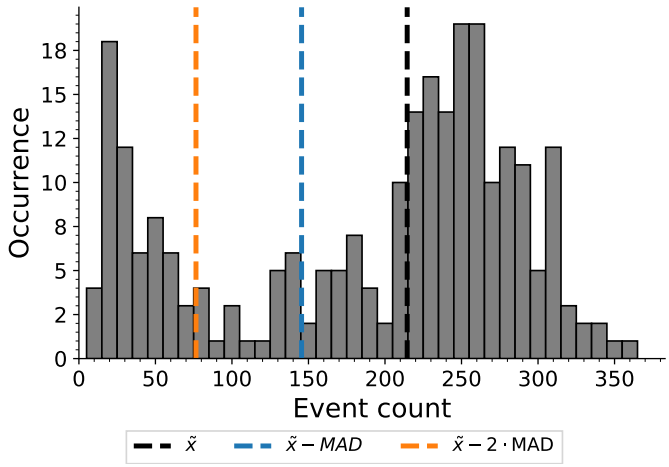

Figure 4.12: Distribution of count values with median and MAD limits.

If we apply this limit to all pixels in the measured data, it can identify most of the mask candidates as seen in fig. 4.13. This data is from 2018 when the fourth detector was broken (see section 3.1.3), hence the lack of data on the right side of the figure. On the left and right sides the identification is not optimal due to lack of data in these region from limitations on how far the beam could be moved. With a reduction in data it becomes more difficult to distinguish between the two populations in fig. 4.12.

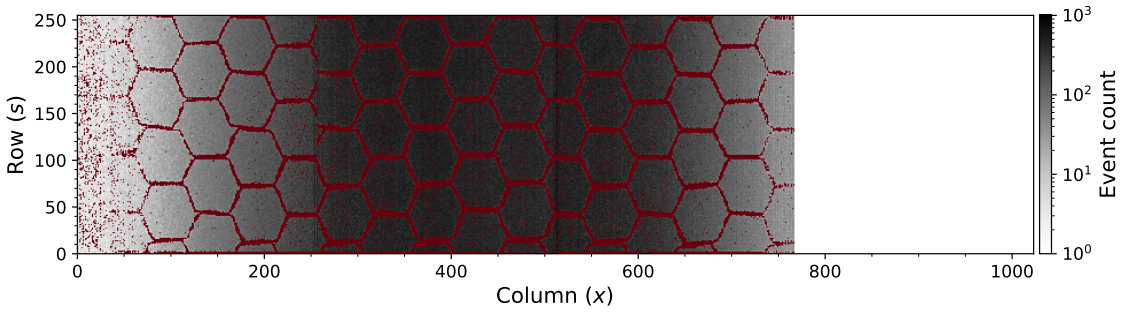

Figure 4.13: Identification of the mask candidates (red) from measured data using an event count cut of $\tilde{x} - 2 \cdot MAD$.

Source of Ionisation Electrons For Identification of Mask Candidates

The data that is used to identify the mask candidate pixels should ideally cover all four detectors evenly in all directions. As demonstrated, ionisation electrons from beam gas ionisation in the accelerator can be used, which has the advantage that the instrument is mechanically stable and protected by the vacuum from new debris. If the source of ionisation electrons is located in a test laboratory, the instrument has to be moved, during which the RF shield might shift or debris and dust may land on the sensor. The disadvantages of using the accelerator beam as the source is the necessary access to a beam, which might not be possible all the time, and limitations in the detector area that can be intentionally exposed evenly to ionisation electrons. Due to the limitations in the beam aperture more high energy beam loss particles will be created when the beam is moved closer to the beam pipe which reduces the quality of the test data due to these particles passing through the RF shield and any debris on the sensor.

Mask Stability Over Time

It is important to understand how stable the mask is over time, due to more dust collecting on the sensor surface or if the dust that is already there moves around. The RF shield is held in place mechanically but one could argue that over time and with heating and cooling that the shield might move or change shape with respect to

the pixel detectors. Ideally, an experiment to test this potential issue would use the exact same source of ionisation electrons and collect data over a long time period.

From the measured data in 2018 using the Mk II instrument there are only a handful of data sets which can be used as a source for identifying mask candidates due to limited beam time where the beam could be moved over all detectors. Further, the result is not exactly reproduced for these data sets due to variations in the measurement conditions, such as vacuum pressure. With that in mind, in fig. 4.14 the masks extracted from two proton beam measurements taken five months apart are shown. On the right in the figure, the difference between the first data source and second data source is shown. Both data sets consists of 24 beam cycles recorded during 1 h, where the 2018-05-25 data had an average of 21.8 million events per cycle and the 2018-11-06 data an average of 1.7 million events per cycle. To avoid any bias due to this difference in event count, 20 million events were chosen at random from all cycles for both data sets. For the RF shield, the centre of each horizontal and diagonal line is stable as indicated by the lack of dark pixels in these regions in the difference plot. This is the important information as it indicates a stable position of the RF shield.

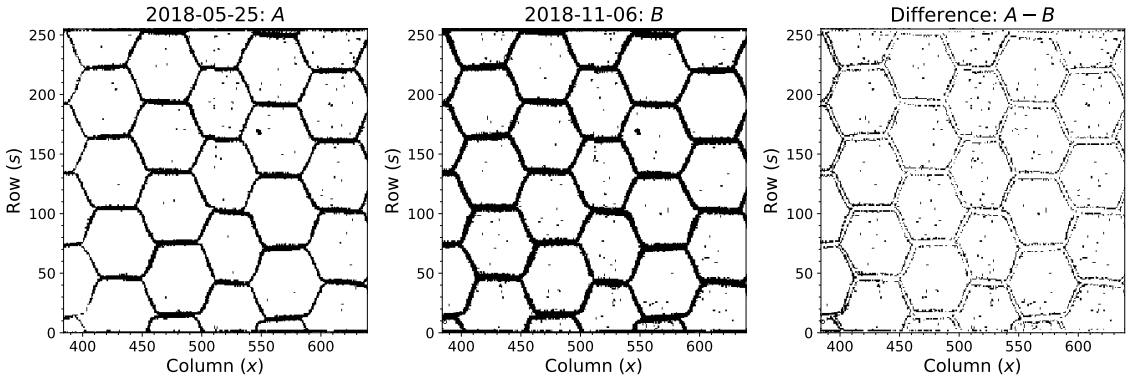

Figure 4.14: Masks from data taken using an accelerator beam on 2018-05-25 (left) and another beam on 2018-11-06 (middle). The difference between them is shown on the right.

We can also see in fig. 4.14 that the shadow of the honeycomb shield is broader for the 2018-11-06 data set. This is due to a difference in the field cage potential

between the two, where it was −10 kV on 2018-05-25 and −20 kV on 2018-11-06. The lower potential in the former translates to a lower energy of the ionisation electrons as they are accelerated down towards the shield and leads to a higher number of electrons blocked by the shield. For the discussion of stability over time, this difference in width of the mask is not of a major concern. It would only be problematic if the shield degrades over time which could result in a thinning of the walls in the shield.

The dust particle highlighted in fig. 4.10 at around column 550, in the green coloured box, can also be seen in fig. 4.14 and it stays the same in both cases. There is also no indication of new dust getting stuck on the sensor.

Based on these findings, we can assume that a mask will stay constant during the operation of the instrument and can otherwise be calibrated in place with an accelerator beam as the source of ionisation electrons. With any large modifications, such as disassembly of the instrument or installation of new detectors, a new mask has to be measured which can be considered part of the commissioning of a new or updated instrument.

Final Mask

Based on the discussion thus far we can construct a mask by combining the result of multiple measurements with an accelerator beam as the source of ionisation electrons. We will use data recorded on six different occasions, each composed of 10 to 48 beam cycles. For each measurement, the mask candidates are identified using the median method described earlier and the resulting mask can be seen at the top in fig. 4.15. This mask finds the outline of the RF shield and it is identifying many pixels in the open apertures as well. Some of these pixels should be masked, such as the piece of dust in column 550 and pixels that are disabled, but the combination of all the data sets could have resulted in false positives. A set of disabled pixels is identified during the equalization of the detectors (see section 3.3.2) which we

can use directly in the mask. Furthermore, we can make the assumption that dust particles have a diameter greater than the size of one pixel and therefore remove all single pixel clusters from the mask unless they are part of the disabled set. The final mask after these operations can be seen at the bottom in fig. 4.15 and will be used in the beam profile reconstruction process that follows.

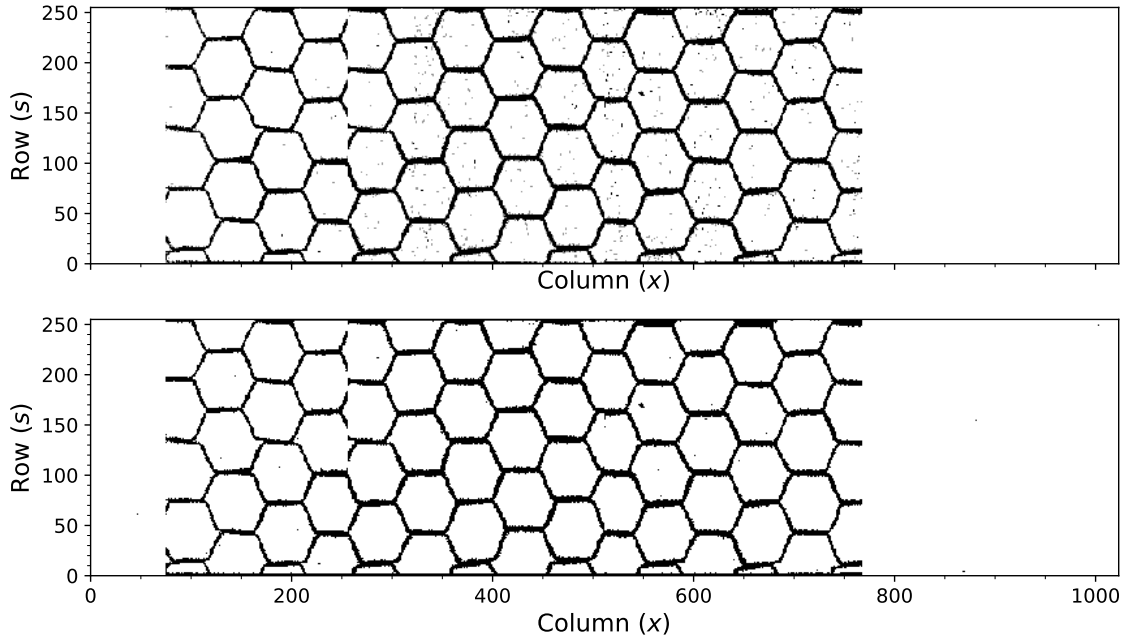

Figure 4.15: Mask after combining data from several data sets (top). Mask after identifying disabled pixels and removing dust particles smaller than the size of one pixel (bottom).

Applying Mask to Pixel Image

At the top of fig. 4.16 a pixel image is shown for a 5 ms integration window which has not been masked yet. Using the mask from before we can identify the pixels and visualize them using a different colour as seen in the middle of fig. 4.16. The next step is to remove the masked pixels which will create a non-uniform efficiency for ionisation electron detection across columns, where a column with fewer pixels removed will have a higher efficiency and vice versa. To ensure a constant detection efficiency for all columns, a random set of unmasked pixels is chosen for each column

such that all columns have the same number of pixels as the column with the lowest number of unmasked pixels. All other pixels that have not been selected will be discarded. The resulting pixel image is shown at the bottom of fig. 4.16.

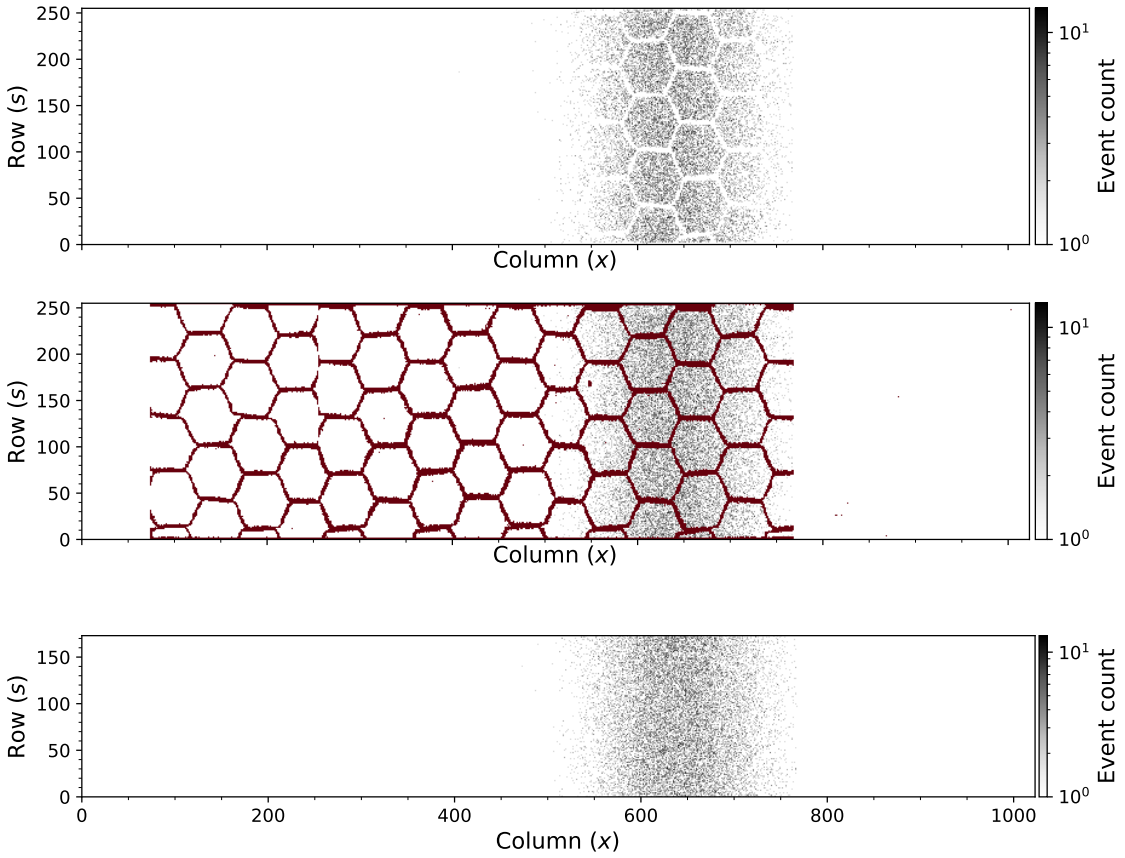

Figure 4.16: Example of pixel image without masking of pixels (top), with mask applied and visible (middle) and with only unmasked pixels (bottom).

4.3 Beam Profile and Size Measurement

We now have a pixel image containing predominantly ionisation electron events for a specific time window. The next step is how to interpret these events and make a meaningful beam profile and beam size measurement from them. A single particle hitting the sensor of the Timepix3 hybrid pixel detector triggers an event, which suggests we can treat events from the pixel detector as a counting experiment. The detected ionisation electrons correspond to samples from the transverse beam

distribution, i.e. the original population which is what we want to measure. The counting approach suggests that we can apply Poisson statistics to analyse all the events in a methodical way. The motivation and applicability behind this will be detailed in section 4.3.1.

From the counted events we can construct a beam profile which contains information about the shape of the transverse beam distribution in a given time window (section 4.3.2). We could stop at this point and use the profile as the final measurement, but it is more common to take it one step further and extract a beam size from the profile. The beam size measurement can also be linked to the emittance of the beam, following the procedure outlined in section 2.2. The beam size can be measured either by fitting a model to the beam profile, such as a Gaussian model which we do in section 4.3.3, or by directly calculating the statistical properties of the samples in section 4.3.4. For the fitting it is important to have a good understanding of potential mismatch between the model and the actual beam distribution. We will also consider what impact noise could have on the beam size measurement.

4.3.1 Counting Ionisation Electrons

The Poisson distribution describes the probability of a given number of events k occurring in a fixed interval when the events occur with a known constant rate λ independent of other events. The fixed interval could be time, distance, area or volume. This can be used to describe processes where a quanta is counted, which is the case for this application of an ionisation profile monitor using hybrid pixel detectors. Individual ionisation electrons are counted over the area of the detector as described earlier in section 3.4. The probability mass function (PMF) for the Poisson distribution is defined as

$$f(k \mid \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}. \tag{4.2}$$

The event count of an individual pixel at column c and row r in the detector should therefore, for a fixed beam profile, follow a Poisson distribution with an average event rate given by

$$\lambda_{\text{pixel,c,r}}. \tag{4.3}$$

Given a fixed beam profile, we hypothesise that pixels in the same column have the same average event rate. This is because they are located at the same transverse position relative to the beam and are therefore exposed to the same flux of ionisation electrons, under the assumption that the columns are aligned to the beam.

Counting Using Measured Data

Measuring multiple accelerator beams under the same settings will present difficulties due to fluctuations in the equipment that is controlling the beam. This makes it impractical to measure beams with the exact same settings multiple times. From a single beam measurement we can get an indication that our hypothesis is correct. On the left in fig. 4.17 a pixel image is shown for a measured data set. The masked pixels have been removed using the method described in section 4.2.3. Due to the removal, the pixel image contains 184 rows of pixels instead of the original 256. A histogram

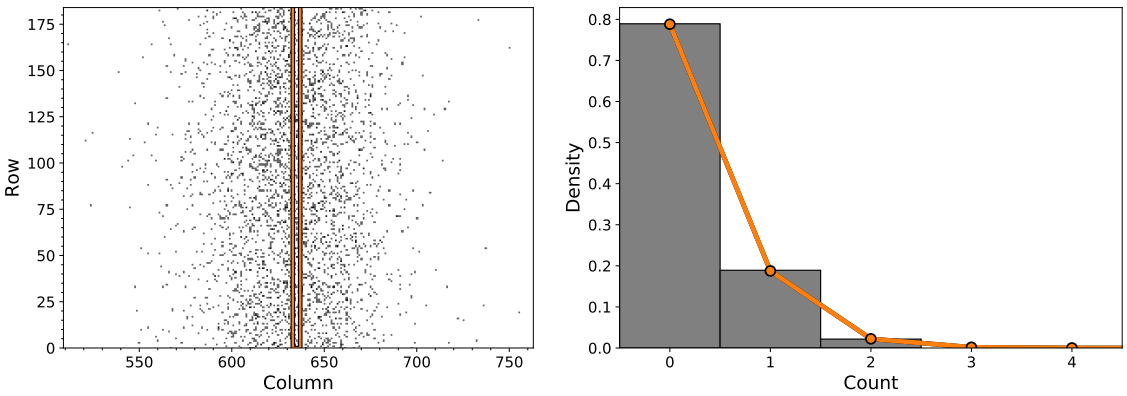

Figure 4.17: The distributions of event counts for the highlighted column on the left is calculated from one measured beam profile. A Poisson function is fitted to the distribution on the right and gives an event rate of $\lambda_{column,635} = 0.238 \pm 0.001$.

of the event count for all the pixels in the highlighted column is constructed, as seen

on the right in fig. 4.17. More than 75 % of the pixels have a count of 0, about 20 % have a count of 1 and there are also a few with a count of 2. A function following the Poisson distribution is fitted to the histogram and the average event rate for the highlighted column is estimated by the fit to be $\lambda_{\text{column},635} = 0.238$. If we multiply this rate by the number of pixels in the column (184) we get $\tilde{k}_{635} = 44$, which is the estimated total event count for this column. The actual total event count for the column is $k_{635} = 43$, which is in agreement with the estimate.

Even though this is only one example of measured data it indicates that treating events from the hybrid pixel detectors, corresponding to individual ionisation electrons, as a counting experiment is reasonable. We will use this property of the events in the following steps of the analysis to first create a beam profile and later to calculate a beam size. First we will discuss standard errors in the context of Poisson statistics, which we will use in our beam profile and size measurements.

Calculating Standard Error of The Mean For Poisson

The standard error of the mean can be expressed as

$$\sigma_\mu = \frac{\sigma}{\sqrt{n}}, \tag{4.4}$$

where n is the number of observed samples and σ is the standard deviation of the population. For a Poisson distribution the mean is equal to the variance and we can therefore express the standard deviation of a Poisson distribution as

$$\sigma = \sqrt{\lambda}. \tag{4.5}$$

For the beam profile, the best estimate of the mean λ for a specific column is the sum of counts k_c and we only have one observed sample for each column, i.e. $n = 1$.

The standard error of the mean can therefore be expressed as

$$\sigma_\mu = \frac{\sqrt{\lambda}}{\sqrt{n}} = \sqrt{\frac{\lambda}{n}} \approx \sqrt{\frac{k_c}{1}} = \sqrt{k_c}. \tag{4.6}$$

We will use this expression to calculate error bars for the beam profile measurement.

4.3.2 Beam Profile Measurement

The sum of counts k_c can be calculated for each column c, where each column will have an average event rate $\lambda_{\text{column,c}}$. Each column can also be assigned a physical transverse position x_c, taking into account any gap between detectors. The quantities x_c and k_c for all columns c define the beam profile. In fig. 4.18 a beam profile is shown from a simulation on the left and from a measurement on the right. The error bars are calculated using eq. (4.6). We know the properties of the population from which the simulated data was sampled. For the measured data we do not know the true values of the transverse beam distribution (this is the reason to have a beam profile monitor, to estimate the population of the particles in the beam). We can conclude that the measured data, in this example, is similar to the simulated data, but a quantifiable size measurement is needed to compare two profiles in a general case. One approach that involves fitting a model to the profile will be detailed in the next section.

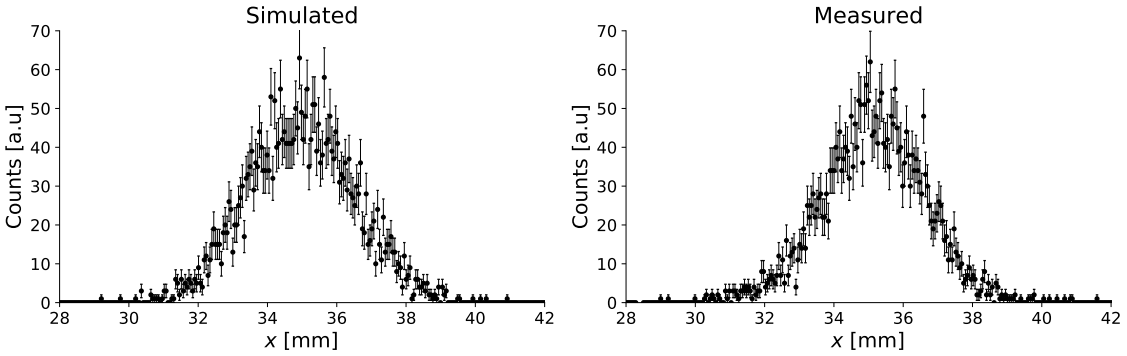

Figure 4.18: Beam profiles with error bars from a simulation with an assumed beam size of $\sigma_{\text{true}} = 1.55\,\text{mm}$ (left) and from a measurement (right).

An important detail that can be seen outside the tails of the measured beam profile in fig. 4.18 is the addition of noise. The beam profile from the simulation is noiseless. For the measured data there are additional events detected that cannot be identified and removed by the method described in section 4.2.2 due to their similarities with ionisation electron events. We will discuss how this additional noise affects the measurements and how to compensate for it throughout this chapter.

4.3.3 Beam Size Measurement Using Gaussian Distribution

The beam profile can be modeled using a function. A normal (Gaussian) distribution can for example be used as the model function and the beam size can be specified as the standard deviation σ of the distribution. In reality, the distribution of particles in the beam can be the convolution of multiple models which, when projected onto the transverse planes, require a more complex function than a Gaussian. To get an accurate measurement of the beam size care must be taken when defining the model for the profile.

Definition of the Gaussian Distribution

In general, the function that is used to model the beam profile can be written as $f_X(x \mid \theta)$, where x are the observed values and θ are the parameters that describe the model. In the case of a Gaussian distribution the parameters are the mean μ and the standard deviation σ and they are confined by the dimensions of the detector, i.e. a mean value greater than 56 mm would be invalid since it is greater than the width of the detector in the transverse direction. To model the noise in the measured beam profile, seen in fig. 4.18, an offset parameter a is added to the Gaussian distribution. The final function that we will use to model the beam profile is therefore:

$$f_X(x \mid \theta) = f_X(x \mid I, \mu, \sigma^2, a) = \frac{I}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} + a, \tag{4.7}$$

where I is a scaling factor that is related to the number of ionisation electron events. The next step is to find the estimates of μ and σ for a given beam profile which will give the best agreement between the observed values and the modeled Gaussian distribution.

Minimisation of Likelihood Function

The likelihood function is a function of the parameters θ given the observed values x. The parameters are μ, σ and a in this case for the Gaussian distribution with an offset as defined in eq. (4.7). We will use the likelihood function to fit the model to the observed data by minimising the likelihood function. The parameter values we get after the minimisation are the best estimates for the model but does not necessarily mean that the model fits the data well. A goodness of fit test is therefore required to evaluate the fit validity, which will be detailed after this discussion on minimising the likelihood function.

If we sample n independent identically distributed (i.i.d.) random variables $X = (X_1, ..., X_n)$ we get the observed values $x = (x_1, \ldots, x_n)$. The probability of these n events together can be described using the joint probability distribution

$$\prod_{i=1}^{n} f_{X_i}(x_i \mid \theta) = L(\theta \mid x), \tag{4.8}$$

where f_{X_i} is the model for a specific observation x_i and L is the likelihood function. As described earlier in section 4.3.1, the count k_c in a column at the transverse position x_c can be modeled as a Poisson distribution, after discarding masked pixels. To construct the full beam profile the individual columns have to be combined. The probability of observing $\{k_c\}$ given the parameters $\{\lambda_c\}$ for all columns c can therefore be described as the joint probability

$$P(\{k_c\} \mid \{\lambda_c\}) = \prod_{c=1}^{m} P_c = \prod_{c=1}^{m} \frac{\lambda_c^{k_c} e^{-\lambda_c}}{k_c!}, \tag{4.9}$$

where m is the number of columns and P_c is the Poisson probability for observing k_c counts with expectation of λ_c counts (eq. (4.2)). Since the likelihood function is the joint probability for all observed values we can use eq. (4.9) to write the likelihood function as

$$L(\lambda_c \mid k_c) = \prod_{c=1}^{m} \frac{\lambda_c^{k_c} e^{-\lambda_c}}{k_c!}. \tag{4.10}$$

This product can make the equation difficult to compute. By taking the natural logarithm of the likelihood function we get the log-likelihood function which is more convenient to use:

$$\mathcal{L}(\lambda_c \mid k_c) = \ln(L(\lambda_c \mid k_c)) = \sum_{c=1}^{m} \Big(k_c \ln(\lambda_c) - \lambda_c - \ln(k_c!) \Big). \tag{4.11}$$

The parameters λ_c are defined by the chosen model. This could in theory be any type of model, but in this case we will use the Gaussian with an offset as defined in eq. (4.7). The Gaussian model is often a good approximation for transverse beam profiles but could also be extended or replaced if it does not fit well with the transverse distribution of the beam. Therefore, for a given transverse position x_c, eq. (4.7) can be used to write

$$\lambda_c = \frac{I}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_c-\mu)^2}{2\sigma^2}} + a. \tag{4.12}$$

Combining eq. (4.11) and eq. (4.12) results in the log-likelihood function

$$\mathcal{L}(\mu, \sigma, a, I \mid x_c, k_c) = \sum_{c=1}^{m} \Big(k_c \ln(\frac{I}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_c-\mu)^2}{2\sigma^2}} + a) - \frac{I}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_c-\mu)^2}{2\sigma^2}} - a - \ln(k_c!) \Big). \tag{4.13}$$

The parameter I can be estimated as the total number of ionisation electron events under the assumption that $a << I$, which follows from the low amount of noise events shown earlier in section 4.3.2. The values x_c and k_c are known from the observed data, while μ, σ and a are the unknown parameters we want to esti-

mate. Maximising the log-likelihood function gives the estimated parameters for the selected model given the observed data. This is called the maximum likelihood estimation (MLE) method. As an alternative, the negative log-likelihood function ($-\mathcal{L}$) can instead be minimised. Doing the optimisation analytically involves calculating the partial derivative of the function with respect to each parameter and solving the equations

$$\frac{\partial \mathcal{L}}{\partial \mu} = 0, \frac{\partial \mathcal{L}}{\partial \sigma} = 0, \frac{\partial \mathcal{L}}{\partial a} = 0. \tag{4.14}$$

A manual approach would be to select a range of values from the parameter space and calculate the negative log-likelihood value for each case. The combination of parameter values which results in the lowest log-likelihood value are the best estimates of the true parameters. In fig. 4.19 a beam profile from simulation with a σ_{true} of 1.0 mm and a μ_{true} of 7.0 mm is analyzed. The best estimate of the parameters can be seen at the centre of the plot where the negative log-likelihood $-\mathcal{L}(\mu, \sigma)$ value is at its minimum. In this example the offset a was set to 0.

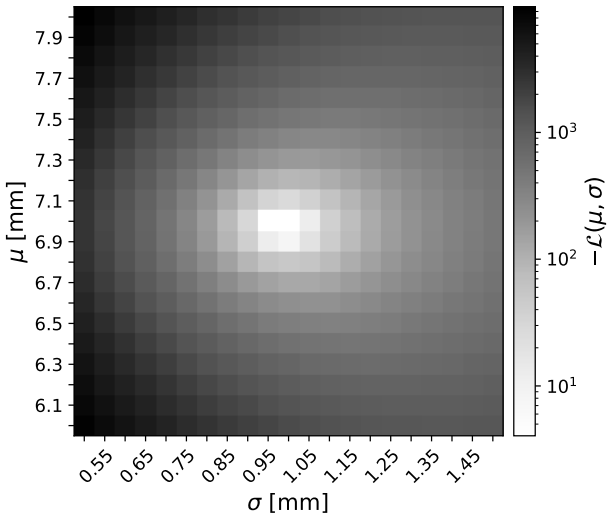

Figure 4.19: Manual minimisation of the negative log-likelihood function by scanning a range of σ and μ values.

This approach might not give the optimal estimate for the parameter values and with smaller step sizes the computation takes a long time. There are several software packages with numerical minimisers which can be used to minimize the negative log-

likelihood function. One of these is Minuit which is written in C++ [76]. A Python package is available to interface with Minuit called iminuit [77]. Another package is probfit [78] and it provides a set of predefined functions, called cost functions, that can be minimised using iminuit. One of these cost functions is the binned likelihood estimator which is what we will use for beam profile fitting with a Gaussian model.

An example of the result from the minimisation using Minuit can be seen in fig. 4.20. On the left is a simulated beam profile and on the right is a measured profile, same as was shown in fig. 4.18. From Minuit the best estimates of the parameters σ, μ and a can be extracted after the minimisation. The solid line is computed using eq. (4.7) and the estimated parameter values. The true values for the simulated data were $\sigma_{\text{true}} = 1.55\,\text{mm}$, $\mu_{\text{true}} = 35.00\,\text{mm}$. We can see in the figure for the simulation that the estimated values from the Minuit minimisation are in good agreement with the true values.

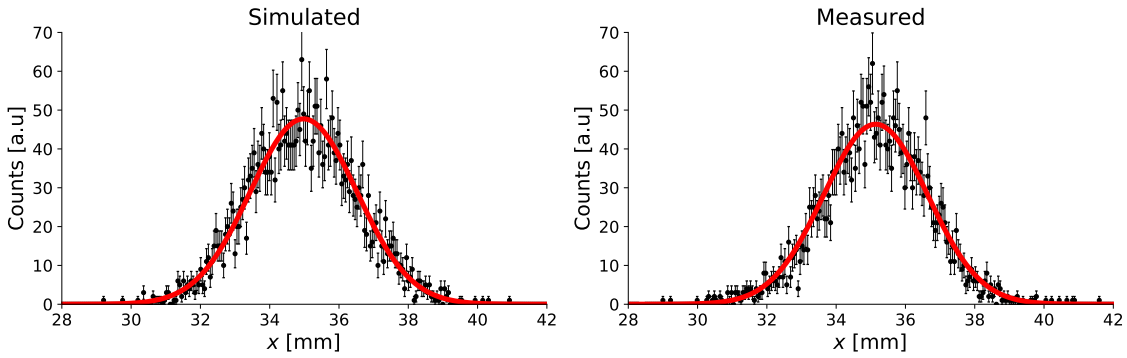

Figure 4.20: Beam profile data fitted with a Gaussian distribution scaled by the sum of counts. On the left is simulated data with an assumed beam size $\sigma_{\text{true}} = 1.55\,\text{mm}$, beam position $\mu_{\text{true}} = 35.00\,\text{mm}$ and offset $a_{\text{true}} = 0$. The result from the fitting gives $\sigma = 1.54 \pm 0.02\,\text{mm}$, $\mu = 34.98 \pm 0.03\,\text{mm}$ and $a = 0.000000 \pm 0.000004$. On the right is measured data with the result $\sigma = 1.50 \pm 0.02\,\text{mm}$, $\mu = 35.14 \pm 0.03\,\text{mm}$ and $a = 0.00003 \pm 0.00002$.

In the simulated data the offset a is 0 due to the noiseless nature of the simulation. For the measured data there is a non-negligible contribution from the noise, i.e. events that are identified as ionisation electrons which might not be part of the beam distribution as discussed in section 4.2.2. We can model this noise as a constant offset

due to the uniform distribution of these background events, under the assumption that they originate from beam loss. Without the offset a, the Gaussian model does not fit well to the beam profile. We can see in fig. 4.20 that the estimated value of a is small for this specific measurement. A simulation in section 4.4.4 will illustrate how sensitive the minimisation is to noise and how the offset a reduces the errors on the estimated parameters.

Minuit also provides errors on the estimated parameters as seen in the legend at the bottom of fig. 4.20. There is a lower limit on the error on an estimated parameter based on the number of samples used. For the sample standard deviation s we can calculate a fractional error η as

$$\eta = \frac{1}{\sqrt{2(n-1)}}, \tag{4.15}$$

where n is the total number of samples [79]. In fig. 4.21 the fractional error η is shown as a line for a range of sample sizes n. The errors on the estimated parameter σ from Minuit for ten simulated beam profiles are also shown in the figure. The errors from Minuit are marginally greater than the expected fractional errors which indicates that the errors from Minuit are consistent with expectation.

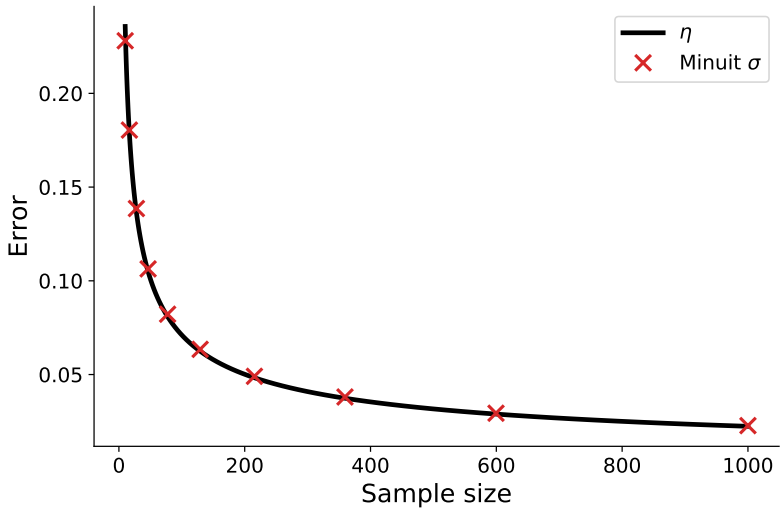

Figure 4.21: Fractional error η and extracted error from Minuit for different sample sizes.

Testing the Fit

The final question to ask is if the fit is good or not, i.e. is the beam profile consistent with a Gaussian model. With the likelihood function as defined in eq. (4.13) an analytical expression that quantifies the goodness of fit is complex compared to an alternative approach which will be discussed below.

To assess if the measured profile is consistent with a Gaussian, we will compare the likelihood value from the measured profile against the likelihood value of a test profile. From the fit we can extract the parameters σ_m, μ_m and offset a_m, as well as the number of samples N_m and the minimised likelihood value $\mathcal{L}_m$. The parameters and the number of samples are then used to create the test profile, where the offset a_m is simulated by adding noise uniformly over the detector surface. The samples from the test profile are then used to calculate the likelihood $\mathcal{L}_t$ and a ratio can be defined as

$$R = \frac{\mathcal{L}_m}{\mathcal{L}_t}. \tag{4.16}$$

This ratio quantifies how well the measured profile is in agreement with a Gaussian distribution where the noise is uniformly distributed. Ideally, if $R = 1$ the measured profile is in perfect agreement with a Gaussian, but due to sampling we expect R to deviate from 1. We can quantify the expected deviation with a Monte Carlo experiment, where we generate over 500 000 measured profiles with random parameter values of σ_m, μ_m, a_m and N_m and perform the test as outlined above and calculate the ratio. The result of the experiment can be seen in fig. 4.22 where the mean of the distribution is 0.96 and the standard deviation is 0.058. This confirms our expectation that $R \neq 1$ and that the test likelihood $\mathcal{L}_t$ is greater on average. When we create the test profile the Gaussian population that we sample from have the properties of the measured profile that we set, but because we only select relatively few samples (from 100 to 10 000) the properties of the sampled distribution will deviate from the population. As we then use these samples to calculate the test

likelihood $\mathcal{L}_t$ it will be different from the measured likelihood $\mathcal{L}_m$, which represents the minimum likelihood given the measured samples.

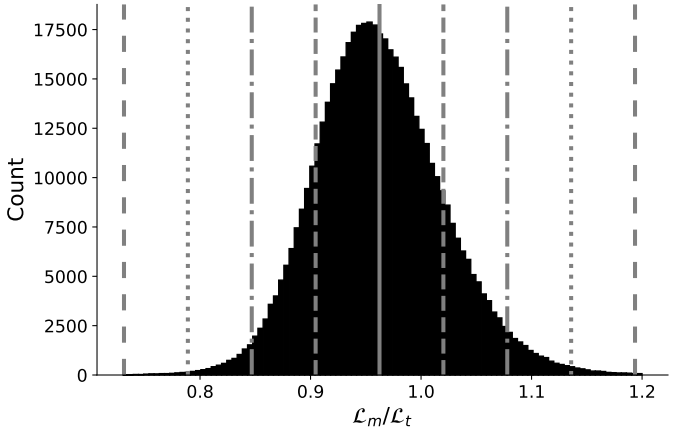

Figure 4.22: Distribution of likelihood ratio values from Monte Carlo experiments. The mean of the distribution is indicated by the solid grey line and the dashed lines indicates the one to four standard deviations from the mean.

The result from the Monte Carlo simulation enables us to set a bound on the ratio R that indicates if the fit of a Gaussian to a measured profile is good or not. In fig. 4.22 we can see that the dashed lines at four standard deviations from the mean covers almost all possible cases. We therefore define a low limit on the ratio at 0.73 and a high limit at 1.19. If the ratio from a test of a measured profile is within these limits it indicates a good fit, otherwise a Gaussian with uniformly distributed noise is not a good model for the measured beam profile.

4.3.4 Beam Size Measurement Using Statistical Moments

The ultimate goal is to quantify the transverse size of the beam from a beam profile. If the underlying distribution of the profile is well understood and stable in time, a robust and methodical analysis of the profile can be done, as discussed in the previous section. If the data is not modeled well, by a Gaussian distribution for example, we can instead determine the beam size using statistical moments. Both methods achieve the same goal of determining the beam size with different advantages and disadvantages. One drawback of using the statistical moments is that the data has

to be noise free, otherwise the noise will add up and result in an overestimation of the beam size. When fitting with the Gaussian we added the offset a to the model which accounted for noise in the data.

The statistical moments are calculated from the beam profile using the column positions x_c and the total count in each column k_c. This is a weighted calculation due to the binning of the data and because we are dealing with samples of a population we have to use the sample moments. The first moment is the weighted sample mean

$$\bar{x} = \frac{\sum_{c=1}^{n} x_c k_c}{\sum_{c=1}^{n} k_c}. \tag{4.17}$$

The second moment is the sample variance and it is calculated about the mean as

$$s^2 = \frac{\sum_{c=1}^{n} (x_c - \bar{x})^2}{\sum_{c=1}^{n} k_c - 1}. \tag{4.18}$$

By using the mean in the calculation for the variance, one degree of freedom is taken from the total. For this reason it is necessary to divide by $\sum_{c=1}^{n} k_c - 1$ instead of $\sum_{c=1}^{n} k_c$ to make sure the sample variance is an unbiased estimator of the population variance. The weighted sample standard deviation is the square root of the weighted sample variance in eq. (4.18).

The root mean square (RMS) is sometimes used in literature as a synonym for beam size, i.e. the sample standard deviation of the beam profile. Mathematically this only holds true if the population mean is 0. If the RMS is quoted for a beam profile measurement it implies the RMS deviation from the mean [79]. This nomenclature of RMS will be used in the remaining part of the thesis.

There are higher order moments which provides more information about the distribution. The third order moment, called skewness, describes the symmetry of the distribution as compared to a symmetric normal distribution. This could be of interest to measure the symmetry for the beam profile distribution. In the context of this thesis, only the first and second moment will be used.

In fig. 4.23 the same simulated, with the assumed values $\sigma_{\text{true}} = 1.55\,\text{mm}$ and $\mu_{\text{true}} = 35.00\,\text{mm}$, and measured beam profiles that have been used earlier in this chapter are shown. The beam sizes have been calculated using the statistical moments $s = \sigma_{\text{rms}}$ and $\bar{x} = \mu_{\text{rms}}$. We can see that for the simulated beam profile there is an agreement with the assumed values, which is due to the noiseless simulation that was used here. The measured data shows an overestimation of the beam size compared to the values obtained from the fitting method seen in fig. 4.20, which were $\sigma_{fit} = 1.50 \pm 0.02$ and $\mu_{fit} = 35.14 \pm 0.03$. This is due to the presence of noise in the measured data and it demonstrates that the RMS calculation is sensitive to noise. We will further explore how the noise affects the beam size measurements in section 4.4.4.

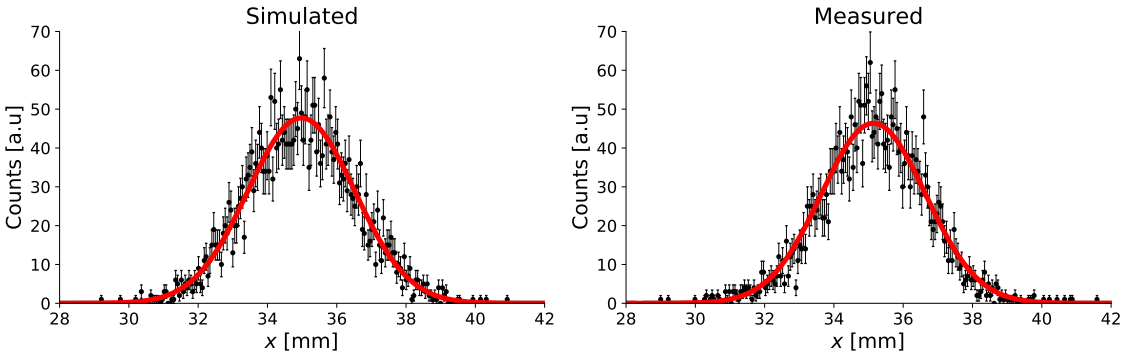

Figure 4.23: Beam profile data for which the beam size has been calculated using statistical moments. The result for the simulated profile (left) is $\sigma_{rms} = 1.54 \pm 0.02\,\text{mm}, \mu_{rms} = 34.98 \pm 0.02\,\text{mm}$ with the assumed values $\sigma_{\text{true}} = 1.55\,\text{mm}$ and $\mu_{\text{true}} = 35.00\,\text{mm}$. The result for the measured profile (right) is $\sigma_{rms} = 1.70 \pm 0.02\,\text{mm}, \mu_{rms} = 35.13 \pm 0.02\,\text{mm}$.

Sample Mean and Sample Standard Deviation Errors

With the sample mean and sample standard deviation we can now quantify the beam size without assumptions about the underlying distribution of the beam profile. But it is not enough to only use these two estimates as a description of the beam size. They also need to be accompanied with appropriate errors to facilitate a proper understanding of the estimates. This is especially important when comparing

estimates from different measurements. We start with the error for the sample mean. If we had an infinite number of independent samples, the sample mean would converge to the true population mean. The standard error of the mean (SEM) is defined as

$$\mathrm{s}_{\bar{x}} = \frac{s}{\sqrt{n}}. \tag{4.19}$$

From the simulated data in fig. 4.23 we calculated that the sample standard deviation s was 1.54 mm and n was set to 5000. This means that the SEM in this case is 0.022 mm. We can express this as a percentage of the sample mean $\bar{x}$, which was 34.98 mm, which gives us a value of 0.06 %.

It can be shown for that large $n > 20$ [80] the fractional error on the sample standard deviation can be estimated as

$$\eta = \frac{1}{\sqrt{2(n-1)}}. \tag{4.15 revisited}$$

For $n = 20$ the error is 16 %, but already at $n = 100$ it drops down to 7.1 %. In the example data from fig. 4.23, the fractional error of the sample standard deviation is 1.0 %. It is important to evaluate the magnitude of these statistical errors compared to expectations, which will conclude if they are significant or not. This was for example done to evaluate the error from the Minuit fitting procedure as was shown in fig. 4.21.

4.4 Measurement Precision and Accuracy

We can take advantage of the simulation tool detailed earlier in section 4.1 to run Monte Carlo simulations to better understand the limitations of a beam gas ionisation profile monitor based on hybrid pixel detectors. If we have a good understanding of the precision and accuracy of the measured values we can increase the confidence in the result. In this last part of the chapter, three Monte Carlo simulations will be

shown. The first two deal with the precision and accuracy limits that we can expect from this type of instrument and the last simulation is looking into how sensitive the beam size measurement is to noise. We will start by defining what is meant by precision and accuracy.

As indicated in eq. (4.15), there is a correlation between the statistical error on the beam size measurement and the number of samples n. There are also systematic errors from the instrument itself and potential systematic errors in the processing of the data, detailed in the next section (section 4.4.1). Monte Carlo simulations can be used to study the contribution of the data processing to the systematic error. A residual can be calculated for each simulated profile as

$$r = \frac{\sigma_{\text{true}} - \sigma_{\text{measured}}}{\sigma_{\text{true}}}, \tag{4.20}$$

where σ_{measured} is the beam size from the fit or the calculated RMS. The result from an example experiment is shown in fig. 4.24 where all the residual values are put into a histogram. A total of 10 000 beam profiles (each with a sample size $n = 10\,000$) were created at random and processed using the method described earlier. A residual value of 0 is the true value according to eq. (4.20). From the residuals, a mean value μ_r and a standard deviation σ_r is calculated and can be seen as the vertical lines in the figure. Two important parameters can be defined from this distribution. The difference between the true value and the mean μ_r is a measure of how far from the true value the measurements are and is therefore defining the accuracy. The width of the distribution calculated from σ_r quantifies the spread of the residuals and ultimately the spread in the measured beam size which defines the precision of a single measurement.

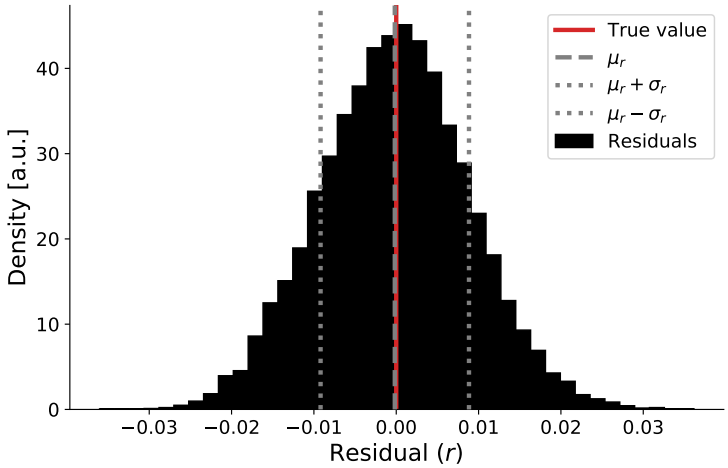

Figure 4.24: Distribution of residuals for a Monte Carlo experiment with a sample size of $n = 10\,000$.

4.4.1 Sources of Systematic Uncertainty

The sources of systematic uncertainty can come from either the instrument itself or from the processing of the Timepix3 events for the beam profile reconstruction. In the first case, alignment of the individual components with respect to each other and to the beam can be a major source of uncertainty. If the magnetic and electric fields are not parallel, the forces exerted on the ionisation electrons will displace them and and the uncertainty to determine their original position increases. The uniformity of the fields is also crucial to minimize the systematic uncertainty as was discussed earlier in sections 3.1.1 and 3.1.2. A misalignment of Timepix3 detectors relative to the beam axis can introduce a systematic error due to ionisation electron counts being spread between columns in the detector. With a pixel width of 55 µm and a length along the beam axis of 14 080 µm the rotational misalignment of a detector must be below 0.22°. The measured gaps between the detectors in the Mk II assembly resulted in a rotational misalignment below 0.17° for all detectors.

The second source of systematic uncertainty, the processing of Timepix3 events, involves the reconstruction steps described earlier in this chapter. Starting with the integration time (section 4.2.1), there are no major uncertainties as long as the time

is greater than the revolution period of the accelerator (about 2.2 µs). Below this, when integrating for turn-by-turn and bunch-by-bunch measurements, the start and stop of each integration period is critical. If we, for example, want to integrate over one turn but the triggers are not aligned or not even synchronized to the beam, the integration might start in the middle of one turn and stop in the middle of the next. The result would be a turn-by-turn beam profile measurement where a portion of the events are from turn N and the rest are from turn $N + 1$. Fortunately, the timing information that is available ensures that we can verify where the turns in the accelerator are in relation to the detected events as was shown in fig. 4.5.

Another systematic uncertainty in the reconstruction is the removal of background signal (section 4.2.2), where the parameters of the clustering algorithm and subsequent filtering steps have to be tuned to ensure that the maximum amount of background is removed, while minimising removal of the ionisation electron signal. Related to this is the compensation for the remaining background using the offset a. The assumption of a uniform distribution of background that we have used in this chapter is based on observations for the measurements that have been taken with the instrument. This does not rule out that a uniform distribution is not a suited model under all conditions and there could be a correlation with the position of the beam and the amount of background signal detected. This thesis project has not gone further into the modeling of the background signal.

4.4.2 Precision Simulation

Beam profiles were simulated for 20 different sample size values n, chosen between 50 and 10 000. For each sample size, the standard deviation σ_r of the residual values of 10 000 beam profiles is calculated, i.e. the precision. The σ_r values are multiplied by 100 to get the precision as a percentage. The result of the experiment can be seen in fig. 4.25 where the fit and RMS results are shown separately. The first observation is that the residuals from the fit procedure are consistent with those obtained from the

RMS calculation. Considering that the generator is sampling from a true Gaussian distribution, there should be no difference between a fit and RMS which holds true in this experiment. The theoretical limit from eq. (4.15) is also shown in the figure and for large values of n the values from the experiment converge to the limit. For smaller n the precision is slightly worse which indicates that there are errors in the processing of the data or the detector setup which decrease the overall precision in this region. According to the experiment we should expect a 9 % precision on a beam size measurement given a sample size of 100 ionisation electrons and to get below 1 % more than 5000 electrons are needed.

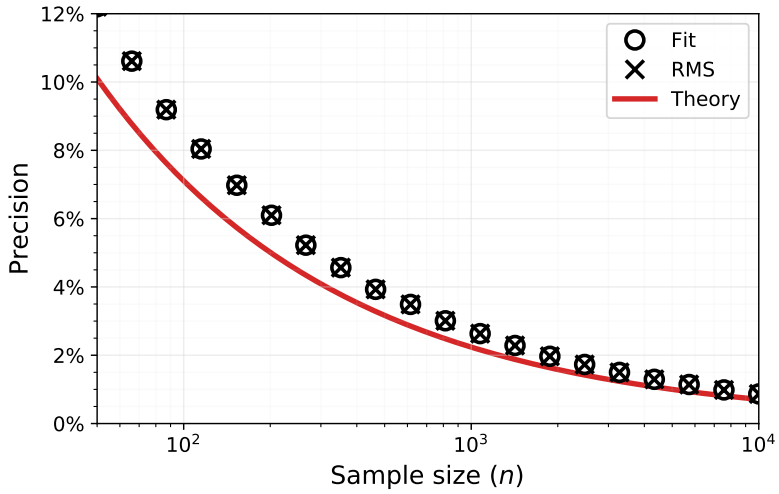

Figure 4.25: Beam size precision extracted from Monte Carlo experiments for different sample size values.

4.4.3 Accuracy Simulation

In fig. 4.24, where $n = 10\,000$, the difference between μ_r and the true value (the accuracy) is close to 0. The true beam size and position were chosen at random for this experiment and the distribution therefore reflects the accuracy for all these random values. To get a better understanding of the accuracy it is necessary to scan a range of values of beam size and position and for each fixed case average a large selection of profiles. It is also important to use a large sample size n to avoid

statistical fluctuations. For this Monte Carlo experiment $n = 500\,000$ was used and 1000 profiles were averaged for each case. The beam size was scanned from 1.0 mm to 6.0 mm and the beam position from 7.0 mm to 49.0 mm which covers a large area of the detectors, including the gaps between them.

The result can be seen in fig. 4.26 with the accuracy using the fit on the left and using the RMS on the right. For both methods there are corner cases when the beam size is large and the beam position is close to the edges of the detector area, 0 mm and 56 mm. This is due to parts of the beam profile being outside the detector region which will lead to less accurate estimation of the beam size. The RMS calculation is more sensitive to a wide beam near the edges which is due to parts of the profile missing. In contrast, the fit is able to reconstruct the true profile with part of the beam profile undetected. In a real world case the information at the tails of the beam profile could be crucial. For the discussion of accuracy these corner cases will be ignored.

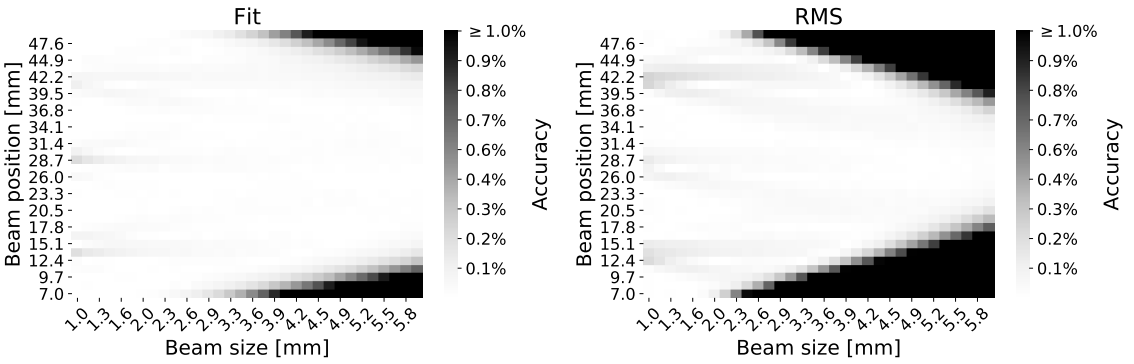

Figure 4.26: Accuracy extracted from Monte Carlo experiments for a range of beam positions and beam sizes.

In the case of small beam sizes there are three beam positions where there is a reduction in accuracy. These are around 14 mm, 28 mm and 42 mm which is where the gaps are between the detectors. It is worth noting that the accuracy in these regions is in the order of 0.3 % which is insignificant. Overall, we see that the beam size can be determined with an accuracy better than 1 %.

4.4.4 Background Signal Simulation

Earlier in section 4.3.4 it was mentioned that the beam profile is sensitive to noise and how the offset a in the Gaussian model (eq. (4.7)) mitigates this. With Monte Carlo simulations we can add noise to a simulated beam profile and calculate the fractional error on the estimated fit parameter σ. This will verify if the offset a improves the model given our assumption of the inherent noise in the beam profile data.

The noise is due to detected particles that are evenly distributed over the whole detector area and look similar to ionisation electrons. We can therefore add noise events in the simulation that are evenly distributed over the beam profile and specify the number of additional noise events as a percentage of the events that make up the beam profile. The noise seen during measurements have a correlation with how the beam behaves and how intense it is and will therefore vary from one measurement to another. We will see later in section 5.2.4 that the noise percentage for a measurement can also vary during the beam cycle in the range of 5 % to 20 %.

In fig. 4.27 the result from the Monte Carlo simulation is shown. For each case,

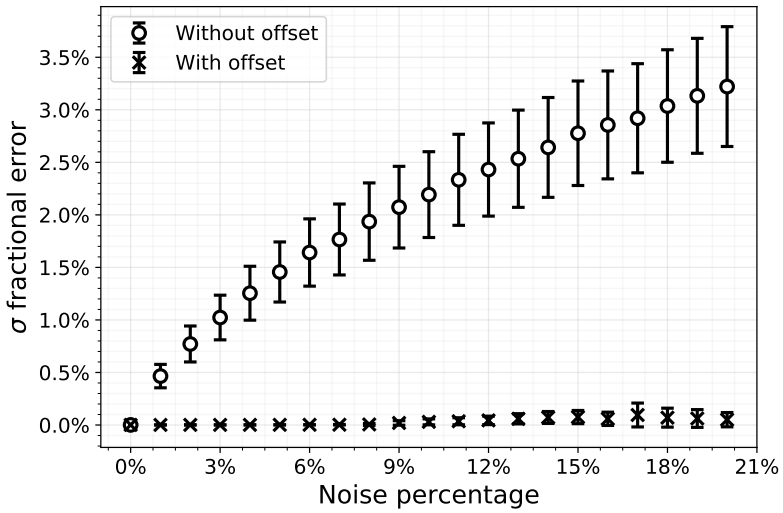

Figure 4.27: Fractional error on the beam size measurement with and without offset when different amount of noise is added to the beam profile.

a total of 10 000 profiles were simulated and both a Gaussian model with offset and without was fitted to the beam profile. Without the offset a there is a significant error on σ and also a large spread in the values, as indicated by the tall error bars for the data series without offset. We can therefore conclude that noise affects the fitted beam profile unless the offset a is taken into account and it is therefore an important parameter to model.

5 | Beam Dynamic Studies

The previous chapter described in detail how events from the Timepix3 pixel detectors can be processed to extract the ionisation electron events, from which a beam profile can be reconstructed. From the profile we can then calculate the beam size and beam position using a Gaussian model or statistical moments. We will now apply these reconstruction methods to a selection of measurements, recorded with beam in the CERN PS accelerator using the horizontal PS-BGI instrument during 2018.

In section 5.1 an emittance measurement based on a beam profile measurement will be shown. This will demonstrate the steps outlined in section 2.2 on how to go from a transverse beam profile to an emittance measurement, which is one of the main use cases of the PS-BGI instrument in the context of the LHC Injector Upgrade (LIU) project.

One of the strengths of the PS-BGI, compared to for example a wire scanner or a SEM grid, is that it can record data continuously throughout a complete beam cycle without disturbing the beam. This enables studies of how the beam evolves within the cycle, such as changes to the beam width, position and the amount of losses, which are important parameters to monitor as they provide input to the tuning and performance improvements of the accelerator. Results from such measurements will be presented in section 5.2.

Another strength is the 1.5625 ns time resolution in the Timepix3 detectors, which is in the same order of magnitude as the spacing and the length of the bunches in the beam. The PS-BGI instrument can therefore provide bunch-by-bunch measurements, where the beam profile of each individual bunch is monitored over several turns, and turn-by-turn measurements, where the beam profile is measured every turn in the accelerator. Two examples of such measurements will be detailed in section 5.3. Additionally, the turn-by-turn capabilities open up the possibility to study

the behaviour of the beam during the very first turns in the accelerator. This is important as the injected beam must be well matched to the accelerator, otherwise the size of the beam will blow up. In section 5.4, beam size and position measurements over the first 30 turns in the PS will be shown as an example of an injection mismatch measurement using the PS-BGI.

Technical upgrades of the accelerator injectors aim to reach the goals of the LIU project, with for example a new source linear accelerator (Linac4) with higher energy and with renovation of the RF systems for the PS and SPS. The project also requires new beam instrumentation that can handle the higher brightness, intensity and reduced beam size and to allow for performance measures of the accelerators, such as a brightness curve measurement. In section 5.5 a brightness curve measurement will be demonstrated using the PS-BGI where the beam size is measured for beam cycles with different intensities.

Lastly, a use case for the instrument will be shown in section 5.6 where a short integration window and the continuous recording will be used to study undesirable oscillations occurring during a transition crossing in the PS with ion beams.

5.1 Emittance Measurement

One of the main objectives of the PS-BGI instrument is to use the beam profile measurements to monitor the emittance of beams in the PS machine. This section will present an emittance measurement using the PS-BGI for three cycles of a LHC INDIV beam, which is a single bunch proton beam with a relatively low intensity of 10×10^{10} protons that is used for commissioning and machine development of the LHC. We will measure the beam profile between 178 ms and 182 ms in the beam cycle, which allows for comparison with operational wire scanner measurements measured approximately between 184 ms and 186 ms. If we calculate the normalised emittance using eq. (2.24) it requires that the beam is Gaussian in the longitudinal

and transverse planes. We will assume it is Gaussian in the longitudinal and in the transverse we will verify if the fit of our measured beam profile is consistent with a Gaussian using the procedure detailed in section 4.3.3.

5.1.1 Beam Profile Measurement

In fig. 5.1, the measured beam profiles from the BGI are shown together with Gaussian fits. On average, 1200 ionisation electron events were detected per profile, which according to the precision simulation in fig. 4.25 allows for determining the beam profile with a precision of about 2 %. From the fit parameters in fig. 5.1 we can see that both the beam size (σ) and beam position (μ) is consistent between the three measurements.

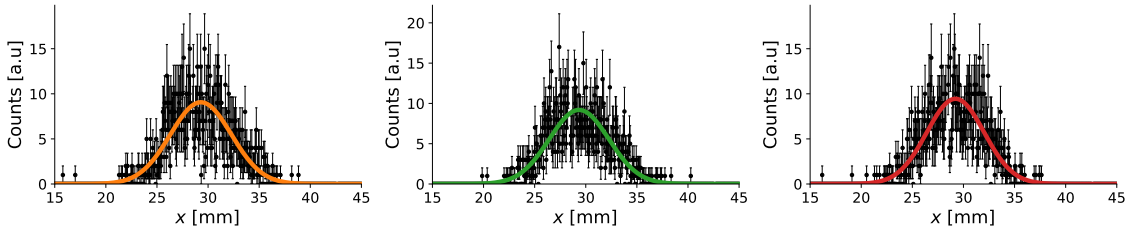

Figure 5.1: Three beam profiles measured between 1394 ms and 1396 ms for three cycles with the same type of beam. The parameters extracted from the fit are: [$\sigma = 2.98 \pm 0.06$ mm, $\mu = 29.30 \pm 0.08$ mm] (left), [$\sigma = 2.91 \pm 0.06$ mm, $\mu = 29.34 \pm 0.08$ mm] (middle) and [$\sigma = 2.81 \pm 0.06$ mm, $\mu = 29.22 \pm 0.08$ mm] (right).

The beam sizes $\sigma_{x,n}$ for the PS-BGI are shown in table 5.1 together with the beam sizes measured with the wire scanner for the same cycles. The wire scanner beam sizes were obtained using the operational software for that instrument, which fits a Gaussian model to the measured beam profile. The errors for the PS-BGI are extracted from the fit and similar errors of 2 % have been assumed for the wire scanner. From the table we can see a difference in the beam size between the PS-BGI and the wire scanner which is due to the different physical locations of the instruments in the accelerator and the different optic function values because of this.

Table 5.1: Measured beam sizes for the LHC INDIV cycles using the PS-BGI (PR.BGI82) and the operational wire scanner (PR.BWSH65).

Instrument	$\sigma_{x,1}$ [mm]	$\sigma_{x,2}$ [mm]	$\sigma_{x,3}$ [mm]
PR.BGI82	2.98 ± 0.06	2.91 ± 0.06	2.81 ± 0.06
PR.BWSH65	3.95 ± 0.08	3.95 ± 0.08	3.92 ± 0.08

5.1.2 Gaussian Fit Test

Even though the Gaussian fits look consistent with a Gaussian model in fig. 5.1, we have a tool to quantify the goodness of the fit that we will use, as detailed in the previous chapter (section 4.3.3). In fig. 5.2, the distribution of likelihood ratios $\mathcal{L}_m/\mathcal{L}_t$ is shown for the three beam profiles from the PS-BGI. A total of 1000 test profiles were generated for each measured beam profile and the mean of the ratio distributions are within the limits of 0.73 to 1.19 we defined earlier. We can therefore conclude that the measured beam profiles with the PS-BGI, shown in fig. 5.1, are consistent with a Gaussian model.

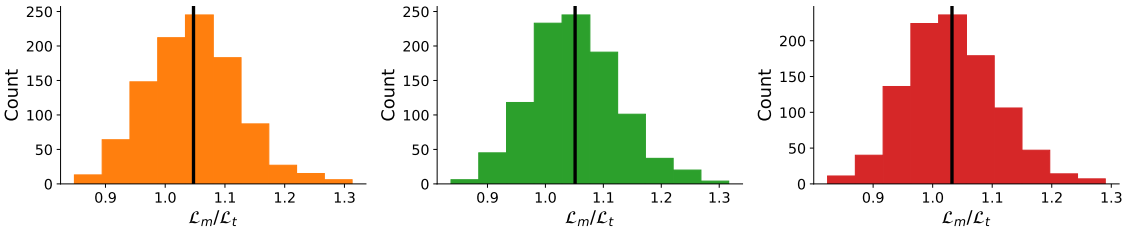

Figure 5.2: Distribution of the likelihood ratio $\mathcal{L}_m/\mathcal{L}_t$ from 1000 test profiles for three measured beam profiles.

5.1.3 Beam Emittance Calculation

Based on our observation that the transverse beam profiles are Gaussian and under the assumption that the longitudinal beam profile is also Gaussian we can calculate the normalised emittance as

$$\epsilon_{x,N} = \frac{\gamma_r \beta_r}{\beta_x} \left(\sigma_x^2 - D_x^2 \sigma_{\delta_p}^2 \right). \qquad \text{(2.24 revisited)}$$

The error on the emittance can be calculated using error propagation, derived in [27], as

$$\sigma_{\epsilon_{x,N}} = \gamma_r \beta_r \sqrt{\left(\frac{2\sigma_x}{\beta_x}\sigma_{\sigma_x}\right)^2 + \left(\frac{2D_x\delta_p}{\beta_x}\right)^2 \left(\delta_p^2\sigma_{D_x}^2 + D_x^2\sigma_{\delta_p}^2\right) + \left(\frac{\sigma_x^2 - D_x^2\delta_p^2}{\beta_x^2}\sigma_{\beta_x}\right)^2}, \tag{5.1}$$

where σ_P is the error on the parameter $P \in [\sigma_x, \beta_x, D_x, \delta_p]$. The PS-BGI and the operational wire scanner measured the same beam cycles which means that the only difference between the two instruments should come from the optic functions β_x and D_x. In table 5.2, the common beam parameters and the specific values of the optic functions at the location of the instruments are shown.

Table 5.2: Beam parameters and optic function values for LHC INDIV at the PS-BGI (PR.BGI82) and the operational wire scanner (PR.BWSH65). Retrieved from [23].

Instrument	E_{kin} [GeV]	γ_r	β_r	δ_p	β_x [m]	D_x [m]
PR.BGI82	1.4	2.49	0.92	0.49×10^{-3}	11.678	2.386
PR.BWSH65					22.450	3.122

With eqs. (2.24) and (5.1) we can calculate the emittance for each beam cycle and the corresponding error. In table 5.3 two cases are presented, one where a 2 % error is assumed for δ_p, β_x and D_x and another with 5 % error. To determine the actual error for each parameter we would have to measure them, as detailed in section 2.2, for these specific beam cycles.

Table 5.3: Calculated normalised beam emittances for the LHC INDIV cycles using the PS-BGI (PR.BGI82) and the operational wire scanner (PR.BWSH65).

Instrument	$\epsilon_{x,N,1}$ [µm]	$\epsilon_{x,N,2}$ [µm]	$\epsilon_{x,N,3}$ [µm]
2 % error on optics and momentum spread			
PR.BGI82	1.48 ± 0.09	1.40 ± 0.09	1.28 ± 0.08
PR.BWSH65	1.35 ± 0.08	1.35 ± 0.08	1.33 ± 0.08
5 % error on optics and momentum spread			
PR.BGI82	1.5 ± 0.2	1.4 ± 0.1	1.3 ± 0.1
PR.BWSH65	1.4 ± 0.1	1.4 ± 0.1	1.3 ± 0.1

5.1.4 Conclusion

An emittance measurement for a LHC INDIV beam in the PS at CERN has been presented based on beam profile measurements with the prototype PS-BGI and an operational wire scanner. The beam size was measured using the PS-BGI with a precision of about 2 % and a Monte Carlo simulation of the likelihood ratio between the measured profile and test profiles indicated that they were consistent with a Gaussian. Beam parameters and optic function values were obtained from a simulation model of the accelerator optics lattice. In an ideal case, all the parameters used to calculate the emittance should have been measured for the specific beam that was used at the location of both instrument. Such measurements are time consuming and requires dedicated beam cycles, as discussed in section 2.2, and was therefore not done for this measurement.

The emittance measurements are presented in table 5.3, under the assumption of a 2 % and a 5 % error on the beam parameters and optic function values. These results demonstrate that the PS-BGI is capable of measuring the emittance of a LHC INDIV beam and the emittance values are in agreement with the operational wire scanner.

5.2 Beam Evolution During an Acceleration Cycle

Monitoring beam parameters continuously throughout a full beam cycle gives better insight compared to snapshots at single times taken over many cycles of the accelerator. Measurements across many cycles are subject to fluctuations from one cycle to the next and therefore do not provide a complete view of the beam evolution on any individual measurement. In contrast, the Timepix3 detectors in the PS-BGI and the non-invasive measurement enable continuous measurement on a single beam cycle, with a time resolution of 1.5625 ns.

This section will present the measurement of a single beam cycle in the PS

accelerator using the PS-BGI instrument. It will demonstrate the capabilities of the instrument to resolve continuous beam profiles throughout the cycle, from which beam size and position can be calculated. Additionally, it will be shown that the background particles, i.e. events in the pixel detectors not caused by ionisation electrons, can be used to monitor beam loss at the location of the instrument and complements the beam profile measurement.

The measurement described here was taken in the PS on March 2018, for a similar LHC INDIV beam used in the emittance measurement earlier. The vacuum pressure at the location of the instrument was 6×10^{-10} mbar. At injection, the beam had an energy of 1.4 GeV and at extraction 25 GeV. The cycle is 2400 ms long but the measurement was taken between 200 ms and 1600 ms, due to the limited storage in the back-end. The electric field cage was operating with the maximum -20 kV potential difference and the pixel detector sensors were biased with -35 V. The measurement was taken at the start of the 2018 run, at a time when there were some concerns if the instrument magnet would cause an orbit distortion to the beam, due to a mismatch between the main and trim dipole fields. The magnet was therefore limited to a 0.08 T field strength, compared to the nominal 0.2 T. All four Timepix3 hybrid pixel detectors were operational in high power mode at this point in time in the Mk II detector assembly (see section 3.1.3) and the readout electronics was the first version (see section 3.2.1). Due to the limited data rate in this readout, there is data missing in parts of the cycle that will be discussed later in section 5.2.2.

5.2.1 Selection of Integration Time

In section 4.2.1 we noted that there is a compromise between precision and sample rate when deciding on the integration time for the beam profiles. The precision simulation seen in section 4.4.2 showed that with a sample size (number of ionisation electron events) of 4000 we can expect a precision on the beam size measurement of

1 %. Given the beam and detector conditions outlined above, an integration time of 5 ms was chosen to reach 4000 detected ionisation electron events. This means that the beam profile is sampled at a rate of 200 Hz.

In fig. 5.3, the beam profile is shown for three time windows in the cycle. We can see from the beam size (σ), that the beam is shrinking from one profile to the next and from the beam position (μ) we can see that it is moving around. If we only have these three snapshots of the beam cycle we do not know what is happening with the beam in between, which is why the complete evolution of the beam in the cycle is needed. Such a continuous measurement over the full cycle is not possible with a wire scanner or a SEM grid, but requires a non-invasive monitor such as the PS-BGI.

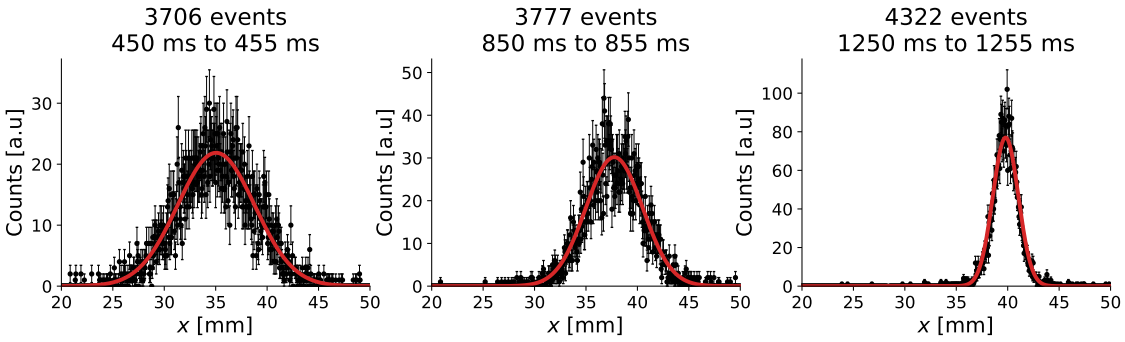

Figure 5.3: Three beam profiles from 5 ms long integration windows in the measured beam cycle with Gaussian fits. The beam size (σ) and beam position (μ) extracted from the fit are: $\sigma = 3.75 \pm 0.05\,\text{mm}, \mu = 35.04 \pm 0.06\,\text{mm}$ (left), $[\sigma = 2.77 \pm 0.04\,\text{mm}, \mu = 37.75 \pm 0.05\,\text{mm}]$ (middle) and $[\sigma = 1.24 \pm 0.02\,\text{mm}, \mu = 39.79 \pm 0.02\,\text{mm}]$ (right).

5.2.2 Beam Profile Evolution

In fig. 5.3, we can see that the beam profile is changing over time. In fig. 5.4, we can visualize all the 280 profiles that were measured in a waterfall plot. Each horizontal slice in the figure is one beam profile and the intensity is proportional to the amplitude of the profile. In this waterfall plot we can see the dynamic evolution of the beam in the cycle. For example, between 400 ms and 800 ms the beam is stable

in size and position and after this the size is shrinking due to adiabatic damping.

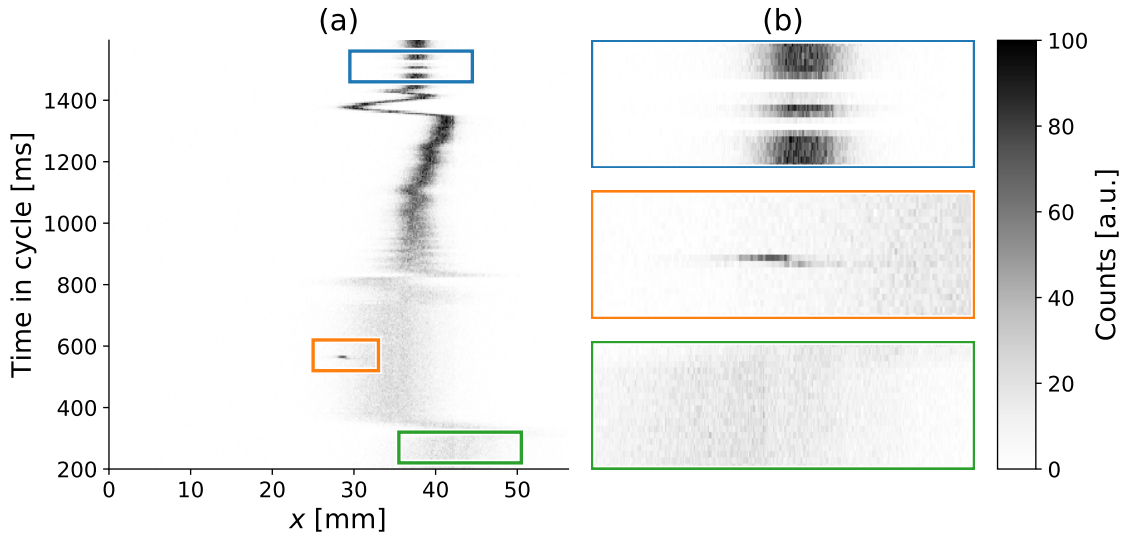

Figure 5.4: Evolution of beam profiles in the measured cycle shown as a waterfall in (a). Occurrences in the cycle that affect the beam profile measurements are highlighted and zoomed in (b), from top to bottom: data loss (blue), beam loss (orange), reduction in detection efficiency (green).

In fig. 5.4 there are three areas highlighted that can have negative effects on the measurement. At around 1500 ms (blue) there are missing beam profiles, caused by the limitations of the first version of the readout electronics, which did not have enough memory in the buffer to store events for a full beam cycle (see section 3.2).

At the highlight at around 550 ms in fig. 5.4 (orange), a high intensity spot is visible at the edge of the beam profile, which was due to a high energy particle that deposited energy over a large area. The limited readout data rate and the large number of events created from this particle caused an overflow in the timestamp for these events, which we will see again in the next section. Some of the events are therefore assigned a later timestamp and the uniformity of the cluster is lost and it cannot be detected by the background removal procedure detailed in section 4.2.2.

Lastly, there is an indication at the beginning of the cycle (green) that some columns have a lower intensity, due to limitations in the detection efficiency. In fig. 5.4 this is not obvious and we therefore need to study the beam loss evolution

in section 5.2.4 before we can draw any conclusions for this part of the cycle.

5.2.3 Beam Size and Position Evolution

The waterfall plot in fig. 5.4 gives an overview of the beam profile evolution. To quantify the dynamic behaviour of the beam we will calculate the beam size and position for each profile. For this measurement, we will use the Gaussian fit method described in section 4.3.3. In fig. 5.5 the evolution of the calculated beam size is shown at the top and the beam position at the bottom. In the same plot as the beam size, a solid red line shows the momentum of the particles in the beam. At around 750 ms, the momentum starts to increase until around 1350 ms, during which the beam size is shrinking from 3.5 mm to 1.2 mm. We also saw this in fig. 5.4 and it

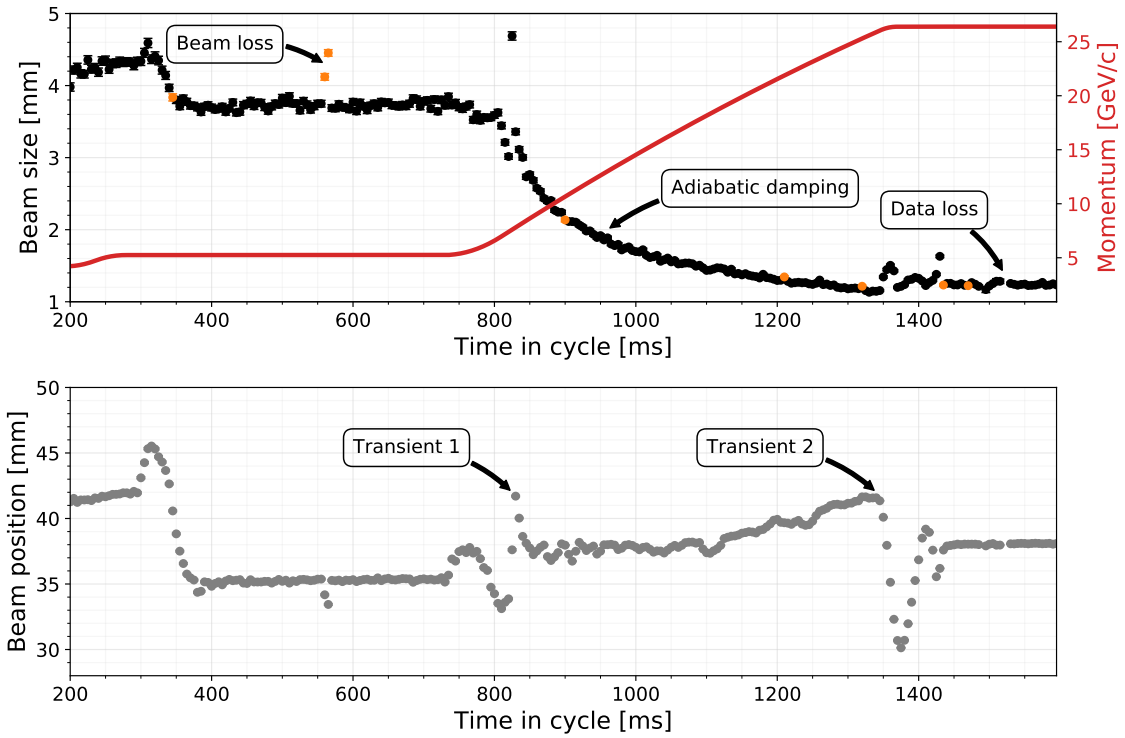

Figure 5.5: Evolution of beam parameters for a measured cycle with beam size (black) and momentum (red) at the top and the beam position (grey) at the bottom. Beam sizes marked in orange failed the goodness of fit test.

is an expected beam dynamics effect known as adiabatic damping, where the beam size will shrink during the acceleration of the beam, during which the momentum

of the particles in the beam increases.

The beam loss we could identify in fig. 5.4 at 550 ms can also be seen in fig. 5.5 as manifesting outliers in the beam size and we could mistake this for a real beam dynamics phenomena. The waterfall figure and the beam size evolution figure therefore provide complementary information which helps us to determine if an event in time is caused by beam dynamics or external disturbances such as beam loss. Another indication shown in the beam size evolution in fig. 5.5 is that the beam profiles during this time in the cycle failed the goodness of fit test as indicated by the orange color. This means that the beam profiles are not consistent with a Gaussian as detailed in section 4.3.3. For the majority of the cycle, the profiles are consistent with the Gaussian model, which means that the PS-BGI provides an independent means of showing if the Gaussian model is broadly appropriate during an acceleration cycle.

At the bottom of fig. 5.5 where the beam position evolution is shown, there are two transient events. The first transient at 825 ms is during the transition crossing of the beam and due to the relatively slow 200 Hz sample rate of the beam profile we cannot resolve the complete dynamic behaviour of the beam during this short time period. At the end of the acceleration, around 1375 ms, there is an overshoot in the main magnetic dipole field in the accelerator due to the finite response of the power converters, which causes a disturbance in the beam position. The increase in beam size at the same time is an effect of this as well since the beam moves more than 5 mm between two sample points, which causes the measured beam profile to appear wider than it is.

Earlier, we defined the fractional error on the standard sample deviation in eq. (4.15) and the standard error of the sample mean in eq. (4.19). Both of these measures are inversely proportional to the sample size n. We can use these equations to compare with the errors calculated by the Gaussian fit to verify that the errors are not underestimated. The data points in fig. 5.5 have error bars but the scaling of the axis make them indistinguishable. We therefore need to zoom in, which is done

in fig. 5.6 where a detailed view of the beam cycle between 400 ms and 500 ms is shown. For each measured beam size and position point, a limit is calculated using the equations mentioned above and the result is shown as a band around the points. The error bars from the fit extend outside the bands for both size and position, which indicates that the errors from the fit are not underestimated.

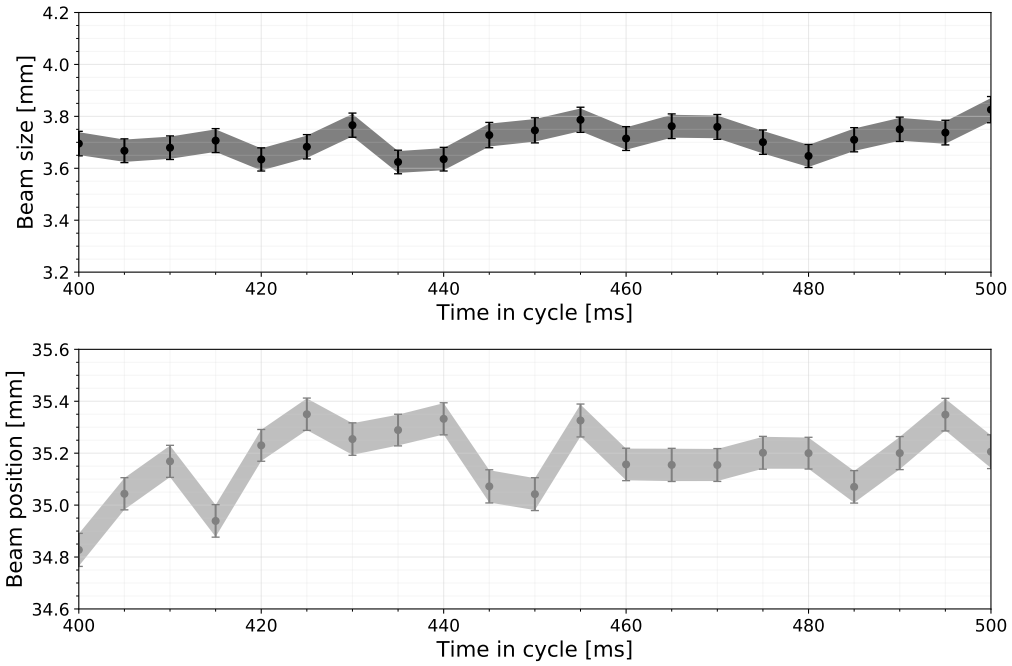

Figure 5.6: Detailed view of beam size (top) and beam position (bottom) between 400 ms and 500 ms in the measured beam cycle. The error bars are taken from the Gaussian fit errors and the grey bands are calculated from the fractional error on the standard sample deviation (top) and the standard error of the sample mean (bottom).

5.2.4 Beam Loss Evolution

As indicated earlier, there is more information about the beam evolution contained within the detected background particle events. These are all the events that were not classified as ionisation electron events during the removal of background signal, detailed in section 4.2.2. We will refer to all background particles here as beam loss.

At the bottom of fig. 5.7 the number of beam loss events (black) is shown together with a cumulative sum (red), which is an accumulation of the beam loss event count

throughout the beam cycle. The beam size evolution is shown at the top of the figure for comparison and identification of highlighted events. We discussed earlier in section 5.2.2 about a potential issue at the beginning of the cycle due to the limited data rate in the readout. One possible cause of this is now visible in the beam loss information, which has a high count value during the first 100 ms of the cycle. The average beam loss event count here is 800, which is significant as it is 20 % of the 4000 detected ionisation electron events on average. This will increase the required data rate from the pixel detectors as both types of events have to be read out at the same time.

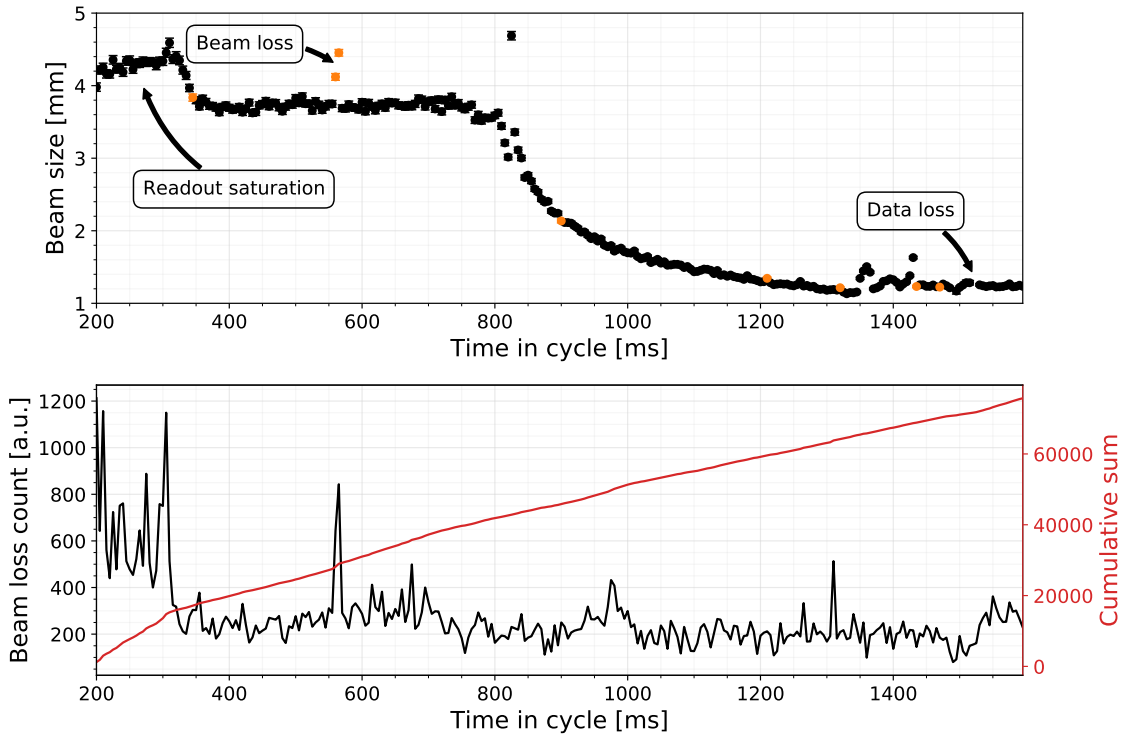

Figure 5.7: Evolution of beam size (black) is shown at the top and the beam loss (black) with cumulative sum of beam loss events (red) is shown at the bottom for the measured data.

The readout was configured such that each of the four Timepix3 pixel detectors could process 10 Mevents/s, which means they should be capable of processing 50 000 events in a 5 ms integration time. An important detail to note is that this event rate is only valid under the condition that all the events are evenly distributed over the

full 14 mm by 14 mm sensor area for one detector. If we look at the beam profile on the left in fig. 5.3, we see that the majority of the ionisation electron events are concentrated around the peak of the distribution within a width of 5 mm, which is about 90 columns or 35 % of the total detector area. As discussed in section 3.5.3, it is difficult to predict the expected event rate for the Timepix3 due to this non-uniform distribution of events. We can expect a factor 8 improvement in the processing rate when the Timepix3 is read out at its full speed instead of the slowest speed as was the case for this measurement.

Lastly, the high intensity spot we identified in fig. 5.4 is correlated with a peak of beam loss events at 550 ms. This shows that the background removal procedure is identifying clusters at this point in the cycle, but as we saw in fig. 5.4 not all of the background events are identified, which disturbs the beam size measurement.

5.2.5 Conclusion

Results have been presented for a measurement of the beam evolution using a beam gas ionisation profile monitor with the Timepix3 hybrid pixel detectors. The Timepix3 enables evolution studies throughout the cycle due to its high event processing rate (10 Mevents/s in this case) and ability to detect individual particles. For a low intensity proton beam with 10×10^{10} protons, operating under nominal vacuum conditions of 6×10^{-10} mbar, an integration time of 5 ms resulted in 4000 detected ionisation electron events per beam profile. The expected uncertainties in the measured beam size and position parameters are achieved.

The evolution of the beam profile was shown as a waterfall plot in fig. 5.4, where the size of the beam can be seen shrinking and the position change during acceleration. It also highlighted limitations in the first version of the readout system that was used during the measurement, which has been improved in later iterations. A Gaussian distribution was fitted to each beam profile and the beam size and position from the fit was extracted together with the errors reported by the fitting

procedure. The result was shown in fig. 5.5 where the dynamic behaviour of the beam size and position could be observed.

Additional information in the measured data is available from the detected background particles, which was shown in fig. 5.7. This measurement was used to identify issues related to the readout and to indicate where the beam profile was affected by the background particles.

The beam evolution measurements demonstrate the unique ability of the PS-BGI instrument to measure and reconstruct beam profiles continuously throughout a beam cycle without disturbing the beam with a time resolution of 1.5625 ns. This is one of the strengths of a beam gas ionisation profile monitor using Timepix3 hybrid pixel detectors.

5.3 Towards Bunch-by-bunch and Turn-by-turn

Bunch-by-bunch beam profile monitoring means that the profile of each bunch in the beam is resolved and integrated over a number of turns. Turn-by-turn means that the profile of the full beam is resolved for each turn in the accelerator. These are two different monitoring modes that a beam profile monitor must be capable of to fulfill the requirements of the LIU project as discussed in section 1.2. This section will present two beam measurements that explore the capabilities of the PS-BGI instrument in this context. In section 5.3.1, a bunch-by-bunch measurement is shown for a beam with multiple bunches and in section 5.3.2 a turn-by-turn measurement is presented for a beam with a single bunch.

The Timepix3 hybrid pixel detectors with their time-of-arrival (ToA) timestamp resolution of 1.5625 ns is what enables these timing critical measurements with the PS-BGI instrument. As was shown in section 3.5.1, the longitudinal length of a bunch in the PS can be below 4 ns and the spacing between two bunches can be down to 25 ns. Already in 1969, with the first IPM installed in the PS, it was

identified that a sampling period of 1 ns was necessary to resolve the profiles of individual bunches [11]. The goal of the measurements in this section is to identify the possibilities and challenges of bunch-by-bunch and turn-by-turn beam profile measurements using the PS-BGI instrument.

5.3.1 Bunch-by-bunch Profile Measurement

For the bunch-by-bunch measurement, a beam with four bunches was used and each bunch was made up of 240×10^{10} protons. The residual vacuum pressure was measured as 2×10^{-10} mbar, which should yield between 5 to 10 ionisation electron events, according to the expected values calculated in table 3.6. On average, 9 events were detected per bunch and per turn, which is in agreement with the expectations.

In fig. 5.8, the histograms of the measured event timestamps are shown for three turns of the beam. We can identify in the figure groups of four peaks, for example the peaks marked 0 to 3, which are the four bunches within one turn. The four bunches are observed again after 2.3 µs, which is the revolution period of the accelerator, and we can identify each bunch for each turn as highlighted with the colors. Peaks 0, 4 and 8 belong to the first bunch, peaks 1, 5 and 9 to the second bunch and so forth. The empty space between the last bunch and the first bunch in the next turn indicates that the bunches are not evenly spaced around the machine.

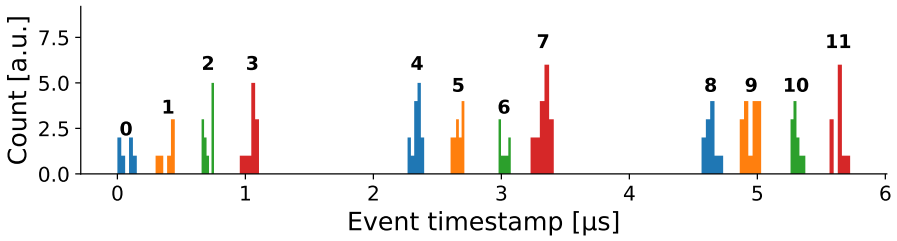

Figure 5.8: Histograms of the measured event timestamps for a beam with four bunches shown for three turns in the accelerator.

The histograms in fig. 5.8 can be extended for more than three turns, which means that we can extract the events that belong to a specific bunch and integrate

the events over several turns. One of the requirements of a beam profile monitor for the LIU project was to measure the beam profile bunch-by-bunch at 1 kHz. This is equivalent to measuring the profile of each bunch over 435 turns for the measurement presented here. We can therefore identify the events from each of the four bunches over 435 turns and reconstruct a beam profile for each bunch. On average, 3400 events were identified for each bunch, which means that we can determine the beam size with a precision of 1.5 % as shown in section 4.4.2. The result can be seen in fig. 5.9 for the four bunches. The beam sizes extracted from the Gaussian fit are in agreement between the bunches and the errors match our expectation of a 1.5 % precision.

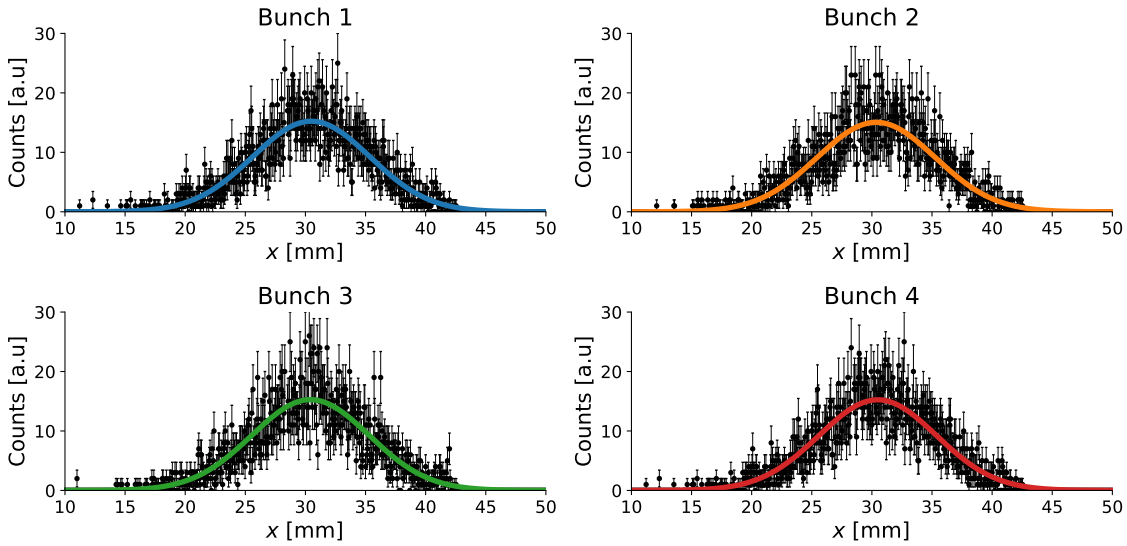

Figure 5.9: Bunch-by-bunch beam profiles over 435 turns. The Gaussian fit parameter values are for bunch 1: $\sigma = 4.88 \pm 0.07\,\text{mm}, \mu = 30.49 \pm 0.09\,\text{mm}$, bunch 2: $\sigma = 4.92 \pm 0.07\,\text{mm}, \mu = 30.36 \pm 0.09\,\text{mm}$, bunch 3: $\sigma = 4.88 \pm 0.07\,\text{mm}, \mu = 30.42 \pm 0.09\,\text{mm}$ and bunch 4: $\sigma = 4.88 \pm 0.07\,\text{mm}, \mu = 30.49 \pm 0.09\,\text{mm}$.

5.3.2 Turn-by-turn Profile Measurement

The beam used for this measurement consisted of a single bunch with an intensity of 70×10^{10} protons and an initial bunch length of 180 ns. In the previous measurement, with more than three times the number of protons per bunch, only 9

ionisation electron events were detected per bunch and per turn. This means that we would expect around 3 events per bunch and turn for this measurement, which is not enough to resolve the beam profile turn-by-turn. Gas injection would solve this issue, but it was not available when the measurement was taken in 2018. An alternative and temporary method was instead explored using sublimation of the vacuum pump mounted next to the instrument, which caused a peak in the vacuum pressure close to the instrument for about 1 min.

We do not have an exact measurement of the pressure inside the PS-BGI instrument at the exact time the measurement was taken, due to the short time the vacuum pressure peaked. A pressure gauge inside the pump measured 2×10^{-10} mbar before the sublimation and a peak value of around 4×10^{-8} mbar during the sublimation, a difference by a factor of 200. The instrument was configured using the nominal field cage voltage of -20 kV, magnetic field of 0.2 T and a sensor bias of -35 V.

The event count over time for all the recorded events in a 5 µs time period is shown in fig. 5.10, with three clusters of events separated by 2.3 µs. The clusters were identified using the clustering method described in section 4.2.2. Proton beams in the PS have a revolution period of 2.29 µs, which is in agreement with what we observe in this measurement. Within the 5 µs time period, the beam passes through the instrument three times and the close proximity of the events during each pass and the knowledge of the machine period indicates that there is only one bunch in the beam, as expected.

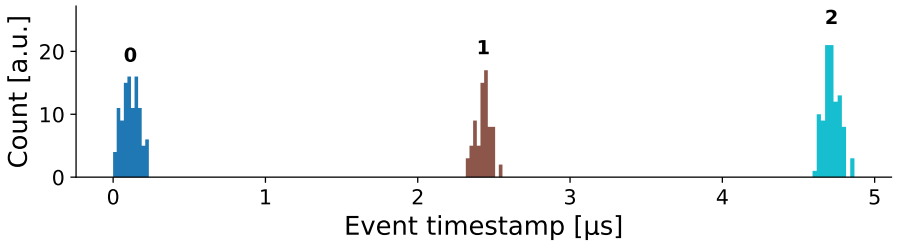

Figure 5.10: Histograms of the measured event timestamps for a single bunch beam in a 5 µs time period.

We can study the longitudinal length of the bunch in more detail by zooming into the histograms for each of the three passes shown in fig. 5.10 and the result can be seen in fig. 5.11. On average, 85 events per turn were detected, which indicates that the sublimation of the vacuum pump increased the vacuum pressure in the ionisation interaction region. Another feature we can see in the figure is that the widths of the distributions are consistent with the expected 180 ns bunch length.

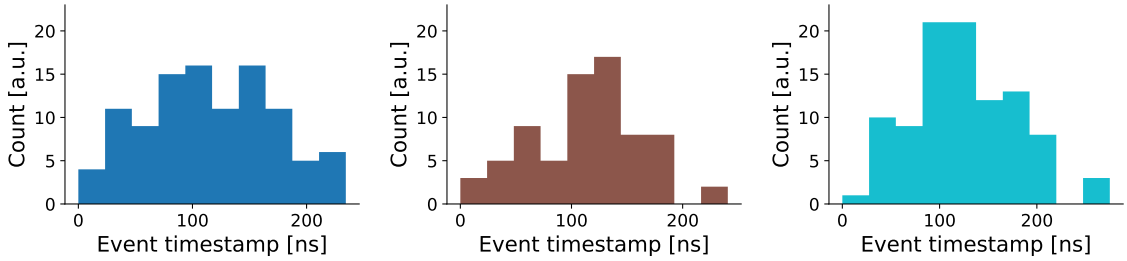

Figure 5.11: Detailed view of the event timestamp histograms for the three turns shown in fig. 5.10.

These results demonstrate that we are able to reconstruct the beam profile from the single bunch for each turn in the accelerator. In the next section (section 5.4) we will use these turn-by-turn beam profiles for an injection mismatch study, which requires turn-by-turn measurements due to rapidly changing beam dynamics at injection.

5.3.3 Conclusion

These measurements have demonstrated the unique capabilities of the PS-BGI profile monitor, based on Timepix3 hybrid pixel detectors, to measure the beam profile bunch-by-bunch and turn-by-turn. The 1.5625 ns timing resolution in the Timepix3 enables precise identification of ionisation electron events in time and is compatible with the short bunch length (4 ns) and the close bunch spacing (25 ns) of the beams in the PS.

The bunch-by-bunch measurement highlighted the identification of individual bunches within each turn of the machine. A space between the last and the first

bunch for this measurement helped with the identification of the first bunch within a turn. If the bunches are distributed evenly around the machine an additional trigger could be used to identify the start of a turn. The 9 ionisation electron events that were detected for each bunch and turn is in agreement with expectations calculated earlier in section 3.5.2. By integrating the individual events from each bunch over 435 turns, bunch-by-bunch beam profile measurements at 1 kHz were achieved, which satisfies one of the requirements for a beam profile monitor in the LIU project.

To achieve turn-by-turn beam profile measurements residual gas is not enough, as the ionisation electron yield is low and therefore limits the precision of the reconstructed beam profile. Ideally, a gas injection system is used, which increases the yield by orders of magnitude as shown in section 3.5.2, but the sublimation technique presented here is an alternative method when gas injection is not available. Even though the yield is lower in the sublimation case, the 85 detected events per turn on average allows for a turn-by-turn measurement of the beam profile. We will therefore in the next section use this turn-by-turn measurement to study the injection mismatch in the accelerator. For an operational system, the sublimation technique is not a viable option due to the lack of pressure control, which is why gas injection is foreseen to be added to the PS-BGI instrument in 2020 and will enable systematic turn-by-turn studies in the PS.

5.4 Injection Mismatch

The ionisation electron events from the turn-by-turn measurement above will, after the beam profile reconstruction, result in profiles reconstructed from an average of 70 ionisation electron events. This is due to the masking of the pixels as detailed in section 4.2.3. We can therefore expect a 10 % precision on the beam size measurement given this number of events. In this measurement we will use the turn-by-turn beam profiles to study the dynamic behaviour of the beam for the first 30 turns after

injection into the accelerator. Injection mismatch is important to study to ensure that the beam that is injected into the accelerator (PS in this case) is well matched and does not lead to a blowup of the beam size. In the PS such measurements have been done using SEM grids, but because of their interaction with the beam only 30 turns could be measured and the interaction itself could have an effect on the beam and cause a blowup during the measurement. The results from the PS-BGI will be compared against results from the SEM grids.

5.4.1 Identification of the First Turn

The measurement is started by an injection trigger (see section 3.2.5), which is propagated through the readout electronics to the shutter signals of the Timepix3 pixel detectors. This trigger has not been adjusted to compensate for the propagation delay from the source to the pixel detectors and there is therefore an offset in the measured data. Histograms of the measured event timestamps for the first 30 turns after the injection trigger at time 0 µs is shown in fig. 5.12. A repeating pattern starts after 11 µs with an average spacing of 2.29 µs, consistent with the revolution period of the PS.

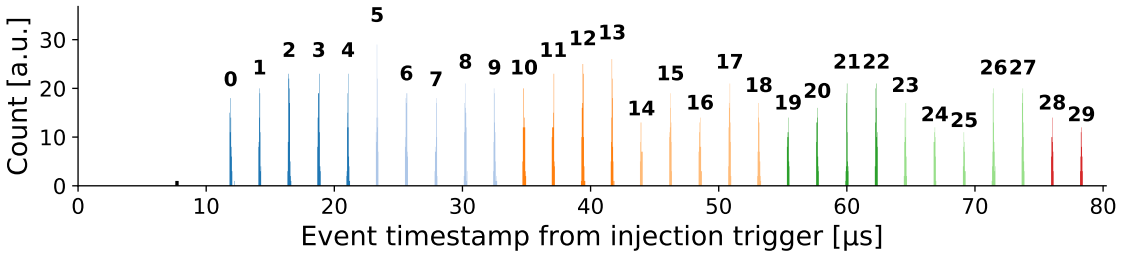

Figure 5.12: Histograms of the measured event timestamps from the injection trigger for measurement of a single bunch proton beam showing the first 30 turns.

5.4.2 Beam Profiles

From each turn we can extract the ionisation electron events, as shown in the previous section, and reconstruct a beam profile. In fig. 5.13, the first four profiles are

shown with Gaussian fits. It can be seen from the figure that the beam profile is changing from one turn to the next, both in size and position. This is due to filamentation of the beam that occurs after injection, which is correlated to how well the beam is matched when it is injected and is therefore an important behaviour of the beam to measure and study.

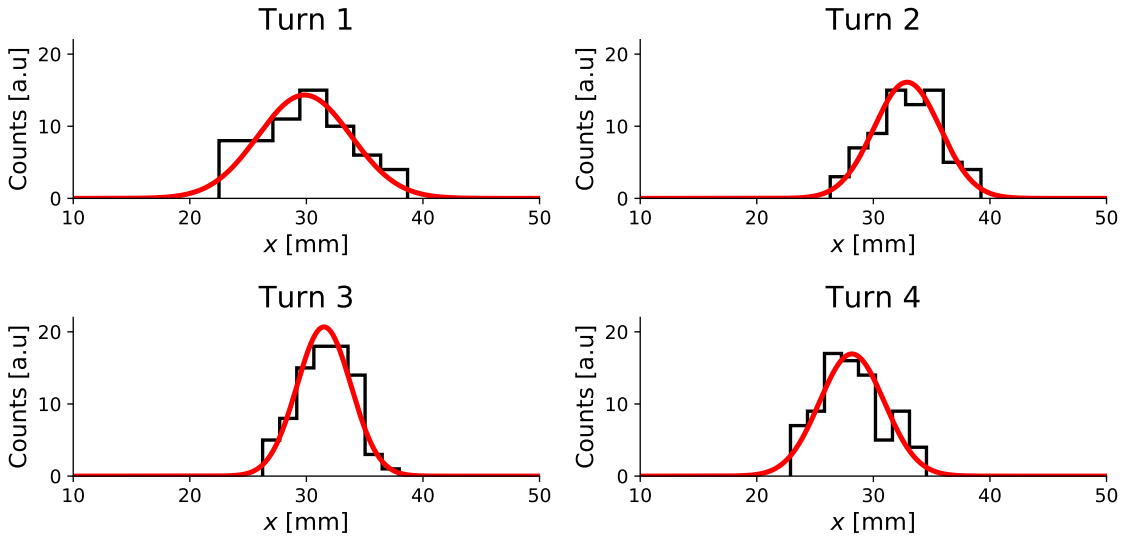

Figure 5.13: Measured beam profiles from the first four turns of a single bunch proton beam at injection.

5.4.3 Fractional Tune Measurement

From the Gaussian fit of the beam profiles we will now extract the beam position and observe how it evolves turn-by-turn. The result can be seen in fig. 5.14, where the position at each turn is shown together with a sinusoidal function that has been fitted to the position points. The frequency of the oscillation is 0.218 oscillations/turn, which is in agreement with the horizontal fractional tune (Q_x) of the PS at injection [81]. This is an expected result and is due to the betatron oscillations of the particles in the beam. The peak-to-peak amplitude of the beam position is 6.58 mm.

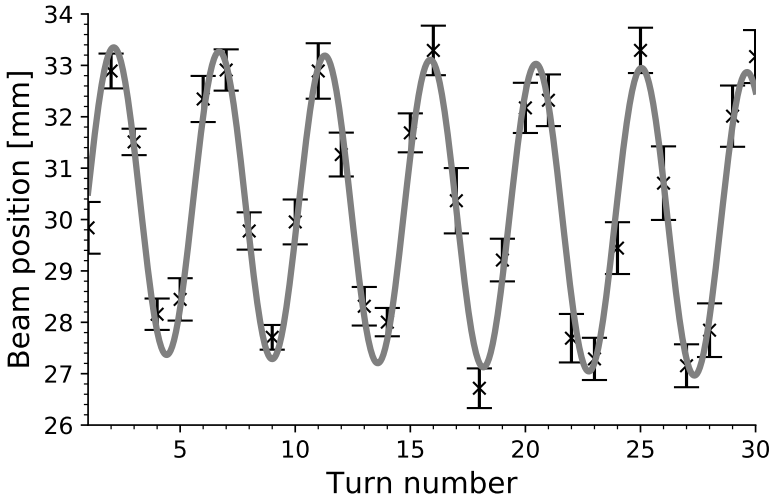

Figure 5.14: Beam position with error bars (black) for the first 30 turns with a fit function (grey) that oscillates with a frequency $\omega = 0.218$ oscillations/turn.

5.4.4 Injection Mismatch

The beam size extracted from the fit is shown in fig. 5.15 for each turn. We can observe from this result that there is also an oscillation in the beam size, with a frequency of 0.186 oscillations/turn. The beam size changes by 2.67 mm between the maximum and the minimum during the 30 turns. When the optics between

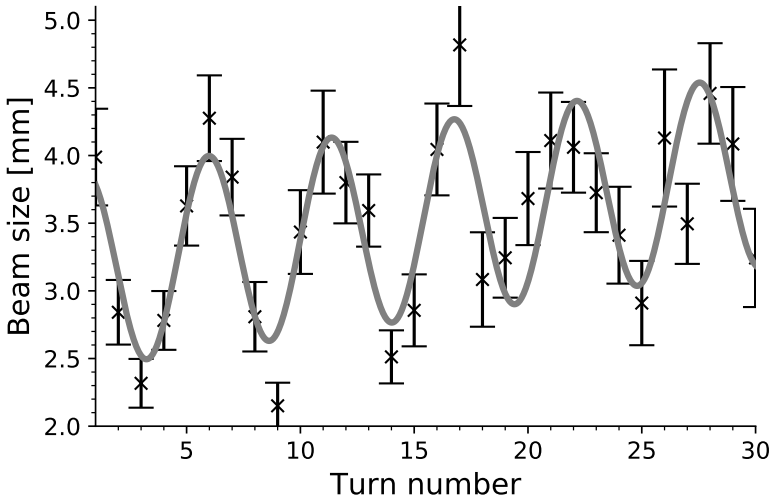

Figure 5.15: Beam size with error bars (black) for the first 30 turns with a fit function (grey) that oscillates with a frequency $\omega = 0.186$ oscillations/turn.

the injection transfer line and the PS is well matched, we expect an oscillation in

the beam size with a frequency that is twice of the fractional tune ($2Q_x$) as seen in [82]. The frequency we measure is closer to the fractional tune and even below it and is therefore not in agreement with expectations. A measurement done with SEM-grids shows similar result, where a similar type of beam used here had a beam size oscillation frequency below the fractional tune [8]. The mismatch between the PS and the transfer line from the PSB was studied in [8] to understand if injection oscillations could be a cause for an unexpected emittance blowup between PSB and PS, which is further discussed in [22]. A beam size oscillation that is closer to Q_x instead of $2Q_x$ indicates that the dispersive component of the beam size is dominant over the betatronic component (see eq. (2.25)) in this measurement. It does not however explain the measured oscillation that is below Q_x, which is now being studied to understand where the discrepancy comes from and if it is the cause of the emittance blowup.

5.4.5 Conclusion

A turn-by-turn measurement of the beam profile using the PS-BGI instrument based on Timepix3 hybrid pixel detectors have been presented with 70 detected ionisation electron events per turn on average. An ion pump was sublimated during the measurement to create a localised peak in the vacuum pressure, which increased the yield of detected ionisation electrons. This demonstrates the potential utility of a gas injection system mounted near the system, with which the amount of residual gas in the interaction region of the instrument can be controlled. The instrument should remain non-invasive to the beam and so the effect of a gas injection system must therefore be quantified and minimised.

The beam profiles were fitted with a Gaussian model to extract the beam position and size turn-by-turn. An observed oscillation in the beam position was in agreement with the expected fractional tune of the accelerator. Additionally, an oscillation in the beam size was observed that was in agreement with an independent

measurement of a similar type of beam using SEM-grids presented in [8]. Limitations in the SEM-grid measurements only allowed for 30 turns to be measured due to their interaction with the beam. The PS-BGI can measure more than the 30 turns presented here, as long as the event processing limit of the Timepix3 detector is not exceeded, as discussed in section 3.5.3. It is therefore of interest to utilise the unique non-invasive property of the PS-BGI instrument for future turn-by-turn injection oscillation studies and to measure in the order of 100 turns or more.

5.5 Brightness Curve

One of the targets for the LIU project is to provide a beam to the LHC that has a 70 % higher intensity per bunch and with a brightness that is more than doubled [3]. Brightness is the relationship between intensity and emittance (or beam size) and can be used to define the performance limitations of an accelerator as done in the LIU specification. A brightness curve measurement using the PS-BGI instrument will be presented here that highlights how the current performance of the PS can be measured. Such measurements will be essential after the LIU upgrades have been completed at the end of 2020, to verify that the accelerator performs as expected. The measurements were taken with the PS-BGI during 2018 at the same time as a wire scanner measurement. The result from the wire scanner is presented in [27], where the transverse emittance is calculated under the assumption that the longitudinal and transverse beam profile are Gaussian, as discussed in section 2.2.

5.5.1 Event Count and Intensity

As shown in section 2.3.1, the number of ionisation electrons is proportional to the number of particles in the beam (intensity) and we therefore expect to see this correlation in our measurements. We measured 20 different beams over the course of 15 min, where the intensity of the beam coming from the PSB was varied randomly

by modifying the beam current as indicated in [27]. It is not clear exactly how the beam current was changed and if it could have affected other beam parameters at the same time. The number of ionisation electron events in 550 ms time windows are shown in fig. 5.16 for the different beams and a linear correlation between event count and intensity can be seen, as expected.

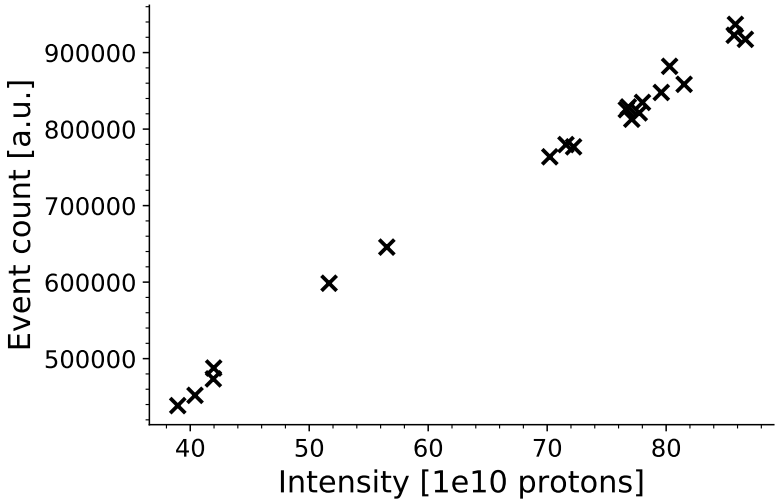

Figure 5.16: Correlation between beam intensity and number of detected events in 550 ms time windows for 20 measured beams with varying intensity.

The vacuum pressure at the instrument was measured at 1.3×10^{-10} mbar and the PS-BGI was operating in nominal conditions with the field cage at a -20 kV potential, the magnet at 0.2 T, the sensor bias at -35 V and the Timepix3 detectors in high power mode. During the measurements, the main magnetic field in the PS was kept constant and the beams were not accelerated to keep the beam size stable. The beams were dumped after circulating in the accelerator for 550 ms.

5.5.2 Beam Size Evolution and Profile

In fig. 5.17, the evolution of the beam size is shown for 3 of the 20 measured beams, with an integration time of 2 ms for each beam profile. The beam cycle starts at 0 ms but the beam is not injected until 170 ms. At 185 ms and 365 ms, highlighted in the figure, the wire scanner passed through the beams for measurements, during which

the wire is interacting with the particles in the beam for about 2 ms. The beam size measurement we will use for the brightness curve should therefore be taken before the wire scanner, at for example 176 ms, to ensure the beam has not been disturbed. We can see in fig. 5.17 that the beam size is increasing during the cycle and that there are oscillations, which could be due to the wire scanner, instabilities in the beam or both.

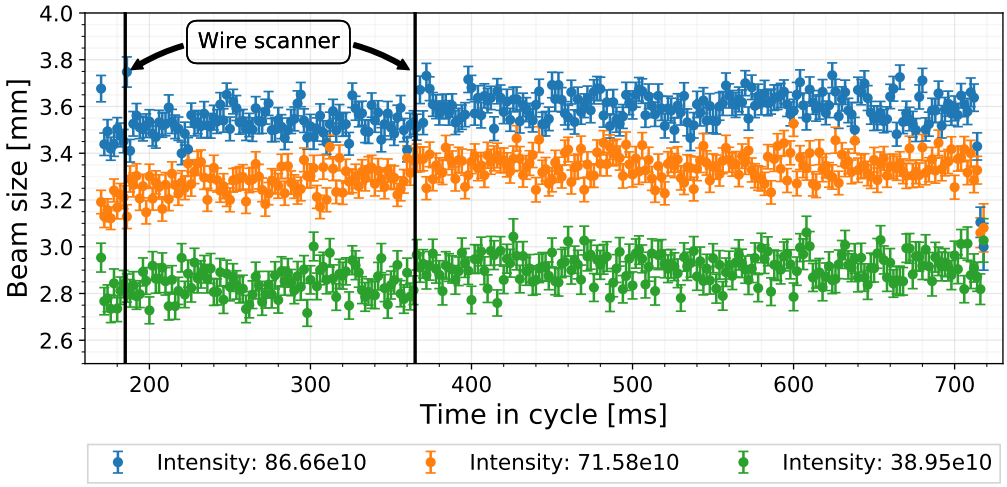

Figure 5.17: Measured beam size evolution for three cycles with different intensities. The beams were injected at 170 ms and the wire scanner was active during these measurements at 185 ms and 365 ms.

The beam profiles at 176 ms for the beams in fig. 5.17 are shown in fig. 5.18. We can see in this figure that the reduction in events for lower intensities affects the amplitude of the beam profile, as expected.

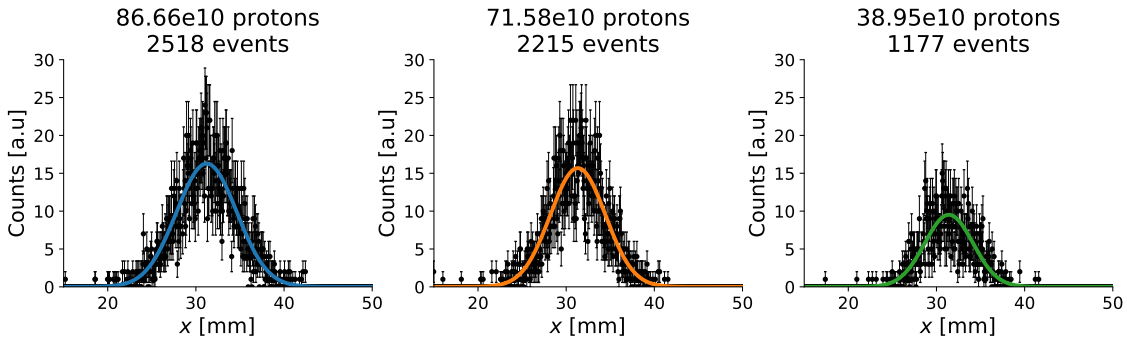

Figure 5.18: Measured beam profiles at 176 ms for three cycles with different intensities.

5.5.3 Brightness Curve

For the 20 beams measured with the PS-BGI, we calculate the beam size by fitting a Gaussian to the beam profiles at 176 ms. The result can be seen in fig. 5.19, where the error bars are the errors from the fit, and the beam size is growing with intensity. We can also see in the figure that there are cycle to cycle fluctuations in the beam size, for example the two data points at the intensity of 42×10^{10} protons are not in agreement.

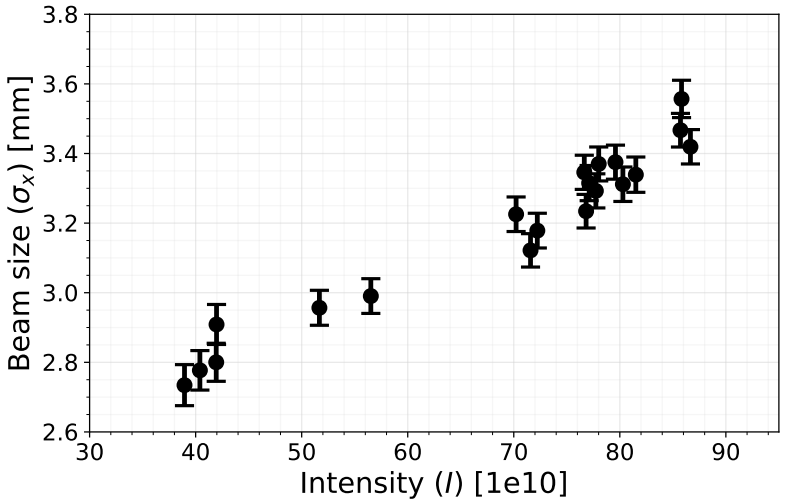

Figure 5.19: Measured beam size at 176 ms in the cycle for different intensity beams.

The next step is to calculate the emittance from the beam size as we did earlier in section 5.1. The 20 beam profiles used in fig. 5.19 all pass the goodness of fit test which means that they are consistent with a Gaussian model. We will also assume that the longitudinal beam profile is Gaussian and calculate the normalised emittance using eq. (2.24). The relativistic parameters of the beam γ_r and β_r, the momentum spread δ_p, and the optic function values β_x and D_x at the location of the PS-BGI are shown in table 5.4.

With the values in table 5.4 we can calculate the normalised emittance from the measured beam sizes shown in fig. 5.19. The result can be seen in fig. 5.20, where the emittance increases as the intensity increases. The error bars represent the errors

Table 5.4: Beam parameters and optic function values at flat bottom in the beam cycle at the location of the horizontal PS-BGI instrument. Retrieved from [23].

E_{kin} [GeV]	γ_r	β_r	δ_p	β_x [m]	D_x [m]
1.4	2.49	0.92	0.9×10^{-3}	11.678	2.386

on the emittance assuming a 5 % error on the optics values and is calculated using eq. (5.1). A linear function has been fitted to the data points in the figure and the 0.016 slope of this line is the "brightness".

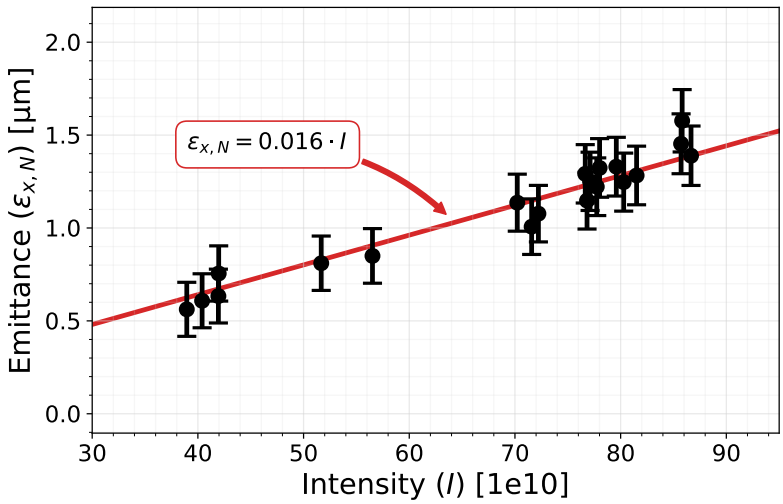

Figure 5.20: Calculated normalised emittance from the measured beam size at 176 ms in the cycle for different intensity beams. The error bars are assuming a 5 % error on the optics values and the slope of the linear curve is the "brightness".

The main limitation in the presented data is the low number of measurements (20), which was due to the operational instability of the detectors (see section 3.1.3). This type of measurement benefits from a high sample size due to fluctuations from one beam cycle to the next. For the wire scanner measurements presented in [27], several hundred beams were measured and a cycle to cycle variation in the beam size can be seen for beams with the same intensity.

Another aspect of the measurement that should not be overlooked is the stability of the beam at the chosen time of our measurement (176 ms). The beam size evolution shown in fig. 5.17 indicates that there are fluctuations throughout the cycle,

which could be due to beam dynamics or the interaction of the two wire scanner measurements in the cycle with the beam. A dedicated measurement without the wire scanner active would eliminate one possible source of disturbance and enable measuring the brightness curve at different points in time in the cycle. This is an added advantage of using an ionisation profile monitor for a brightness curve measurement, as it adds the flexibility to, for example, calculate the average beam size over N points in time or compare the brightness curve throughout the cycle.

5.5.4 Conclusion

A brightness curve measurement using the horizontal PS-BGI instrument, based on Timepix3 hybrid pixel detectors, in the PS has been presented. Brightness curve measurements quantifies the performance of an accelerator. For the LIU project, which the PS-BGI is part of, the aim is to produce a beam with 70 % higher intensity per bunch and with a brightness that is more than doubled. The results presented here demonstrate the capabilities of the horizontal PS-BGI instrument to measure brightness curves and it can therefore be used to characterise the beams for the LIU project.

In the last step of the analysis where the emittance was calculated from the beam size, the assumption of a Gaussian transverse and longitudinal beam profile was used. If the beam shape is observed to have a non-Gaussian shape, this formula for the emittance would introduce a systematic error and a full deconvolution of the betatronic and dispersive contributions to the beam size would be required before calculating the emittance. In [22], such a deconvolution is discussed for measurements done in the PSB accelerator where the beam profile was observed to have a non-Gaussian shape. They also review the impact of systematic errors on the emittance measurement due to beta and dispersion functions, where a dedicated optics measurement is suggested to precisely quantify these values at the location of each instrument. The measured optics parameters were found to differ by about

5 % from the model. A similar dedicated optics measurements is needed to obtain a precise brightness curve measurement using the horizontal PS-BGI instrument.

5.6 Transition Crossing

In the 1.4 GeV to 25 GeV energy range that the PS accelerator operates in, both protons and ions have to cross the transition energy at around 6 GeV. To reduce beam instabilities during the transition crossing, a rapid change in the optics of the accelerator is done to minimize the time spent at the transition energy. This is known as a gamma jump and uses pulsed quadrupole magnets around the accelerator. For protons, the crossing occurs in the cycle when the main magnetic field is low, at the beginning of the acceleration. The quadrupole magnets used for the gamma jump are therefore strong in comparison to the rigidity of the beam and the transition can be crossed efficiently. For ions, the crossing is instead later in the cycle, at the end of the acceleration, where the main dipole field is almost at its maximum. At this time in the cycle, the quadrupoles are weak in comparison to the rigidity of the beam, which results in an inefficient transition crossing.

The momentum spread of the beam increases at transition. This can lead to an increase in the beam size, through dispersion, and subsequently beam loss if the beam grows outside the acceptance region. To quantify the efficiency of the transition crossing we can measure the evolution of the beam size continuously during the crossing or measure the beam loss as done in [83]. The longitudinal beam profile can also be measured just before transition which provides information about instabilities within the bunch as shown in [84]. The result from these measurements are important to understand the safety margins with respect to instabilities during the transition, that could limit the desired 70 % increase in beam intensity for the LIU upgrade [3]. The beam loss and size measurements are complementary to each other and this section will present results of the beam size and position evolution

during transition as measured by the horizontal PS-BGI instrument during 2018.

Two measurements of ions beams will be presented, where the inefficient crossing will be shown to be caused by the magnet limitations outlined above. The transition crossing is at 973 ms in the cycle for these Pb^{54+} ions that were used. Unfortunately, transition could not be measured for proton beams during 2018 as the beam was positioned over the broken detector in the Mk II assembly (see section 3.1.3) at transition.

5.6.1 LHC ion beam May, 2018

Operational ion beams in the PS are mostly produced at the end of the year when the vacuum in the accelerator has stabilized and can be kept low. Before this though, there are ion beams used for configuring the machine and such a beam was measured with the PS-BGI instrument in May, 2018. The intensity of the beam was 5×10^{10} ions and the vacuum pressure was 5.4×10^{-10} mbar. The instrument was operating at the nominal conditions. A short integration time of 0.2 ms was chosen to capture the dynamic behaviour of the beam and resulted in an average of 1200 ionisation electron events per beam profile.

The beam size, position and losses around the transition crossing are shown in fig. 5.21. We can see from the beam size that before the transition at 973 ms there is an increase that is expected due to the momentum spread of the beam increasing just before transition. After the transition the beam size starts to oscillate, which is an indication of instabilities in the beam. From the beam loss data we can observe occasional peaks, but the overall trend is stable during the transition crossing. We should note that this beam loss data is at the location of the instrument, which does not correspond to the beam loss around the whole accelerator. Limitations in the aperture of the beam pipe at another location could cause losses there that are not observed at the PS-BGI.

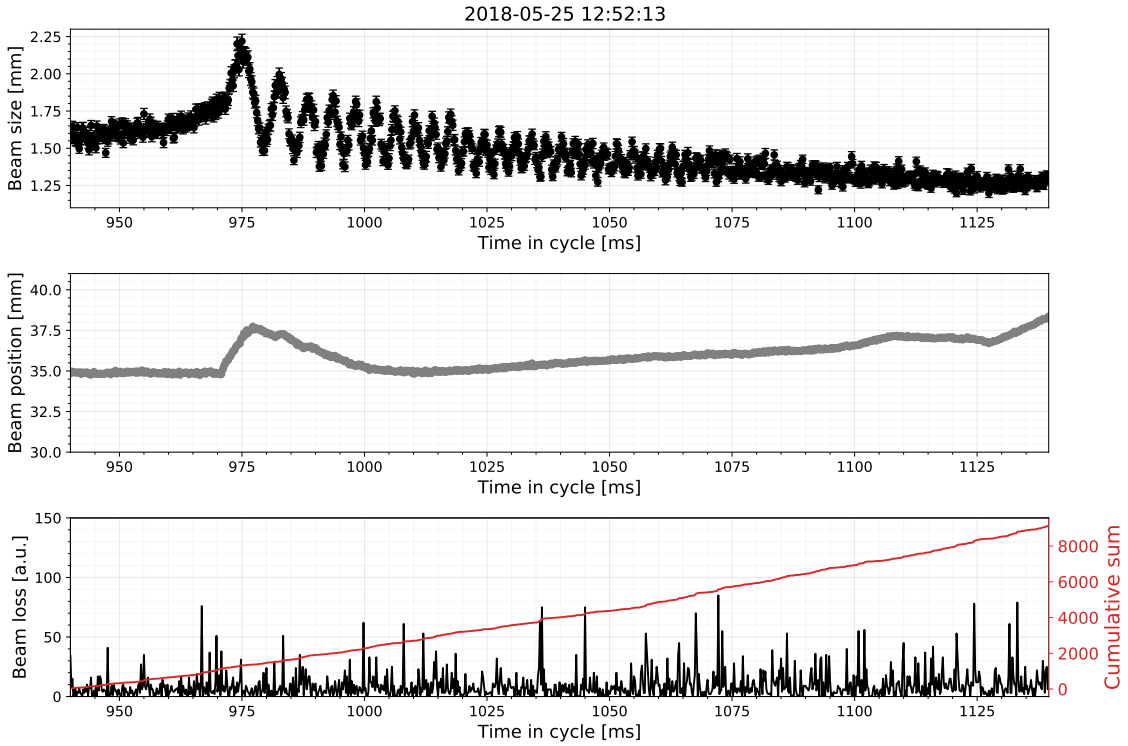

Figure 5.21: Measured beam size, position and losses during a transition crossing for an ion beam in May, 2018.

To verify that the dynamic behaviour of the beam we see in fig. 5.21 is not just from an unstable cycle, we can study the behaviour of multiple cycles. Four consecutively measured cycles are shown in fig. 5.22 and the oscillations in the beam size is consistent for all of them. The beam position is also stable from one cycle to the next. We can therefore conclude that the measured beam dynamics are real and are related to the transition crossing. Since this was an ion beam that was produced at the beginning of the year, it was not yet optimised for operation. The large amplitude of the beam size oscillations is not expected from an operational ion beam in the PS.

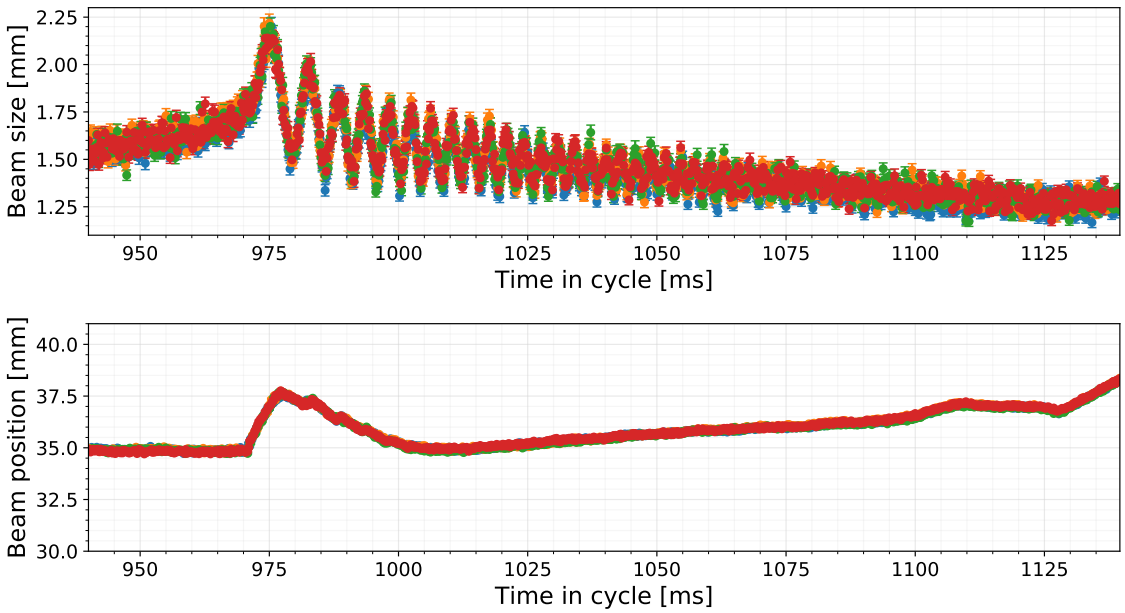

Figure 5.22: Measured beam size and position during a transition crossing for four ion beams in May, 2018.

5.6.2 LHC ion beam December, 2018

The second measurement took place later in the year in December, when the vacuum pressure was lower at 7.9×10^{-11} mbar and the intensity was marginally higher at 7.7×10^{10} ions. The instrument was operated at the same nominal values as the previous measurement, except for the Timepix3 detectors that were now in high power mode. Due to the lower vacuum, an integration time of 0.4 ms was needed to reach 1400 ionisation electron events per profile on average.

The result of the measurement can be seen in fig. 5.23. Similar to before, there is an initial increase in the beam size just before transition and after there is an oscillation, but with a lower amplitude compared to the previous measurement. We can also see that the beam position is changing more here, which is causing some issues in our measurement as the beam moves close to the edge of our detector at 980 ms, and therefore the beam size values between 980 ms and 985 ms should not be considered valid. As soon as the beam moves back over the functional detectors after this time in the cycle, the beam size values are valid again. The increased

integration time leads to a reduction in the sampling rate of the beam profile by half, but despite this we can resolve the oscillations. We also note that the amount of beam loss at the moment of the transition (973 ms) is similar to the rest of the cycle shown in fig. 5.23.

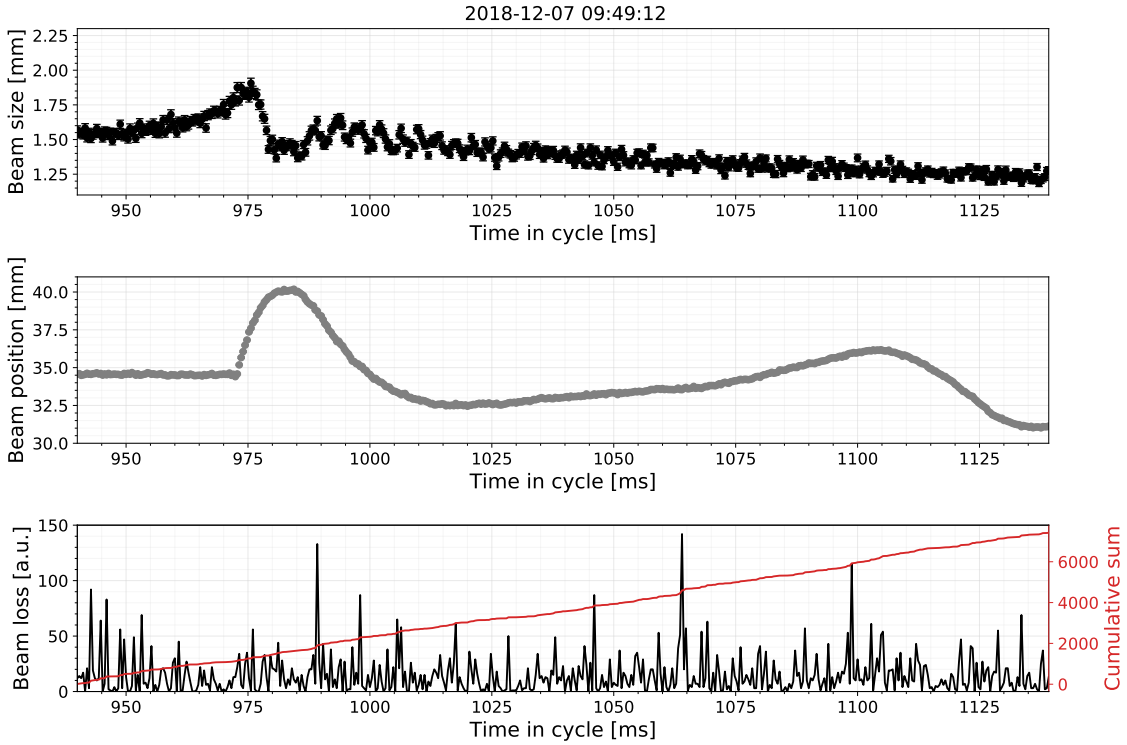

Figure 5.23: Measured beam size, position and losses during a transition crossing for an ion beam in December, 2018.

5.6.3 Conclusion

Measurements using the horizontal PS-BGI instrument in the PS have been presented where the evolution of ion beams during the transition crossing has been detailed. An expected increase in beam size just before the transition is observed, which is due to an increase in the momentum spread that couples to the horizontal transverse beam size through dispersion. After transition there is an oscillation in the beam size with different amplitude for the two measurements. An ongoing study to find the cause of this instability suggests microwave instabilities in the longitu-

dinal plane, due to impedance mismatch in the accelerator, observed for example in [85]. The result from this study is expected to be published later in 2020.

In [25] an ionisation profile monitor based on an anode strip detector is used to measure the evolution of the beam at the Fermilab Booster accelerator. This machine also crosses the transition energy and the measurements showed a similar oscillation after transition of the horizontal beam size as was presented in this measurement with the PS-BGI. The continuous measurement of the beam profile with the PS-BGI instrument provides a detailed observation of the transverse dynamic behaviour of the beam during the transition crossing in the PS. This complements the longitudinal [84] and beam loss measurements [83] that have been used to study instabilities during transition in the PS before. The operational PS-BGI instrument is therefore foreseen to aid the operators of the accelerator and the beam physicists with understanding and improving the beam dynamics during this critical point in the cycle and to aid in fulfilling the goals of the LIU project for a higher intensity and brighter beam.

6 | Future Outlooks

The previous chapters have demonstrated the strengths and weaknesses of using hybrid pixel detectors in ionisation profile monitors (IPMs) and the new insights into the beam dynamics that have already been made possible with the first PS-BGI prototype instrument. In this chapter we will first discuss a future application of such IPMs in the future high luminosity configuration of the LHC (HL-LHC) in section 6.1. The highly energetic and densely packed particle beam in HL-LHC presents new challenges that need to be addressed. Secondly, improvements to the readout electronics and pixel detectors will be addressed in section 6.2, both in the short term, based on available components, but also looking into the future with emphasis on specifications that will improve the performance of pixel detector based IPMs and facilitate the application of IPMs at future accelerator facilities.

6.1 High Luminosity LHC

There is currently no beam profile monitor installed in the LHC that can measure both beams through the full cycle, in particular during the acceleration of the beams to the top energy of 7 TeV. The limitations of the currently installed wire scanners and synchrotron light monitors are discussed in [2]. In summary, the wire scanners are of the linear type with a maximum scan speed of 1 m/s [5], which limits their use to low intensity beams at injection only, and the synchrotron light monitors suffer from large uncertainties due to multiple sources of synchrotron light and low light yield.

A new beam profile monitor was therefore proposed, called the Beam Gas Vertex detector (BGV) [86]. The BGV is based on the reconstruction of inelastic beam gas interactions, where the density of the reconstructed vertices is a direct measure of the beam profile. A demonstrator BGV was built and successful in measuring the

beam size throughout the LHC cycle, but limitations in the detector did not allow for beam profile measurements [87].

IPMs, called LHC-BGIs, were installed in the LHC until 2017 when they had to be removed due to impedance mismatch in the instrument vacuum chamber that caused excessive heating and destruction of the instruments. The detector in the LHC-BGIs was based on a MCP with a phosphor screen that was imaged with a camera, as discussed in section 2.3.3. Each LHC-BGI had a 0.2 T dipole magnet to create a magnetic field parallel to the electric field and compensate for the effect of the beam bunch on the ionisation electrons, but initial results indicated that the measured beam size at high energy was larger than expected [88]. An investigation was launched and the high charge density of the LHC bunch at high energies was found to cause deformations to the measured beam profile in the LHC-BGI. This was due to the distortion of the electron trajectories caused by the beam space charge, which was not completely compensated for by the 0.2 T magnetic field [89]. A method based on machine learning was explored to correct the distorted beam profile in [90], but another solution for this issue would be to increase the strength of the magnetic field which we will explore in the next section.

The beam parameters that affects the beam profile for the HL-LHC differ slightly from the nominal LHC that are currently used and are shown in table 6.1. The total number of bunches (n_b) is decreased but the number of protons in each bunch (N_b) is increased and at the same time the transverse normalised emittance ($\epsilon_{x,y,N}$) of the beam is decreased for HL-LHC. Both of these changes result in an increased charge

Table 6.1: Parameters for nominal LHC and future HL-LHC beams in the LHC, from [2].

Beam:	Nominal LHC	HL-LHC
Peak luminosity [$\mathrm{cm^{-2}\,s^{-1}}$]	1×10^{34}	5×10^{34}
N_b [10^{11} p/b]	1.5	2.2
n_b	2808	2748
$\epsilon_{x,y,N}$ [µm]	3.75	2.50

density of the beam, which as the investigation of the LHC-BGI showed, was one of the main reasons for the distortion of the measured beam profile. It is therefore foreseen that for a HL-LHC-BGI instrument, the measured beam profile will be even more distorted with respect to the true beam profile compared to the LHC-BGI, an effect which has to be mitigated.

6.1.1 Magnetic Field Strength

The magnetic field parallel to the electric field confines the ionisation electrons as they are accelerated towards the detector. A stronger field causes them to spiral with a smaller gyroradius and reduces the effect of the beam on the electrons. Several tools have been developed over the years to simulate the transport of ionisation products in IPMs, one of which is called "Virtual-IPM" [91].

The Virtual-IPM simulation starts from the ionisation process where the initial direction and energy of the ionisation electrons is determined. In the next step, the trajectory of the electrons is tracked through electric and magnetic fields, which includes the external magnetic field from the instrument magnet, the electric drift field and time varying electric fields due to space charge of the beam bunch. Once the tracked particles have reached the final position, where the detector is located, the simulation stops. The beam parameters used for the simulation are summarised in table 6.2. A circular beam is used and the beam size ($\sigma_{x,y}$) is calculated from the normalised emittance in table 6.1 at the location of the previous LHC-BGI instrument.

Table 6.2: Parameters used for Virtual-IPM simulation of HL-LHC beams.

Energy:	Injection (450 GeV)	Collision (7 TeV)
N_b [10^{11} p/b]	2.2	2.2
$\sigma_{x,y}$ [µm]	1000	270
Electric drift field [kV/m]	330	330

If we simulate a HL-LHC-BGI without a magnetic field we can observe the severe

distortion of the measured beam profiled caused by the interaction with the bunch. The result can be seen in fig. 6.1, where the beam profiles measured at the location of the detector (shown in red) are severely distorted both at injection and collision energy.

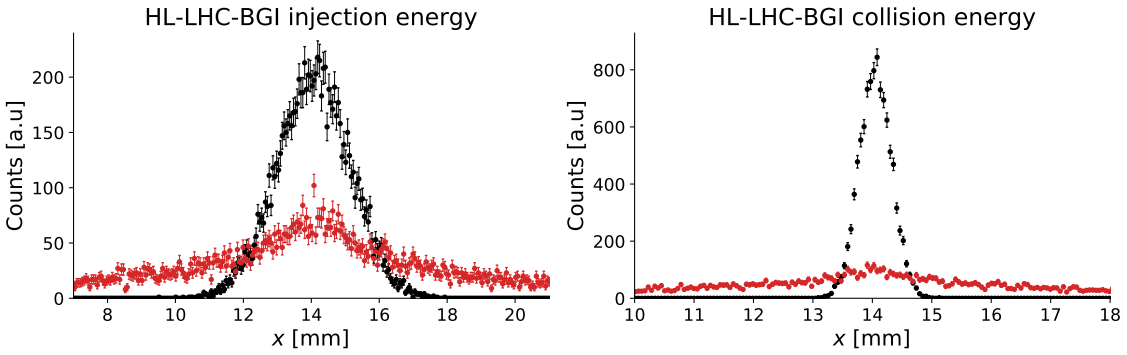

Figure 6.1: Virtual-IPM simulation of the true (black) and measured (red) beam profile for a HL-LHC-BGI instrument without a magnetic field. On the left, the profile shortly after injection (injection energy) is shown and on the right, the profile after acceleration (collision energy) is shown.

If we now add a 0.2 T magnetic field, same as the original LHC-BGI, we can see in fig. 6.2 that at injection energy the measured beam profile (red) is in good agreement with the true profile (black), but at collision energy the measured profile is distorted. It is therefore not accurate to measure HL-LHC beams at 7 TeV using the original 0.2 T LHC-BGI magnets.

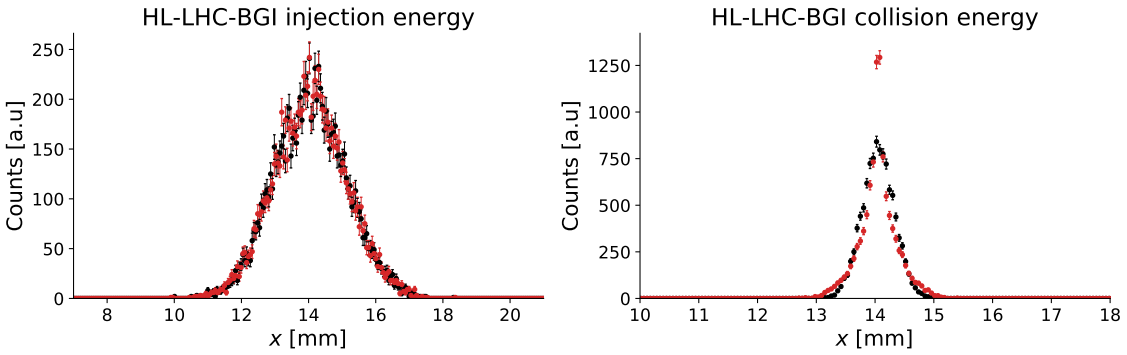

Figure 6.2: Virtual-IPM simulation of the true (black) and measured (red) beam profile for a HL-LHC-BGI instrument with a 0.2 T magnetic field.

In Virtual-IPM we can simulate the response of a HL-LHC-BGI instrument for

a range of magnetic field strengths from, for example, 0 T up to 1 T. To quantify the distortion of the beam profile we can define the ratio $\frac{\sigma_{\text{true}}}{\sigma_{\text{measured}}}$ that compares the beam size of the true profile against the size of the measured profile. A value of 1.0 means that they are in agreement and there is no distortion. The standard deviation is only a simplification considering that that measured beam profile at a low magnetic field is distorted, as seen on the right in fig. 6.2. The result can be seen in fig. 6.3 for the injection energy case on the left and collision energy on the right. As alluded to above, at injection energy the distortion is less severe and the ratio converges to 1.0 at around 0.2 T, while at collision energy a magnetic field strength > 0.5 T is required. Simulation results for the LHC-BGI case can be seen in [89] where the 0.2 T magnet was compared to a 1.0 T magnet.

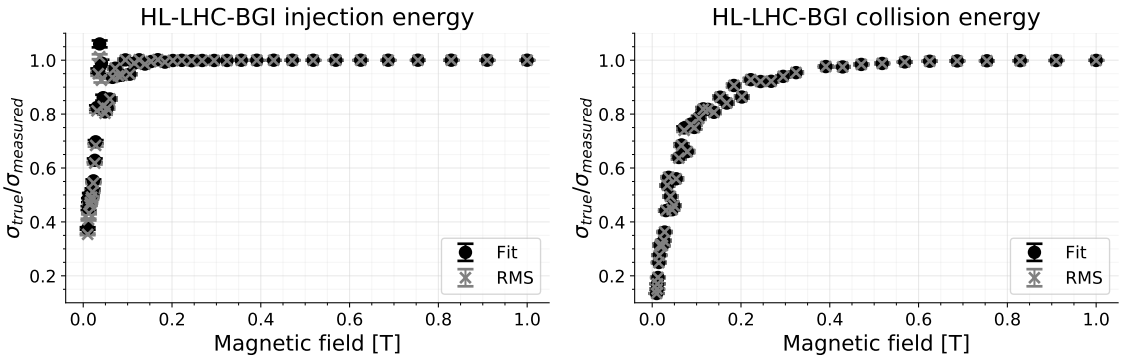

Figure 6.3: Virtual-IPM simulations comparing the beam size of the true and measured beam profile for a range of magnetic field strengths. The beam size are calculated using a Gaussian fit (black) and from the RMS (grey).

A HL-LHC-BGI beam profile monitor therefore requires a dipole magnet with a field strength of at least 0.5 T to reduce the distortion due to the high charge density of the beam. Another aspect of the magnet we need to consider is the cost and size that scales with the required field strength. Electromagnets also require power converters which adds initial cost but more importantly, it increases the operational cost of the instrument due to the energy consumption. An alternative approach could be permanent magnets that do no require any power during operation, which has been successfully realised for an IPM [92]. Due to the importance of the magnetic

field in a HL-LHC-BGI, the choice of magnet design and technology needs to be investigated before implementing the instrument.

A detailed study of the effects that contribute to the profile distortion is conducted in [93] where the Virtual-IPM software is also used. Profile correction methods are proposed in the paper as alternatives to a large magnetic field where for example machine learning is utilised. The HL-LHC-BGI instrument requires further studies to ensure that the profile distortion is understood and can be corrected for.

6.1.2 Expected Ionisation Yield and Hit Rates

In section 3.5.2 the expected ionisation yield was shown for the LHC type beams in the PS for the PS-BGI instrument. Similar calculations can be done for a HL-LHC-BGI instrument and the result can be seen in table 6.3. The number of protons per bunch (N_b) is comparable to the 26×10^{10} p/b at extraction in the PS for HL-LHC type beams. This is because the content of each bunch for the LHC is defined in the PS and the next injector (SPS) will not modify this parameter. The expected ionisation yield was calculated using eq. (2.26) and only the ionisation cross section (σ_ion) and the number of protons per bunch (N_b) were changed compared to the PS calculation. Both of these parameters are comparable between the PS and HL-LHC and the expected ionisation yield from each bunch ($n_\mathrm{e,b,t}$) is therefore similar as well. The big difference is the increase in number of bunches (from 72 to 2748) which leads to 50 times more electrons per turn ($n_\mathrm{e,t}$) for HL-LHC-BGI compared to the PS-BGI.

Based on the expected ionisation yield above, the HL-LHC revolution period (t_rev), the bunch spacing (t_b) and the bunch length (B_l) we can calculate the expected detector hit rates for a HL-LHC-BGI instrument. The result from the earlier PS-BGI calculation can be seen in table 3.8. In general, we can see an increase in all the hit rate values compared to the PS-BGI by a factor of 5 to 20. This is due to the

Table 6.3: Expected ionisation yield for HL-LHC-BGI.

	Injection	Collision
E_k [GeV]	450	7000
N_b [10^{10} p/b]	22	22
n_b	2748	2748
Residual gas (H_2, 5×10^{-10} mbar)		
$\sigma_{\mathrm{ion},H_2}$ [Mb]	0.300	0.371
$n_\mathrm{e,b,t}$	1.14	1.41
$n_\mathrm{e,t}$	3130	3870
Gas injection (Ar, 1×10^{-8} mbar)		
$\sigma_\mathrm{ion,Ar}$ [Mb]	1.61	2.04
$n_\mathrm{e,b,t}$	122	155
$n_\mathrm{e,t}$	335 000	426 000

shorter bunch length and the increase in the yield per turn ($n_\mathrm{e,t}$) for HL-LHC-BGI, which affects the maximum hit rate as defined in eq. (3.8). Additionally, the smaller beam size (σ_x) leads to an increase in the hit rate, as can be seen in eqs. (3.8), (3.10) and (3.11).

Table 6.4: Expected detector hit rates for HL-LHC-BGI.

	Injection	Collision
t_rev [µs]	88.9	88.9
t_b [ns]	25	25
$B_l = 4\sigma_l$ [ns]	1.65	1.04
σ_x [mm]	1.00	0.270
Residual gas (H_2, 5×10^{-10} mbar)		
$R_\mathrm{det}^\mathrm{avg}$ [Mhits/s]	35.2	43.5
$R_\mathrm{det}^\mathrm{max}$ [Mhits/s]	938	1840
$R_\mathrm{pixel}^\mathrm{avg}$ [hits/s]	1890	8660
$R_\mathrm{pixel}^\mathrm{max}$ [khits/s]	17 500	128 000
Gas injection (Ar, 1×10^{-8} mbar)		
$R_\mathrm{det}^\mathrm{avg}$ [Ghits/s]	3.77	4.79
$R_\mathrm{det}^\mathrm{max}$ [Ghits/s]	101	203
$R_\mathrm{pixel}^\mathrm{avg}$ [khits/s]	203	953
$R_\mathrm{pixel}^\mathrm{max}$ [Mhits/s]	1880	14 000

6.1.3 Detector Constraints

Due to the smaller beam size in the HL-LHC, a single Timepix3 pixel detector with a width of 14 mm could be used in a HL-LHC-BGI detector assembly compared to four in the PS-BGI. This requires that the detector is mounted under the centre of the beam and that the beam does not move by more than a few millimeters during the cycle in the HL-LHC. Two pixel detectors would reduce the risk in case the mechanical design is not aligned properly or if the beam moves, but it would also introduce a gap between the detectors. At collision energy in fig. 6.2 we can see that the centre of the beam profile is only covered by 5 pixels in a Timepix3 pixel detector. Earlier in section 3.1.3 it was noted that the gap between the pixel detectors in the PS-BGI was in the order of 1 to 3 pixels, which could therefore be problematic if the gap is located at the beam centre. To remove the gap, a single and wider pixel detector could be used, which will be discussed in section 6.2.2.

If we use the Timepix3 pixel detectors as in the PS-BGI, we can mount two of them facing each other, as shown in fig. 6.4. Compared to the orientation of the

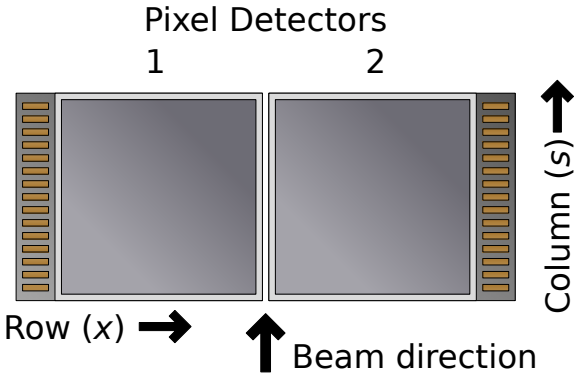

Figure 6.4: Illustration of suggested orientation of pixel detectors for HL-LHC-BGI.

detectors in the PS-BGI (fig. 3.6), where the beam was in the direction of the rows in the detectors, in fig. 6.4 the beam is now in the direction of the columns instead. As discussed in section 3.5.3, it is advantageous for readout performance to orient the Timepix3 columns in the same direction as the beam as shown in fig. 6.4. In the PS-BGI such an orientation was not possible due to the location of the wire bond

pads, where the signals are brought out, in combination with having four detectors. As seen in the previous section, the detector hit rate is expected to increase for a HL-LHC-BGI instrument compared to the PS-BGI and it is therefore important to improve the readout performance in any way possible.

6.1.4 Summary of the HL-LHC-BGI

A HL-LHC-BGI based on Timepix3 hybrid pixel detectors is a viable option and could provide continuous beam profile measurements throughout the full HL-LHC cycle, which is not possible with the currently available beam profile monitors. The complication of profile distortion due to high charge density beam in the LHC is only expected to get worse with the HL-LHC. Simulations indicate that a magnet with a field strength of more than 0.5 T is needed to minimise the distortion to acceptable levels. The detector hit rate is expected to increase by a factor of 5 to 20 compared to the PS-BGI instrument and therefore requires improvements to the Timepix3 detector orientation or an upgrade to a faster pixel detector. Four HL-LHC-BGI instruments would be needed in total for HL-LHC, one for the vertical profile and one for the horizontal profile, for beam 1 and beam 2 respectively. The impedance issues, which caused the original LHC-BGI instruments to be removed in 2017, must also be addressed to minimise excessive heating of a new HL-LHC-BGI instrument. Furthermore, the vacuum requirements of the LHC ($< 1 \times 10^{-10}$ mbar) are more stringent than in the PS ($< 1 \times 10^{-8}$ mbar) and the materials used for the detector assembly must therefore be evaluated to ensure the outgassing is compliant.

6.2 Detector and Readout Improvements

The continued development of microelectronics will increase the performance and add new features to future devices, which a beam profile monitor such as the pixel detector based IPM can benefit from. This section will briefly discuss new devices

that have been announced and will be released in the coming years and also provide a list of features for future development in detector and readout devices that could benefit beam instrumentation.

6.2.1 Readout with Increased Radiation Tolerance

An increase in the radiation tolerance of the readout would enable its use at highly radioactive locations in a machine, such as close to the injection and extraction point, or in machines where the radiation is at a high level in general. The expected radiation levels at the location of the horizontal PS-BGI instrument readout is from 10 Gy/year to 1000 Gy/year as shown earlier in section 3.2. The FPGAs, which are part of the front-end board in the current readout electronics for the PS-BGI (see section 3.2), are the weakest points in the current system in the context of radiation tolerance. These devices have not undergone dedicated radiation qualification where part of a batch of devices are irradiated in a controlled environment to determine at which point they fail, in order to quantify the tolerance of a system with devices remaining from the same batch. Such procedures are both time consuming and add additional cost, but are necessary to quantify the radiation tolerance of commercial-off-the-shelf (COTS) components due to the lack of radiation testing by the manufacturers. There are FPGAs available that have been designed specifically for radiation tolerance, which are qualified by the manufacturer and could replace the current FPGAs in the PS-BGI readout [94, 95].

Another option for the future would be to reduce the number of components in the readout where exposure to radiation is inevitable. This could be done today by removing the FPGAs on the front-end board and directly connect the Timepix3 detectors to the GBTx transceiver chips as illustrated in fig. 6.5(b). The GBTx have been designed specifically for use in radiation and no performance degradation has been observed with a TID up to 1 MGy [96]. A downside to this is that only half of the 5.12 Gbit/s maximum bandwidth of the Timepix3 can be used due to limitations

in the GBTx bandwidth. A new version of the GBTx chip is in development, called

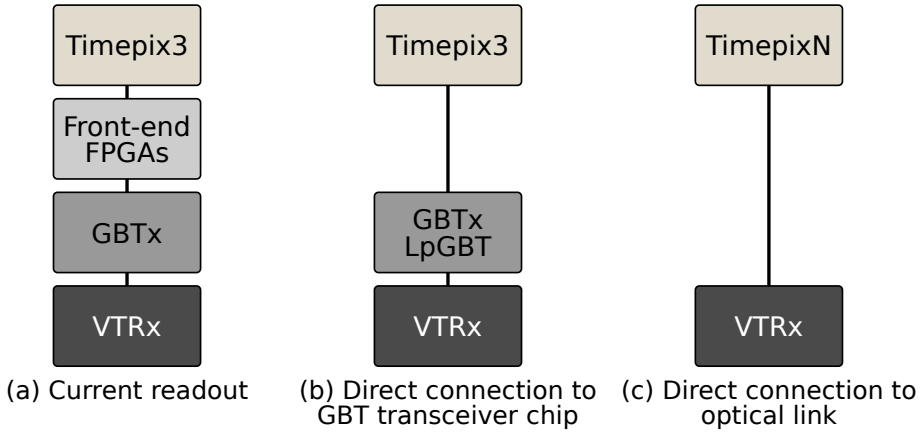

Figure 6.5: Illustration of future improvements to the readout.

LpGBT, which removes this bandwidth limitation and could therefore be used with the Timepix3 to provide the full bandwidth. The production of the LpGBT is foreseen to start at the end of 2020 [97].

We can take this idea of removing components one step further by directly connecting the pixel detectors to the VTRx optical transceiver modules, which convert between electrical and optical signals, illustrated in fig. 6.5(c). The maximum bandwidth of one Timepix3 data channel is 640 Mbit/s which may not be compatible to directly interface with a VTRx module due to a requirement of a minimum switching frequency of the signals. Additionally, the control signals to the Timepix3, such as the shutter and power enable, cannot be connected this way and would therefore require an additional controller, for example a GBTx. This would still reduce the total number of components as one GBTx could for example control all four Timepix3 detectors in the PS-BGI. To take full advantage of this last readout concept, the interface to the pixel detector would have to be changed to data channels operating at data rates compatible with the optical module, for example at 10.24 Gbit/s, and to allow for control signals to be transmitted over a high data rate serial channel instead of individual control signals. The next generation of the Timepix detector family, which will be discussed in the next section, satisfy some of these require-

ments but not all, which is why "TimepixN" is used in fig. 6.5(c) to indicate that an even later generation of the Timepix detector family would be required for this readout solution.

Another aspect of the readout which has not been discussed yet is the routing of the signals from the detector assembly inside the instrument vacuum to the front-end readout outside. In the current readout, all signals are brought out in a chain of flexible cables, connectors, feedthroughs and shielded Ethernet cables over a total distance of more than 2 m, as discussed in section 3.2. This signal path is capable of transmitting the 640 Mbit/s data channels from the Timepix3. For option (c) in fig. 6.5 the data rate would increase by an order of magnitude which will require a detailed study of the signal integrity from the detector to the optical module. If the materials in the optical module are vacuum compatible, it could be mounted inside the vacuum close to the detectors with an optical fiber connection passed through the flange of the instrument outside, if such a feedthrough is available. If the high speed electrical signals must be routed from inside the vacuum to outside, the cables, connectors and feedthroughs must be able to handle the high frequencies.

6.2.2 Future Generation of Pixel Detectors for IPM

In the earlier chapters of this thesis, limitations of using pixel detectors in a beam gas ionisation profile monitor have been highlighted such as the need for cooling and the limited experience of using pixel detectors for this application. Additionally, we discussed in section 3.5 about the expected detector response and how the Timepix3 limits the operation of the instrument in the PS and for future applications in the HL-LHC (section 6.1). In this section we will discuss features and specifications for a future pixel detector that would improve the application of pixel detectors in IPMs.

One of the major limitations of the Timepix3 was shown in section 3.5.3 where it became clear that the maximum processing rate of 85 Mevents/s for one Timepix3

was exceeded in cases with high ionisation yield and when short bunches resulted in high bursts of events. Based on the calculated hit rates, we would need a detector that can process more than 170 Gevents/s to cover all cases, a rate which is 2000 times more than the Timepix3 can handle. This would in turn increase the needed bandwidth of the readout to more than 10 Tbit/s (using the same 60 bit encoded data format as Timepix3 uses), which would be challenging to route from the detectors to the readout as discussed earlier.

One way to reduce the bandwidth is to improve the encoding by replacing the 8b/10b, which is only 80 % efficient, by for example 64b/66b, which is 97 % efficient. Another way would be to modify the content of the event packet (detailed in section 2.4.2), where 4 bits are for the header, 16 bits for pixel address, 18 bits for ToA and 10 bits for ToT. The header is used in Timepix3 to differentiate between data packets and different control packets which are all sent through the same interface. By separating data and control packets into dedicated interfaces we could remove the need for the header for the data. In fig. 3.25 we saw that the majority of the ToT values for the PS-BGI are below 100 and we can therefore reduce the ToT information to 7 bits. If we reduce the pixel address information we would loose the ability to identify background particles using clustering (section 3.4.2) and if we reduce the number of ToA bits we increase the risk of overflowing the counter, which happens after 409.6 µs for Timepix3.

So far we have reduced the packet size by 7 bits and together with more efficient encoding, if we ignore the fact that 64b/66b requires the 41 bit packets to be split and combined to form 64 bit packets, the required bandwidth is now reduced to about 7.2 Tbit/s. This is not a very significant reduction in bandwidth and thus far we have only improved the processing rate of the pixel detector. As touched upon earlier, the overflow of the ToA counter after 409.6 µs could cause issues if all events are not read out from the pixels fast enough. To avoid this issue we could increase the processing capabilities of the detector or increase the number of ToA bit that are

read out, with both of these options requiring an increase in the readout bandwidth. Additionally, in section 3.5.1 we saw that beams in the PS can have a longitudinal bunch length down to 4 ns which is close to the ToA resolution of 1.5625 ns in the Timepix3. To enable further insight into the bunch, the ToA resolution would have to be improved to around 100 ps, which requires more bits for the ToA if the overflow should be kept the same or be improved.

The dimension of the pixel detector and the pixel size could also be optimized for the IPM application. We noted previously in section 6.1 that a smaller pixel size than 55 µm x 55 µm could be beneficial to improve the resolution of the beam profile measurement for small beam sizes. A wider detector would remove the gaps that are present in the PS-BGI due to the four detectors mounted side by side on the detector assembly (section 3.1.3). The depth of the detector, along the beam direction s, could be reduced as the only benefit of a longer detector is a higher ionisation electron yield, as described in eq. (2.26). Additionally, a longer detector comes at the cost of requiring a bigger volume in which the electric and magnetic field has to be uniform.

In table 6.5 a suggestion of specifications is shown for a future pixel detector designed for an IPM application. Many of these parameters are correlated and the discussion above hopefully highlights the compromises engineers have to make when designing a pixel detector. As of the beginning of 2020 there is no pixel detector that fulfills these specifications, but the next generation of the Timepix family, called Timepix4 will be available at the end of 2020 and improves upon the Timepix3.

Table 6.5: Proposed specifications for a future pixel detector for IPM.

Event rate	> 170 Gevents/s
Readout bandwidth	> 10 Tbit/s
ToA overflow	> 1 ms
ToA resolution	< 100 ps
Dimension	30 mm x 5 mm [width x length]
Pixel size	< 55 µm x 55 µm

Timepix4

The design specification, summarised in table 6.6, for the new Timepix4 hybrid pixel detector were finalised in 2017 and the first prototype was sent for production in 2019 [98]. In comparison to the Timepix3, the Timepix4 is wider and has a 3.5 times greater detector area and an 8 times increase in the maximum event rate it can process, both of which were highlighted above as necessary improvements for a future pixel detector in an IPM. The ToA and the ToT have better resolutions and wider range which allows for more precise timing and energy information and reduces the risk of overflow for the ToA. The increased width of the Timepix4 is advantageous for an IPM application and the increased length leads to a higher ionisation electron yield at the cost of requiring a larger uniform field region, as discussed earlier.

Table 6.6: Timepix4 specifications.

Event rate	2.48 Gevents/s
Readout bandwidth	$\leq$ 163.84 Gbit/s
ToA overflow	1.6384 ms
ToA resolution	195 ps
Dimension	29 mm x 24 mm [width x length]
Pixel arrangement	512 x 448 [width x length]
Pixel size	55 µm x 55 µm

A big change in the Timepix4 is that the chip will interface to the outside using Through Silicon Via (TSV) technology to create a ball grid array (BGA) package, which is common in modern chips. This means that all the signals are routed through the bottom of the chip instead from one side using wire bond technology as is the case for the Timepix3. The main reason for this new interface technology is to allow all four sides of the detector to be free of connections and enable multiple detectors to be placed on all sides to create a pixel detector array, for example a 3 x 3 array. For IPM applications only two sides have to be free to allow for a wide detector array, such as the one used in the PS-BGI (section 3.1.3), but the new BGA

type package could simplify the manufacturing of the detector assembly. With wire bonding the detectors have to be mounted and connected to the assembly using specialised wire bonding machines where each wire is mounted one by one, while a BGA package only requires placing the chip down in the correct position and then using a reflow soldering process.

One potential issue with the BGA package is power dissipation. For the Timepix3 the whole bottom area of the detector can be attached to the detector assembly using glue or thermally conductive adhesive film as detailed in section 3.1.3, which leads to a large conduction area. In Timepix4 with the BGA, the bottom of the chip is no longer in contact with the detector assembly because of the balls in between and the power therefore has to be dissipated through the balls. There is a technique called BGA underfill, used in commercial manufacturing, where epoxy is dispensed around the chip after it has been mounted and subsequently fills the gap under the chip. This provides mechanical support for the chip and increases the thermal conduction, but in the IPM application where the detector is in vacuum the epoxy must be approved and pockets of air must be eliminated. The Timepix4 alleviates this thermal issue slightly as the power requirement is expected to be about $0.36\,\mathrm{W/cm^2}$, compared to $0.76\,\mathrm{W/cm^2}$ for the Timepix3.

Lastly, the readout bandwidth for the Timepix4 of 163.84 Gbit/s is distributed over 16 data links which each operate at up to 10.24 Gbit/s. This increases the requirements on the readout as discussed in section 6.2.1 and the only viable option for the IPM application would be to directly connect the data links to optical transceiver modules as illustrated in fig. 6.5(c).

6.2.3 3D Integrated Pixel Detector

Based on the recent development and the identified requirements of using pixel detectors in IPM applications, it is clear that the data rate will increase. From the introduction of the Timepix3 in 2013 [18] to the new Timepix4, the readout

bandwidth has increased by a factor of 32. For applications, such as the IPM, where the detector is located inside vacuum the main challenge will be to route these high speed data links out to the readout.

A different approach is to process the high speed data as close to the pixel detector as possible and output a low speed data stream of the result, for example the reconstructed beam profile for an IPM application. This requires that a processing chip is located as close to the pixel detector as possible to minimise the length of the high speed signals. Ideally, the hybrid pixel detector would be mounted on top of the processing chip as illustrated in fig. 6.6 to create a 3D integrated pixel detector. Such 3D integration has been in development since before year 2000, for example as real time image processing chips or integration of memory with processors [99]. These 3D chips are now commercially available and utilises different techniques, including TSV technology, to integrate several silicon wafers or even 3D integration inside a single wafer [100].

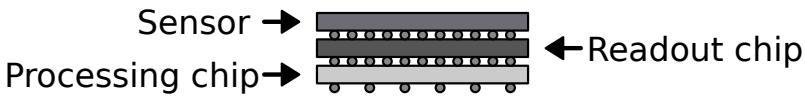

Figure 6.6: Conceptual diagram of a 3D pixel detector with local processing.

A 3D integrated pixel detector is therefore a natural progression in the development of pixel detectors and can hopefully be accomplished once the technology has matured and the price of production has decreased enough, to allow research institutes such as CERN to take advantage of this new technology. For IPM applications, the local processing would include real time clustering and removal of background events, similar to the software procedure discussed in section 3.4.2, with tunable parameters to choose what to filter out. The equalisation of the detector (section 3.3.2) could also benefit from local processing due to the high data rate it requires and the need to monitor voltages and temperatures close to the detector. Lastly, the task of reconstructing the beam profile (chapter 4) is a prime candidate for local processing as the reconstruction decimates the amount of data by several orders of

magnitude. For example, if we have a beam profile with 1000 bins, each represented by a 16 bit number, and a measurement is recorded at 1 kHz, the data rate out of the detector would be 16 Mbit/s. With a 3D integrated pixel detector and IPM specific local processing, a low speed data interface could be used and therefore reduce the complexity of the instrument, the readout design and the operation.

7 | Conclusion

This thesis project presented the implementation of a non-invasive beam profile instrument that can monitor the beam size and emittance throughout the accelerator cycle in the PS. The performance of the PS machine is expected to improve, based on requirements set by the LIU project, and the capabilities of the beam instrumentation must therefore be compatible with the improvements to enable monitoring and diagnostics of the machine. A beam profile monitor in the PS for the LIU must therefore fulfill the following requirements:

- Non-invasive measurements of the transverse beam profile.
- Continuous measurements throughout the beam cycle.
- Bunch-by-bunch beam profile measurement at 1 kHz.
- Turn-by-turn beam profile measurement for up to 100 turns.
- Reliable operation to minimise interventions and down time.

Beam gas ionisation profile monitors (IPMs) are based on ionisation of residual gas in the primary accelerator vacuum and allows for continuous monitoring of the beam throughout the cycle. The implementation of the PS-BGI was detailed in chapter 3 and is in many aspects similar to traditional IPMs, such as the use of parallel electric and magnetic fields to confine the ionisation electrons in a helical motion towards the detector. The novelty comes from the application of Timepix3 hybrid pixel detectors as the detector in the instrument, detailed in section 3.1.3, and the development of a radiation tolerant readout system, detailed in section 3.2. This is the first time pixel detectors have been used inside the primary vacuum of an operational accelerator.

The use of Timepix3 hybrid pixel detectors in the PS-BGI eliminates detector issues seen in traditional IPMs. The Timepix3 enables the detection of individual ionisation electrons for the first time in an IPM as shown in section 3.4, without using a micro-channel plate (MCP) for amplification. This leads to smaller point

spread compared to traditional MCP based detectors and allows for an independent beam profile measurement without requiring cross-calibration with other instruments. Individual ionisation electrons can be counted with the Timepix3 and Poisson statistics can therefore be applied to analyse the ionisation electron events in a methodical way as shown in chapter 4. Additionally, the 1.5625 ns time-of-arrival (ToA) timestamp resolution of each event allows for a precise determination of when each ionisation electron was detected by a pixel.

The Timepix3 event includes time-over-threshold (ToT) information that indicates the length of the input signal pulse in the pixel and is correlated to the amount of energy deposited in the sensor by the incoming particle. The ToT and ToA information in combination with the size and shape of event clusters enables efficient separation of ionisation electron events from background particle events as shown in section 4.2.2. This leads to the unique capability of a Timepix3 based IPM to distinguish between different particles and provides a more accurate determination of the ionisation electrons' transverse positions and improves the reconstruction of the beam profile compared to traditional IPMs.

Section 4.3 detailed a method to combine ionisation electron events in a fixed period of time to construct a beam profile, from which a RMS beam size can be calculated or a Gaussian model can be fitted using the maximum likelihood estimation method. A goodness of fit test compares the likelihood value of the measured profile against generated test profiles from a Gaussian distribution. Measurements on LHC INDIV cycles with the PS-BGI presented in section 5.1 indicated that the measured profiles were consistent with a Gaussian beam profile and the beam emittance was therefore calculated. The emittance values from the PS-BGI were in agreement with operational wire scanner measurements on the same cycle.

The 1.5625 ns time resolution of the Timepix3 is less than the 25 ns minimum spacing between two bunches in the PS and facilities bunch-by-bunch beam profile measurements. A measurement with the PS-BGI in section 5.3.1 showed that by

integrating the events from each bunch separately over 435 turns, bunch-by-bunch profile measurements at a rate of 1 kHz were achieved. The turn-by-turn measurements presented in section 5.3.2 demonstrated a method to increase the ionisation electron yield by sublimating an ion pump, which caused a short localized pressure spike. The yield increased from an expected 3 per bunch based on residual gas to 70 due to the sublimation, which made turn-by-turn beam profile measurements achievable. Injection mismatch into the PS was observed with the turn-by-turn measurements detailed in section 5.4, which confirmed the beam size oscillation and growth also observed from measurements with the invasive SEM grid profile monitor. With a gas injection system, foreseen to be installed with the operational PS-BGI instrument at the end of 2020, turn-by-turn measurements from injection will be possible on a regular basis, without disturbing the beam or damaging the instrument.

The beam gas ionisation profile monitor based on the Timepix3 hybrid pixel detectors presented in this thesis project is the only monitor to date that can provide non-invasive and continuous measurements of the beam throughout a full accelerator cycle in the PS. Beam dynamic studies presented in chapter 5 have demonstrated the capabilities of this new instrument. The non-invasive measurements using the Timepix3 based PS-BGI instrument will provide for new beam dynamics insights and serve as input to the tuning of the accelerator to reach the goals of the LIU project and for future improvements of the beams produced by the PS accelerator at CERN.

Bibliography

[1] E. Mobs, "The CERN accelerator complex - August 2018," Aug. 2018, general Photo, available at https://cds.cern.ch/record/2636343. [Online]. Available: https://cds.cern.ch/record/2636343

[2] G. Apollinari, I. Béjar Alonso *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. [Online]. Available: https://cds.cern.ch/record/2116337

[3] H. Damerau, A. Funken *et al.*, "LHC Injectors Upgrade, Technical Design Report, Vol. I: Protons," Tech. Rep. CERN-ACC-2014-0337, Dec. 2014. [Online]. Available: https://cds.cern.ch/record/1976692

[4] O. R. Sander, "Transverse emittance: Its definition, applications, and measurement," vol. 212, no. 1, 1990, pp. 127–155. [Online]. Available: https://doi.org/10.1063/1.39706

[5] G. Baud, B. Dehning *et al.*, "Performance Assessment of Wire-Scanners at CERN," in *Proceedings of IBIC2013*, no. CERN-ACC-2013-0308, Sep. 2013, p. 4 p. [Online]. Available: https://cds.cern.ch/record/1638354

[6] C. Field, D. McCormick *et al.*, "Wire breakage in SLC wire profile monitors," vol. 451, no. 1, 1998, pp. 440–445. [Online]. Available: https://doi.org/10.1063/1.57029

[7] F. Roncarolo and B. Dehning, "Transverse emittance blow-up due to the operation of wire scanners, analytical predictions and measurements," in *Proceedings of the 2005 Particle Accelerator Conference*. IEEE, 2005, pp. 437–439. [Online]. Available: https://cds.cern.ch/record/872259

[8] M. Fraser, Y. Dutheil *et al.*, "Matching studies between CERN PSB and PS through multi-turn beam profile acquisitions," in *Proc. 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019*, ser. International Particle Accelerator Conference, no. 10, June 2019, paper WEPMP025, pp. 2367–2370. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2019-WEPMP025

[9] F. Roncarolo, J. Allica Santamaria *et al.*, "Beam Instrumentation for the CERN LINAC4 and PSB Half Sector Test," in *8th Int. Particle Accelerator Conf.(IPAC'17), Copenhagen, Denmark, 14-19 May, 2017*, 2017, p. MOPAB120. 4 p. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2017-MOPAB120

[10] W. H. DeLuca, "Beam Detection Using Residual Gas Ionization," *IEEE Transactions on Nuclear Science*, vol. 16, no. 3, pp. 813–822, 1969. [Online]. Available: https://cds.cern.ch/record/888895

[11] C. D. Johnson and L. Thorndahl, "The CPS Gas-Ionization Beam Scanner," *IEEE Transactions on Nuclear Science*, vol. 16, no. 3, pp. 909–913, 1969. [Online]. Available: https://doi.org/10.1109/TNS.1969.4325399

[12] O. Keller, B. Dehning *et al.*, "Feasibility Study of a Novel, Fast Read-out System for an Ionization Profile Monitor Based on a Hybrid Pixel Detector," in *Proceedings of HB2014*, 2015, p. MOPAB41. 5 p. [Online]. Available: http://accelconf.web.cern.ch/HB2014/papers/mopab41.pdf

[13] R. Klanner, "Silicon detectors," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 235, no. 2, pp. 209 – 215, 1985. [Online]. Available: https://doi.org/10.1016/0168-9002(85)90555-8

[14] E. Heijne, P. Jarron *et al.*, "The silicon micropattern detector: A dream?" *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 273, no. 2, pp. 615 – 619, 1988. [Online]. Available: https://doi.org/10.1016/0168-9002(88)90065-4

[15] F. Anghinolfi, P. Aspell *et al.*, "A 1006 element hybrid silicon pixel detector with strobed binary output," *IEEE transactions on nuclear science*, vol. 39, no. 4, pp. 654–661, 1992. [Online]. Available: https://cds.cern.ch/record/228361

[16] P. Delpierre, "A history of hybrid pixel detectors, from high energy physics to medical imaging," *Journal of Instrumentation*, vol. 9, no. 05, pp. C05 059–C05 059, May 2014. [Online]. Available: https://doi.org/10.1088/1748-0221/9/05/c05059

[17] R. Ballabriga, M. Campbell, and X. Llopart, "Asic developments for radiation imaging applications: The medipix and timepix family," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 878, pp. 10 – 23, 2018, radiation Imaging Techniques and Applications. [Online]. Available: https://doi.org/10.1016/j.nima.2017.07.029

[18] T. Poikela, J. Plosila *et al.*, "Timepix3: a 65K channel hybrid pixel readout chip with simultaneous ToA/ToT and sparse readout," *Journal of Instrumentation*, vol. 9, no. 05, pp. C05 013–C05 013, May 2014. [Online]. Available: https://iopscience.iop.org/article/10.1088/1748-0221/9/05/C05013

[19] W. Wong, J. Alozy *et al.*, "Introducing Timepix2, a frame-based pixel detector readout ASIC measuring energy deposition and arrival time,"

Radiation Measurements, vol. 131, p. 106230, 2020. [Online]. Available: https://doi.org/10.1016/j.radmeas.2019.106230

[20] W. Andrzej, *Beam dynamics in high energy particle accelerators*. World Scientific, 2014. [Online]. Available: https://doi.org/10.1142/p899

[21] E. Courant and H. Snyder, "Theory of the alternating-gradient synchrotron," *Annals of Physics*, vol. 3, no. 1, pp. 1 – 48, 1958. [Online]. Available: https://doi.org/10.1016/0003-4916(58)90012-5

[22] F. Antoniou, S. Albright *et al.*, "Transverse Emittance Studies at Extraction of the CERN PS Booster," in *Proc. 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019*, ser. International Particle Accelerator Conference, paper MOPTS087, pp. 1058–1061. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2019-MOPTS087

[23] CERN Optics Repository. Accessed: 2020-10-05. [Online]. Available: http://acc-models.web.cern.ch

[24] S. Gilardoni and D. Manglunki, *Fifty years of the CERN Proton Synchrotron: Volume 1*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2011. [Online]. Available: http://cds.cern.ch/record/1359959

[25] X. Huang, S. Y. Lee *et al.*, "Emittance measurement and modeling for the Fermilab Booster," *Phys. Rev. ST Accel. Beams*, vol. 9, p. 014202, Jan. 2006. [Online]. Available: https://doi.org/10.1103/PhysRevSTAB.9.014202

[26] P. Skowroński, M. Giovannozzi, and A. Huschauer, "Linear and Non-Linear Optics Measurements in PS using Turn-by-Turn BPM Data," in *Proc. 10th International Partile Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019*, ser. International Partile Accelerator

Conference, paper MOPTS102, pp. 1114–1117. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2019-MOPTS102

[27] E. Senes, J. Emery *et al.*, "Transverse Emittance Measurement in the CERN Proton Synchrotron in View of Beam Production for the High-Luminosity LHC," in *Proc. 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019*, ser. International Particle Accelerator Conference, paper MOPTS100, pp. 1106–1109. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2019-MOPTS100

[28] S. Hancock, M. Lindroos, and S. Koscielniak, "Longitudinal phase space tomography with space charge," *Phys. Rev. ST Accel. Beams*, vol. 3, p. 124202, Dec. 2000. [Online]. Available: https://doi.org/10.1103/PhysRevSTAB.3.124202

[29] F. F. Rieke and W. Prepejchal, "Ionization Cross Sections of Gaseous Atoms and Molecules for High-Energy Electrons and Positrons," *Phys. Rev. A*, vol. 6, pp. 1507–1519, Oct. 1972. [Online]. Available: https://doi.org/10.1103/PhysRevA.6.1507

[30] M. Sapinski, P. Forck *et al.*, "Ionization Profile Monitor Simulations - Status and Future Plans," in *5th International Beam Instrumentation Conference*, Barcelona, Spain, 2016, p. TUPG71. [Online]. Available: https://doi.org/10.18429/JACoW-IBIC2016-TUPG71

[31] C. Fischer and J. Koopman, "Ionisation profile monitor tests in the SPS," in *Proceedings of DIPAC 1999*, 1999, pp. 119–121. [Online]. Available: https://cds.cern.ch/record/398770

[32] A. Jansson, T. Fitzpatrick *et al.*, "The Tevatron Ionization Profile Monitors," vol. 868, no. 1, 2006, pp. 159–167. [Online]. Available: https://doi.org/10.1063/1.2401401

[33] J. Zagel, M. Alvarez *et al.*, "Third generation residual gas ionization profile monitors at Fermilab," in *Proceedings of IBIC2014*, 2015. [Online]. Available: https://accelconf.web.cern.ch/IBIC2014/papers/tupd04.pdf

[34] T. Giacomini, P. Forck *et al.*, "Ionization profile monitors-IPM GSI," in *Proceedings of DIPAC2011*, 2011, pp. 419–421. [Online]. Available: https://accelconf.web.cern.ch/DIPAC2011/papers/tupd51.pdf

[35] K. Satou, N. Hayashi *et al.*, "A prototype of residual gas ionization profile monitor for J-PARC RCS," vol. 6, 2006. [Online]. Available: http://accelconf.web.cern.ch/e06/papers/tupch065.pdf

[36] G. Fraser, "The gain, temporal resolution and magnetic-field immunity of microchannel plates," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 291, no. 3, pp. 595 – 606, 1990. [Online]. Available: https://doi.org/10.1016/0168-9002(90)90009-U

[37] K. Wittenburg, "Experience with the residual gas ionization beam profile monitors at the DESY proton accelerators," in *Proc. of 3rd European Particle Accelerator Conference (EPAC 92)*, vol. 24, no. 28.3, 1992, pp. 92–12. [Online]. Available: https://accelconf.web.cern.ch/e92/PDF/EPAC1992_1133.PDF

[38] H. Breede, H.-J. Grabosch, and M. Sachwitz, "A new compact Design of a three-dimensional Ionization Profile Monitor (IPM)," in *Proceedings of IBIC2013*, 2013. [Online]. Available: https://accelconf.web.cern.ch/IBIC2013/papers/wepf04.pdf

[39] R. Connolly, P. Cameron *et al.*, "A prototype ionization profile monitor for RHIC," in *Proceedings of the 1997 Particle Accelerator Conference (Cat. No. 97CH36167)*, vol. 2. IEEE, 1997, pp. 2152–2154. [Online]. Available: https://doi.org/10.1109/PAC.1997.751138

[40] H. Harada, K. Yamamoto, and M. Yoshimoto, "Upgrade of ionization profile monitor (IPM) in the J-PARC 3-GeV RCS," *IPAC*, vol. 12, p. 840, 2012. [Online]. Available: https://accelconf.web.cern.ch/IPAC2012/papers/moppr029.pdf

[41] Y. Sato, K. Agari *et al.*, "Development of residual gas ionization profile monitor for high intensity proton beams," in *IEEE Nuclear Science Symposium Conference Record, 2005*, vol. 2. IEEE, 2005, pp. 1043–1046. [Online]. Available: https://doi.org/10.1109/NSSMIC.2005.1596431

[42] D. Turecek, J. Jakubek *et al.*, "Application of timepix3 based cdte spectral sensitive photon counting detector for pet imaging," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 895, pp. 84 – 89, 2018. [Online]. Available: https://doi.org/10.1016/j.nima.2018.04.007

[43] B. Bergmann, P. Azzarello *et al.*, "Detector response and performance of a 500 um thick GaAs attached to Timepix3 in relativistic particle beams," *Journal of Instrumentation*, vol. 15, no. 03, pp. C03 013–C03 013, Mar. 2020. [Online]. Available: https://doi.org/10.1088/1748-0221/15/03/c03013

[44] L. Rossi, P. Fischer *et al.*, *Pixel detectors: From fundamentals to applications*. Springer Science & Business Media, 2006. [Online]. Available: https://doi.org/10.1007/3-540-28333-1

[45] C. J. Z. M. Berger, M.J. and J. Chang. (2005) ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.2.3). National Institute of Standards and Technology, Gaithersburg, MD. Accessed: 2020-05-27. [Online]. Available: http://physics.nist.gov/Star

[46] M. A. Green, "Intrinsic concentration, effective densities of states, and effective mass in silicon," *Journal of Applied Physics*, vol. 67, no. 6, pp. 2944–2954, 1990. [Online]. Available: https://doi.org/10.1063/1.345414

[47] B. Maximilien, "Timepix3 Photo, CERN Knowledge Transfer," Feb. 2017, general Photo, available at https://cds.cern.ch/record/2253263. [Online]. Available: https://cds.cern.ch/record/2253263

[48] M. D. Gaspari, J. Alozy *et al.*, "Design of the analog front-end for the timepix3 and smallpix hybrid pixel detectors in 130 nm CMOS technology," *Journal of Instrumentation*, vol. 9, no. 01, pp. C01 037–C01 037, Jan. 2014. [Online]. Available: https://doi.org/10.1088/1748-0221/9/01/c01037

[49] A. Miarnau Marin, G. Schneider, and R. Veness, "Development and test of a rectangular CERN ConFlat-type flange," *Vacuum*, vol. 121, pp. 202 – 206, 2015. [Online]. Available: https://doi.org/10.1016/j.vacuum.2015.08.018

[50] S. Levasseur, B. Dehning *et al.*, "Development of a rest gas ionisation profile monitor for the CERN Proton Synchrotron based on a Timepix3 pixel detector," *Journal of Instrumentation*, vol. 12, no. 02, pp. C02 050–C02 050, Feb. 2017. [Online]. Available: https://doi.org/10.1088/1748-0221/12/02/c02050

[51] J. Storey, D. Bodart *et al.*, "Development of an Ionization Profile Monitor Based on a Pixel Detector for the CERN Proton Synchrotron," in *Proc. of International Beam Instrumentation Conference (IBIC2015), Melbourne, Australia, 13-17 September 2015*, ser. International Beam Instrumentation Conference, no. 4, JACoW. Geneva, Switzerland: JACoW, Jan. 2016, paper TUPB059, pp. 470–473. [Online]. Available: https://doi.org/10.18429/JACoW-IBIC2015-TUPB059

[52] K. Satou, J. W. Storey *et al.*, "A novel field cage design for the CPS IPM and systematic errors in beam size and emittance," *Journal of Physics: Conference Series*, vol. 1067, p. 072008, Sep. 2018. [Online]. Available: https://doi.org/10.1088/1742-6596/1067/7/072008

[53] D. Bodart, "CERN ATS/NOTE/2016/001," 2016, internal report.

[54] J. Jakubek, M. Jakubek *et al.*, "Large area pixel detector WIDEPIX with full area sensitivity composed of 100 Timepix assemblies with edgeless sensors," *Journal of Instrumentation*, vol. 9, no. 04, pp. C04 018–C04 018, Apr. 2014. [Online]. Available: https://doi.org/10.1088/1748-0221/9/04/c04018

[55] S. Levasseur, W. Bertsche *et al.*, "Time-Resolved Transverse Beam Profile Measurements with a Rest Gas Ionisation Profile Monitor Based on Hybrid Pixel Detectors," in *Proc. 9th International Particle Accelerator Conference (IPAC'18), Vancouver, BC, Canada, April 29-May 4, 2018*, ser. International Particle Accelerator Conference, no. 9. JACoW Publishing, June 2018, paper WEPAL075, pp. 2361–2364. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2018-WEPAL075

[56] J. Storey, W. Bertsche *et al.*, "First Results From the Operation of a Rest Gas Ionisation Profile Monitor Based on a Hybrid Pixel Detector," in *Proc. of International Beam Instrumentation Conference (IBIC'17), Grand Rapids, MI, USA, 20-24 August 2017*, ser. International Beam Instrumentation Conference, no. 6. JACoW, Mar. 2018, paper WE2AB5, pp. 318–322. [Online]. Available: https://doi.org/10.18429/JACoW-IBIC2017-WE2AB5

[57] S. Michelis, B. Allongue *et al.*, "DC-DC converters in 0.35 μm CMOS technology," *Journal of Instrumentation*, vol. 7, no. 01, pp. C01 072–C01 072, Jan. 2012. [Online]. Available: https://doi.org/10.1088/1748-0221/7/01/c01072

[58] CERN. (2014) Radiation and magnetic field tolerant 10W DC/DC converter module. Accessed: 2020-10-06. [Online]. Available: https://espace.cern.ch/project-DCDC-new/Shared%20Documents/FEAST2Mod_Datasheet_gb2016.pdf

[59] M. B. Marin, A. Boccardi *et al.*, "The Giga Bit Transceiver based Expandable Front-End (GEFE)—a new radiation tolerant acquisition system for beam instrumentation," *Journal of Instrumentation*, vol. 11, no. 02, pp. C02 062–C02 062, Feb. 2016. [Online]. Available: https://doi.org/10.1088/1748-0221/11/02/c02062

[60] P. Moreira, R. Ballabriga *et al.*, "The GBT Project," in *Proceedings of the Topical Workshop on Electronics for Particle Physics*, 2009. [Online]. Available: https://cds.cern.ch/record/1235836

[61] M. J. Wirthlin, H. Takai, and A. Harding, "Soft error rate estimations of the Kintex-7 FPGA within the ATLAS Liquid Argon (LAr) Calorimeter," *Journal of Instrumentation*, vol. 9, no. 01, pp. C01 025–C01 025, Jan. 2014. [Online]. Available: https://doi.org/10.1088/1748-0221/9/01/c01025

[62] L. van Harten, R. Jordans, and H. Pourshaghaghi, "Necessity of fault tolerance techniques in Xilinx Kintex 7 FPGA devices for space missions: A case study," in *2017 Euromicro Conference on Digital System Design (DSD)*. IEEE, 2017, pp. 299–306. [Online]. Available: https://doi.org/10.1109/DSD.2017.45

[63] A. Caratelli, S. Bonacini *et al.*, "The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments," *Journal of Instrumentation*, vol. 10, no. 03, pp. C03 034–C03 034, Mar. 2015. [Online]. Available: https://doi.org/10.1088/1748-0221/10/03/c03034

[64] H. Sandberg, "Radiation Hardened System Design with Mitigation and Detection in FPGA," Master's thesis, 2016. [Online]. Available: http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-132942

[65] M. B. Marin, S. Baron *et al.*, "The GBT-FPGA core: features and challenges," *Journal of Instrumentation*, vol. 10, no. 03, pp. C03 021–C03 021, Mar. 2015. [Online]. Available: https://doi.org/10.1088/1748-0221/10/03/c03021

[66] C. G. Larrea, K. Harder *et al.*, "IPbus: a flexible Ethernet-based control system for xTCA hardware," *Journal of Instrumentation*, vol. 10, no. 02, pp. C02 019–C02 019, Feb. 2015. [Online]. Available: https://doi.org/10.1088/1748-0221/10/02/c02019

[67] Protocol Buffers. Accessed: 2020-02-21. [Online]. Available: https://developers.google.com/protocol-buffers

[68] X. Llopart, J. Alozy *et al.*, "Study of low power front-ends for hybrid pixel detectors with sub-ns time tagging," *Journal of Instrumentation*, vol. 14, no. 01, pp. C01 024–C01 024, Jan. 2019. [Online]. Available: https://doi.org/10.1088/1748-0221/14/01/c01024

[69] M. Ester, H.-P. Kriegel *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *Proceedings of KDD-96*, vol. 96, no. 34, 1996, pp. 226–231. [Online]. Available: http://www.aaai.org/Library/KDD/1996/kdd96-037.php

[70] S. Kumar, B. A. Reshi, and R. Varma, "Comparison of Silicon, Germanium, Gallium Nitride, and Diamond for using as a detector material in experimental high energy physics," *Results in Physics*, vol. 11, pp. 461 – 474, 2018. [Online]. Available: https://doi.org/10.1016/j.rinp.2018.08.045

[71] J. Jakubek, "Precise energy calibration of pixel detector working in time-over-threshold mode," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 633, pp. S262 – S266, 2011, 11th International Workshop on Radiation Imaging Detectors (IWORID). [Online]. Available: https://doi.org/10.1016/j.nima.2010.06.183

[72] J. Coupard, H. Damerau *et al.*, "LHC Injectors Upgrade, Technical Design Report, Vol. II: Ions," Tech. Rep. CERN-ACC-2016-0041, Apr. 2016. [Online]. Available: https://cds.cern.ch/record/2153863

[73] M. Sapinski. Feasibility study of BGI detector for PS. Accessed: 2020-04-03. [Online]. Available: https://edms.cern.ch/document/1396633

[74] U. Fano, "Ionizing Collisions of Very Fast Particles and the Dipole Strength of Optical Transitions," *Phys. Rev.*, vol. 95, pp. 1198–1200, Sep. 1954. [Online]. Available: https://doi.org/10.1103/PhysRev.95.1198

[75] M. INOKUTI, "Inelastic Collisions of Fast Charged Particles with Atoms and Molecules—The Bethe Theory Revisited," *Rev. Mod. Phys.*, vol. 43, pp. 297–347, July 1971. [Online]. Available: https://doi.org/10.1103/RevModPhys.43.297

[76] F. James and M. Roos, "MINUIT-a system for function minimization and analysis of the parameter errors and correlations," *Comput. Phys. Commun.*, vol. 10, no. 6, pp. 343–67, 1975. [Online]. Available: http://cds.cern.ch/record/873119

[77] iminuit, MINUIT from Python - Fitting like a boss. Accessed: 2019-10-24. [Online]. Available: https://github.com/scikit-hep/iminuit

[78] probfit, Cost function builder. For fitting distributions. Accessed: 2019-10-24. [Online]. Available: https://github.com/scikit-hep/probfit

[79] L. Lyons, *A practical guide to data analysis for physical science students*. Cambridge University Press, 1991. [Online]. Available: https://doi.org/10.1017/CBO9781139170321

[80] S. Ahn and J. A. Fessler, "Standard errors of mean, variance, and standard deviation estimators," *EECS Department, The University of Michigan*, pp. 1–2, 2003. [Online]. Available: https://web.eecs.umich.edu/~fessler/papers/files/tr/stderr.pdf

[81] P. Zisopoulos, M. Gąsior *et al.*, "Fast Bunch by Bunch Tune Measurements at the CERN PS," in *Proc. of International Particle Accelerator Conference (IPAC'17),Copenhagen, Denmark, 14-19 May, 2017*, ser. International Particle Accelerator Conference, no. 8, May 2017, paper MOPAB122, pp. 415–418. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2017-MOPAB122

[82] M. Benedikt, C. Carli *et al.*, "Injection Matching Studies using Turn by Turn Beam Profile Measurements in the CERN PS," no. CERN-PS-2001-052-BD, p. 4 p, July 2001. [Online]. Available: https://cds.cern.ch/record/510662

[83] S. Aumon, S. Gilardoni, and M. Martini, "Study of Beam Losses at Transition Crossing at the CERN PS," in *11th European Particle Accelerator Conference (EPAC08), Genoa, Italy*, no. CERN-AB-2008-015, Aug. 2008, p. 4 p. [Online]. Available: https://cds.cern.ch/record/1122738

[84] M. Migliorati, S. Aumon *et al.*, "Instability studies at the CERN Proton Synchrotron during transition crossing," *Phys. Rev. Accel. Beams*, vol. 21, p. 120101, Dec. 2018. [Online]. Available: https://doi.org/10.1103/PhysRevAccelBeams.21.120101

[85] K. Takayama, D. Arakawa *et al.*, "Microwave Instability at Transition Crossing: Experiments and a Proton-Klystron Model," *Phys. Rev. Lett.*, vol. 78, pp. 871–874, Feb. 1997. [Online]. Available: https://doi.org/10.1103/PhysRevLett.78.871

[86] P. Hopchev, M. Kuhn *et al.*, "A Beam Gas Vertex detector for beam size measurement in the LHC," in *Proc. 5th International Particle Accelerator Conference (IPAC'14), Dresden, Germany, June 15-20, 2014*, ser. International Particle Accelerator Conference, no. 5. JACoW, July 2014, paper THPME175, pp. 3680–3683. [Online]. Available: https://doi.org/10.18429/JACoW-IPAC2014-THPME175

[87] A. Alexopoulos, C. Barschel *et al.*, "Noninvasive lhc transverse beam size measurement using inelastic beam-gas interactions," *Phys. Rev. Accel. Beams*, vol. 22, p. 042801, Apr. 2019. [Online]. Available: https://doi.org/10.1103/PhysRevAccelBeams.22.042801

[88] M. Sapinski, W. Andreazza *et al.*, "The first experience with LHC beam gas ionization monitor," in *Proceedings of IBIC2012, Tsukuba, Japan*, no. CERN-ATS-2012-286, Oct. 2012. [Online]. Available: https://cds.cern.ch/record/1495085

[89] D. Vilsmeier, B. Dehning, and M. Sapinski, "Investigation of the Effect of Beam Space-charge on Electron Trajectories in Ionization Profile Monitors," in *Proceedings of HB2014, East-Lansing, MI, USA*, 2015, p. MOPAB42. 5 p. [Online]. Available: https://cds.cern.ch/record/2158997

[90] M. Sapinski, R. Singh *et al.*, "Application of machine learning for the IPM-based profile reconstruction," in *Proc. 61st ICFA Advanced Beam Dynamics Workshop (HB'18), Daejeon, Korea, 17-22 June 2018*, ser. ICFA Advanced Beam Dynamics Workshop, no. 61. JACoW

Publishing, July 2018, paper THA2WE02, pp. 410–415. [Online]. Available: https://doi.org/10.18429/JACoW-HB2018-THA2WE02

[91] D. Vilsmeier, P. Forck, and M. Sapinski, "A modular Application for IPM Simulations," in *Proc. of International Beam Instrumentation Conference (IBIC'17), Grand Rapids, MI, USA, 20-24 August 2017*, ser. International Beam Instrumentation Conference, no. 6. JACoW, Mar. 2018, paper WEPCC07, pp. 355–358. [Online]. Available: https://doi.org/10.18429/JACoW-IBIC2017-WEPCC07

[92] J. R. Zagel, D. J. Harding *et al.*, "Permanent magnet ion profile monitor at the Fermilab main injector," in *PACS2001. Proceedings of the 2001 Particle Accelerator Conference (Cat. No.01CH37268)*, vol. 2, 2001, pp. 1303–1305 vol.2. [Online]. Available: https://doi.org/10.1109/PAC.2001.986661

[93] D. Vilsmeier, M. Sapinski, and R. Singh, "Space-charge distortion of transverse profiles measured by electron-based ionization profile monitors and correction methods," *Phys. Rev. Accel. Beams*, vol. 22, p. 052801, May 2019. [Online]. Available: https://doi.org/10.1103/PhysRevAccelBeams.22.052801

[94] N. Rezzak, J.-J. Wang *et al.*, "TID and SEE Characterization of Microsemi's 4th Generation Radiation Tolerant RTG4 Flash-Based FPGA," in *2015 IEEE Radiation Effects Data Workshop (REDW)*. IEEE, 2015, pp. 1–6. [Online]. Available: https://doi.org/10.1109/REDW.2015.7336739

[95] K. Maragos, V. Leon *et al.*, "Evaluation Methodology and Reconfiguration Tests on the New European NG-MEDIUM FPGA," in *2018 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. IEEE, 2018, pp. 127–134. [Online]. Available: https://doi.org/10.1109/AHS.2018.8541492

[96] P. Leitao, S. Feger *et al.*, "Test bench development for the radiation Hard GBTX ASIC," *Journal of Instrumentation*, vol. 10, no. 01, pp.

C01 038–C01 038, Jan. 2015. [Online]. Available: https://doi.org/10.1088/1748-0221/10/01/c01038

[97] P. Moreira. (2016) The LpGBT project status and overview. ACES 2016 - Fifth Common ATLAS CMS Electronics Workshop for LHC Upgrades. Accessed: 2020-10-06. [Online]. Available: https://indico.cern.ch/event/468486/contributions/1144369

[98] X. Llopart. The design of the Timepix4 chip: a 230 kpixel and 4-side buttable chip with 200ps on-pixel time bin resolution and 15-bits of TOT energy resolution. Accessed: 2020-04-17. [Online]. Available: https://indico.cern.ch/event/788037/

[99] H. Kurino, K. Lee *et al.*, "Intelligent image sensor chip with three dimensional structure," in *International Electron Devices Meeting 1999. Technical Digest (Cat. No. 99CH36318)*. IEEE, 1999, pp. 879–882. [Online]. Available: https://doi.org/10.1109/IEDM.1999.824289

[100] K. Sakuma, P. S. Andry *et al.*, "3D chip-stacking technology with through-silicon vias and low-volume lead-free interconnections," *IBM Journal of Research and Development*, vol. 52, no. 6, pp. 611–622, 2008. [Online]. Available: https://doi.org/10.1147/JRD.2008.5388567