FFAGs and Synchrotrons for Proton Therapy: A Comparative Study

Thesis submitted in accordance with the requirements of
the University of Liverpool for the degree of Doctor in Philosophy by

Chris Edmonds

November 2015

Contents

Illustrations

List of Figures

List of Tables

Abstract

Accelerators play a key role in the delivery of radiotherapy for treatment of cancer and other medical conditions. Proton therapy has the benefit of more localised delivery of dose to deep seated tumour volumes in comparison to treatment using x-rays or electrons. The accelerators currently used for proton therapy are cyclotrons and synchrotrons, which each have certain advantages and disadvantages. It has been proposed that accelerators of a fixed field alternating gradient (FFAG) design may combine some of the advantages and avoid some of the disadvantages of the existing machines. This thesis looks at the use of synchrotrons as a benchmark for the delivery of proton therapy, and then at how FFAGs may improve upon treatment delivery. Particular attention is paid to the beam dynamics issues, including comparisons between simulations and experimental data taken with the EMMA non-scaling FFAG at Daresbury. The results of the comparisons show that simulation is able to predict the behaviour of a particle bunch in a real machine. The simulation tools are then used to evaluate the design of FFAGs incorporating resonant extraction techniques. In principle, resonant extraction could overcome some problems of kicker based extraction methods. The design study highlights technical challenges that would need to be overcome before resonant extraction could be implemented as a beneficial method for a proton therapy FFAG.

Acknowledgements

thesis

noun

1. A statement or theory that is put forward as a premise to be maintained or proved:
   'He vigorously defended his thesis on the causes of war.'
2. A penance for spending time doing research that you find interesting:
   'My thesis.'

I would like to thank the people who helped me create so much to write about, with special thanks to (in alphabetical order of first name) Andy Wolski, Bruno Muratori, David Kelliher, Kai Hock and Shinji Machida: you provided me with interesting challenges, good guidance and the time needed to follow up ideas.
The Cockcroft Institute has always been a very welcoming environment; my time in office S30 would not have been the same without the many new friends that I have made whilst here. Thanks to you all.
Finally, but by no means with the least amount of gratitude, I must thank my family and friends. You made this possible.

Chapter 1

Proton therapy and the role of accelerators

1.1 Introduction

Particle accelerators are widely used for radiotherapy for the treatment of cancer. In 1957 a 3 GHz linear accelerator (linac) was used for the first time to produce a beam of x-rays to treat a patient with an ocular tumour [1]. In the years that have followed much research has gone into enabling treatments which deliver the maximum possible radiation damage to a tumour, whilst minimising damage to healthy tissue. This research has led to better imaging techniques for pinpointing and targeting tumours within the body, better understanding of how different biological tissues react to radiation (allowing clinicians to avoid unnecessary damage to radiosensitive organs) and better treatment techniques (such as the shaping of radiation beams to match the profile of a tumour). However, for most patients who receive radiotherapy, the treatment still means a beam of high energy x-rays delivered by a 3 GHz linac in a single room treatment facility. It may be a testament to the efficiency of linacs in producing treatment beams of electrons and x-rays, that, fundamentally, little has changed in terms of clinical accelerator technology since 1957.

In 1946 the potential therapeutic benefits of using protons rather than x-rays were described [2]. Since then the uptake of proton therapy has been slow, in part due to the technical challenges of delivering treatment beams of protons. In order to treat a tumour at a depth of 30 cm, protons with a kinetic energy of 230 MeV are needed (≈ 10 times higher than the electron energy required by conventional treatments). As a result, circular accelerators, which allow the accelerating rf structures to be reused a number of times per acceleration cycle, have become the preferred solution. The mass of the proton, and limitations on the strength of magnetic fields that can be produced for steering a particle around a circular accelerator presently restrict just how compact these proton accelerators can be made. A hospital has to consider different factors depending on whether it wants to start offering proton treatments or treatments using electrons and

x-rays. The electron based facility (which can offer electron or x-ray therapies) usually consists of several treatment rooms, each containing an electron linac. The operation of each linac does not depend on what is happening in other treatment rooms. A proton based facility will also have several treatment rooms; however, due to the size and cost of the proton accelerator [3], one accelerator will typically serve all of the treatment rooms. The quality of treatment a proton therapy centre can offer is dependent upon the accelerator at its heart. To date, hospitals have used either a cyclotron or synchrotron accelerator, both of which offer distinct advantages and disadvantages. There is a demand for accelerator technologies that either help to make proton therapy more widely available, or improve upon the current standards of treatment. It has been proposed that accelerators of a fixed field alternating gradient (FFAG) design may be one way by which standards of treatment may be improved. This thesis evaluates the potential of FFAGs in the field of proton therapy, by making comparisons with the performance of a synchrotron accelerator.

1.2 Particles for cancer therapy

Late in the 19^{th} century, shortly after Röntgen discovered x-rays, it was first observed that radiation could have a damaging effect on tissue and that it could potentially be used in the treatment of cancer. Today it is understood that a successful radiation treatment means balancing the probability of controlling a tumour with the probability of causing the patient significant further harm due to a large exposure of healthy tissue to radiation. Careful choice of the particle used for radiotherapy is one of the ways through which the damage to a tumour can be maximised whilst minimising damage to healthy tissue. In evaluating the suitability of different particles for radiotherapy, it is often the relationship between absorbed dose and depth within water (a good approximation for tissue) that is considered. Here, the absorbed dose is the absorbed energy per unit mass and is measured in units of gray (Gy)[4]:

$$1\,\text{Gy} = 1\,\text{J/kg}.$$

1.2.1 Electrons and photons

Currently, most patients who receive radiotherapy are treated using either beams of electrons or x-rays. Both treatments are based upon the acceleration of electrons by a linear accelerator to energies of between 4 MeV and 25 MeV. When needed, a beam of x-rays can be produced by impacting the electrons upon a heavy metal target (e.g. tungsten), Bremsstrahlung radiation is emitted as the electrons are slowed within the target. Figure 1.1 shows the relationship between dose and depth for 4 MeV electrons and x-rays as well as 20 MeV x-rays. It can be seen that the dose from the electron beam penetrates a few centimetres of tissue before rapidly dropping off, a feature which makes electron beams well suited to treating tumours which are at, or close to, the surface of

a patient [5]. For treating deep-seated tumours, x-rays are the most commonly applied form of radiation. The notable features of the x-ray treatment are that the maximum dose occurs at a depth of several centimetres within the patient (with the depth of maximum dose increasing with beam energy), and that beyond the maximum dose, the dose then decreases exponentially with depth. As can be seen from Fig. 1.1, a potential problem with x-ray beam treatment is that, in the direction of the beam travel, the healthy tissue both before and after a tumour will receive a significant dose of radiation. The risk of complications from damage of normal tissue may be reduced by rotating the x-ray beam around a patient, whilst keeping the tumour volume at the axis of rotation, a process which reduces the maximum dose given to any one area of healthy tissue whilst maintaining a high dose within the tumour volume. However, in many cases, the ratio of the probability of eradicating a cancer to the probability of normal tissue complications (known as the therapeutic ratio) [6] may be improved through use of a different particle.

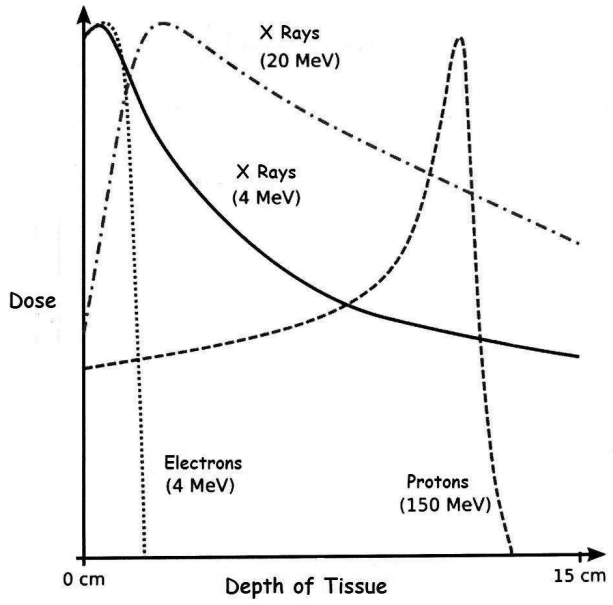

FIGURE 1.1: The relationship between the dose delivered to tissue and depth is dependent upon the particle used for therapy [7].

1.2.2 Protons and other hadrons

A charged hadron travelling through a material will have a finite range that is roughly proportional to the initial kinetic energy of the particle. As the particle travels it will lose energy to the material through inelastic scattering processes, with the rate of the energy loss of the particle being approximately inversely proportional to the kinetic energy of the particle. As the kinetic energy of the particle approaches zero, there is a sharp rise in the energy lost to the surrounding material. This sharp rise in energy loss just before

the maximum range of the proton is reached is known as the Bragg peak, and is the key advantage of using protons, or other ions (e.g. carbon ions), for therapy. The range $R\,(\mathrm{cm})$ of a proton in water is related to the proton kinetic energy $E_k\,(\mathrm{MeV})$ by [8]:

$$R \approx \alpha E_k^p, \tag{1.1}$$

where $\alpha = 0.0022\,\mathrm{cm\,MeV}^{-p}$ and $p = 1.77$.

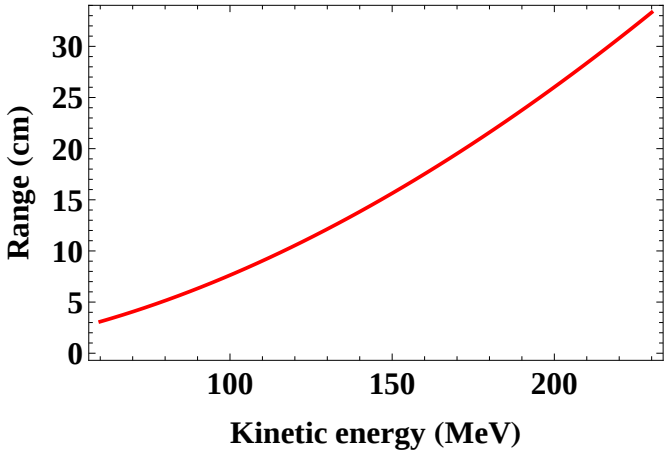

FIGURE 1.2: The approximate range (and hence Bragg peak depth) of a proton in water vs. proton kinetic energy calculated using Eq. 1.1. The energy of the proton beam is selected so that the Bragg peak coincides with a tumour.

If the Bragg peak is positioned at a tissue depth corresponding to the location of a tumour, then a relatively low dose is received by the healthy tissue before the tumour and little or no[1] dose after the tumour (Fig. 1.1). Having been proposed by Wilson in 1946, therapy based on use of the Bragg peak was first applied to a patient in 1960 [9].

For a treatment involving a beam of many particles, the width of the peak in the depth-dose plot depends on the type of charged particle used as well as the total range of the particles. The cause of this is range straggling, which is a result of the statistical variations in energy loss from each interaction of the treatment particle with the medium through which it is travelling. Range straggling means that each particle within a mono-energetic treatment beam will have a slightly different range. The range distribution is Gaussian, with a standard deviation approximately proportional to $R/\sqrt{A}$, where R is the nominal range of the particle, and A the particle mass number [10]. Using heavier ions can lead to a sharper Bragg peak, however, the accelerators required tend to be larger and more expensive than for protons. In this thesis, emphasis is placed on accelerators for proton therapy, although some of the designs discussed are also capable of delivering heavier ions.

[1]Some dose is present beyond the Bragg peak due to neutron production and the additional range of fragmented heavier ions. This exit dose is much lower than the entrance dose.

1.3 Beam requirements of proton therapy

Following diagnosis and an initial assessment of a tumour, a clinician will prescribe a radiation dose to the tumour volume. Typically this dose is measured in tens of gray, and is delivered in fractions of around 2 Gy per day over a number of weeks. Daily treatment times are minimised in order to maximize the patient throughput of a facility, and to reduce the error on dose distribution (which should not vary by more than $\pm 4\,\%$ throughout the tumour volume) due to patient movement; it is expected that the dose rate of a system should be at least 2 Gy/minute/litre which requires a beam current of approximately 0.5 nA [11].

The distal [2] conformation of dose to a tumour is dependent upon the energy spread of the extracted beam as well as scattering in materials before the treatment volume. For a given treatment kinetic energy, E_k, the energy spread of the beam should be less than $\Delta E_k/E_k = \pm 0.1\,\%$ at the point of extraction. The thin 'pencil'-like beams extracted from a proton accelerator typically have too small a cross sectional area and energy spread to deliver a uniform dose throughout an entire target tumour volume. Methods are required to spread the dose both laterally and longitudinally. The earliest treatments were based on a method known as passive scattering, whilst today, the best conformity of dose to a tumour volume is achieved using a method, first developed in the 1980's, known as active scanning.

Passive scattering

Passive scattering uses a treatment field that is shaped to conform as closely as possible to the target tumour volume. To create the treatment field, scattering foils are used to enlarge the pencil beam [12, 13] followed by collimators to shape the enlarged beam so that it has the same lateral profile as the tumour [14]. To ensure that as uniform a dose as possible is delivered to the tumour longitudinally, a spread-out Bragg peak must be formed by delivering protons that have different ranges within a patient (Fig. 1.3) [15]. For a fixed energy particle source, the treatment beam is first passed through a range shifter, consisting of a perspex block with a thickness that positions the Bragg peak at the deepest point of the target volume. Following the range shifter, the beam is passed through a range modulator, which is often a rotating perspex disk of stepped thickness [16], to spread the Bragg peak longitudinally through the full depth of the treatment volume.

Passive scattering allows the required dose to be delivered to a tumour volume within a very short period of time, which has benefits in dealing with tumour movement resulting from, for example, respiratory cycles. However, the disadvantages of passive scattering include the spread out Bragg peak being of fixed width across the tumour volume and range straggling within the range shifter and modulator, both of which lead

[2] Distal refers to the furthermost edge of a tumour as encountered by the beam.

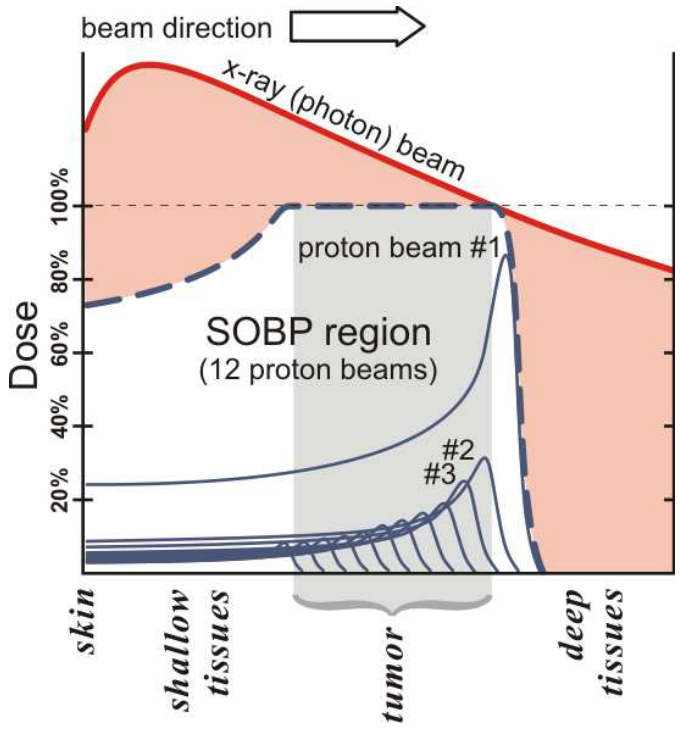

FIGURE 1.3: The Bragg peak of a mono-energetic beam is typically too narrow to give a uniform dose throughout a tumour volume. A spread-out Bragg peak (SOBP) is produced by applying proton beams that have different ranges within the patient. The dashed blue line in the above plot shows the dose distribution calculated by summing the contributions from the individual Bragg peaks (solid blue lines) [17]. The contribution to the SOBP by the entrance dose of protons means fewer shorter range protons are required than longer range protons in order to create a uniform dose distribution.

to a poorer conformation of dose to a tumour volume than is possible when using active scanning methods.

Active scanning

Active scanning [18] can involve visualising a tumour volume as being made up of smaller sub-volumes, which are referred to as voxels. The dose that should be given to each voxel in order to achieve a uniform dose throughout the tumour volume is then calculated [19]. Dose is delivered, voxel by voxel, by scanning the pencil beam laterally and longitudinally through the tumour. The lateral scanning is achieved by steering the beam with magnets, whilst longitudinal scanning is possible by varying the extraction energy from the accelerator or by passing the beam through a range modulator. For a static tumour, active scanning can give a much better conformation of dose to a tumour volume than passive scattering. However, passive scattering may still be preferred in the case of moving tumours (e.g. in the lung), for which the voxels will move significantly during irradiation, resulting in a poor dose distribution [20]. There have been a number of recent studies which have looked at improving the dose distribution to moving tumours when using active scanning techniques. These studies include applying gating (where dose delivery is paused during certain parts of a patients respiratory cycle), rapid scanning

TABLE 1.1: Typical requirements of proton therapy affecting accelerator design (the energy step size and dynamic beam extraction intensity are based on the performance of the Heidelberg Ion Therapy (HIT) synchrotron [26]).

Parameter	Requirement
Extraction energy range	60-250 MeV
Energy spread at extraction $\left(\frac{\Delta E_k}{E_k}\right)$	$\pm 0.1\%$
Energy step size	~ 0.7 MeV (1 mm steps)
Energy switching time	≤ 1 s
Dose rate	2 Gy/min/litre
Dynamic beam extraction intensity	$8 \times 10^7 - 2 \times 10^9$ pps
Stability at extraction point	± 1 mm or ± 1 mrad

(where a beam is swept through the tumour volume in a very short period of time) and multi-painting (where the total dose is given to a tumour volume over several scans by the particle beam). These methods are discussed in more detail in chapter 6.

1.4 Accelerators for proton therapy

To date, the accelerators constructed for proton therapy have been either cyclotron or synchrotron circular accelerators [21]. Linear alternatives have been proposed [22], and ultra compact linear accelerators, which could, in future, see proton therapy centres based upon single room treatment facilities (resembling those offered by electron and x-ray therapy), are being investigated [23, 24]. It has been proposed that a circular accelerator of FFAG design may help to improve upon the quality of proton treatments currently available [25]. In this section, the basic features and principles of circular accelerators are described. Ultra compact linear accelerators are not discussed further, as they are considered to be beyond the scope of this study.

1.4.1 Cyclotrons

A simple cyclotron consists of two D-shaped electrodes within a dipole field (Fig. 1.4). Cyclotrons are well suited to a hospital environment due to their relatively compact size and simplicity.

Charged particles are injected close to the centre of the cyclotron with some velocity towards one of the D's. As the magnetic field within the D is perpendicular to the direction of particle motion, the Lorentz force causes the particle to follow an arc-like path. Once the particle has travelled through 180°, it will then be accelerated as it encounters an electric field between the two D's. The greater particle velocity on entering the second D leads to an arc of greater radius being followed, and the combination of the increased velocity and path length mean that the time taken to travel through the first and second D's is approximately equal. In the time taken to traverse one D the electric field reverses direction, which is achieved by applying a rapidly oscillating radio frequency (rf) field. As a result the particle will be accelerated each time it crosses

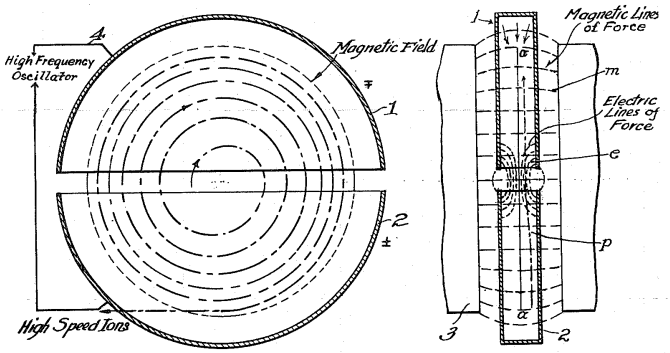

FIGURE 1.4: Diagram of the Lawrence cyclotron in the plane of motion (left) and showing the magnetic fields between the poles of the Ds (right).

the gap between D's. The acceleration cycle continues, with the path of the particle spiralling outwards as the number of interactions with the rf field increases. As the particle approaches the outer edge of the cyclotron, it is extracted. The acceleration process can be completed in approximately 10 μs from injection to extraction. Good dynamic intensity control of the beam may be achieved by monitoring the extracted beam and then adjusting the ion source accordingly.

The rapid acceleration cycle of a cyclotron is a result of the fact that the magnetic fields and rf frequency can remain constant. The structure and operation of a cyclotron mean that the extracted beam has a fixed energy. In order to position the spread-out Bragg peak over the tumour volume, range shifters and modulators must be used. Although these methods allow for very rapid shifting of the position of the Bragg peak, they also introduce the problem of range straggling for lower energy treatments.

1.4.2 Synchrotron

Unlike the cyclotron, synchrotrons (such as the example shown in Fig. 1.5) confine particles to a single path. The layout of a synchrotron can be approximately described by a regular polygon. Dipole magnets are located at the vertices of the polygon in order to bend the particles around a closed circuit. In the straight sections between the dipoles, there are rf cavities for particle acceleration, as well as additional magnets which focus the beam, controlling the beam size and preventing particles from being lost on the walls of the accelerator.

During the acceleration process, the momentum of each particle will increase. In a uniform magnetic field B a particle with momentum p and charge q will follow a circular trajectory of radius ρ, where

$$B\rho = \frac{p}{q}. \tag{1.2}$$

$B\rho$ is called the beam rigidity, and for a given charge is a measure of the momentum. Within a synchrotron, the magnetic field strengths must be ramped up during acceleration in order to keep the path of the beam and the focusing properties of the accelerator

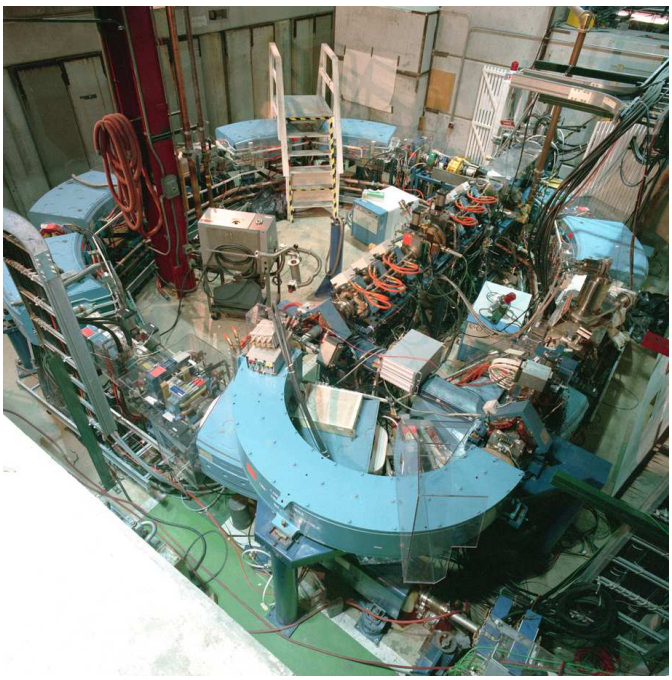

FIGURE 1.5: The synchrotron that is used to accelerate protons for therapy at the Loma Linda Medical Centre [27].

constant. Given the fixed path length, the time taken for a particle to make one revolution of the accelerator will change with the particle velocity during an acceleration cycle. Therefore, the frequency of the rf must also change to maintain synchronicity between the arrival of the particle bunch at a cavity and the desired phase of the rf.

A major advantage of the synchrotron for therapy is that the energy at which the particles are extracted can be varied easily from cycle to cycle, eliminating the need to degrade the beam by the use of range shifters and modulators. For treatments of tumours at small depths within a patient, this can mean better conformation of dose to a tumour volume than is the case when a cyclotron is used.

1.4.3 FFAG

In recent years, fixed field alternating gradient (FFAG) accelerators for hadron therapy have been the subject of a number of design studies [28–30]. It is claimed that accelerators of this design have the potential to offer the benefits of both cyclotrons and synchrotrons, by giving:

- variable energy extraction;
- high repetition rate, leading to rapid dose delivery;
- simple and stable operation due to fixed magnetic fields.

As with a synchrotron, FFAGs have a ring type design; however, like a cyclotron, the magnetic fields remain fixed during the acceleration process. An FFAG is able to confine particles to a ring by having magnetic fields that vary with increasing radius; as

the momentum of a particle increases it follows paths in regions of stronger magnetic fields. FFAGs may be split into two categories, scaling and non-scaling. A scaling FFAG [31] has magnetic fields which vary radially within a plane according to:

$$B_z = B_{z,0} \left(\frac{R}{R_0} \right)^k . \tag{1.3}$$

Here $B_{z,0}$ is the magnetic field perpendicular to the plane at a reference radius, R_0, R is the radial distance and k is known as the scaling index. The term 'scaling' refers to the fact that a high momentum particle will follow a path that is of the same shape but scaled up in size when compared to the path of a low momentum particle; the scaling index determines the proximity of the paths of particles at different momenta. Within a non-scaling FFAG, the field does not obey Eq. 1.3, and the shape of the path of a particle can vary with changing momentum. In both cases the magnetic field varies azimuthally so as to introduce alternating gradient focusing (alternating gradient schemes for FFAG accelerators are discussed further in chapter 2). For the scaling FFAG, the strength of the focusing is independent of the particle momentum, which ensures that the motion of the particles in the beam remains stable throughout an acceleration cycle. However, for the non-scaling FFAG, the strength of focusing can vary significantly with momentum (depending upon the field profile chosen) and there may be times during an acceleration cycle where, if the energy was fixed, the particle motion would be unstable. One advantage of the non-scaling FFAG is that it can be built using magnets that are less complicated than the magnets required for a scaling FFAG. The non-scaling FFAG also allows for the design of accelerators that have a reduced radial aperture and large acceptance when compared to a scaling FFAG (an accelerator with these qualities is discussed in chapter 4).

1.5 Thesis Overview

This thesis is organised as follows:

- Chapter 1 introduces this thesis.
- Chapter 2 provides context by detailing an existing proton therapy facility as well as FFAG design concepts.
- Chapter 3 describes the tracking methods used for characterising the accelerators in computer simulations.
- Chapter 4 looks at the limitations of the design tools used in predicting particle behaviour in a given machine.
- Chapter 5 details the optimization of the synchrotron design which will be used as a basis for comparison with a FFAG accelerator for proton therapy delivery.

- Chapter 6 describes the FFAG design that will be discussed, and compared with the synchrotron design described in chapter 5.
- Chapter 7 gives a summary and conclusion from the studies detailed in previous chapters.

Chapter 2

Beam dynamics and accelerator design

The momentum of a particle of charge q and velocity $\mathbf{v}$ travelling through electric ($\mathbf{E}$) and magnetic ($\mathbf{B}$) fields is influenced by the Lorentz force,

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

where $\mathbf{p}$ is the particle momentum as a function of time (t). This thesis will look at how particles are guided by magnetic fields and accelerated by electric fields within accelerators that deliver particle beams suitable for treating cancer. In this section, the physics of the relevant beam dynamics within accelerators will be introduced and given context through examples of existing accelerator designs.

2.1 Particles for acceleration

The charged particles encountered within the tracking studies in this thesis are electrons (chapter 4) and protons (chapters 5 and 6). The momentum of a particle is given by:

$$\mathbf{p} = m\mathbf{v} = \gamma m_0 \mathbf{v},$$

with γ being the relativistic Lorentz factor and m_0 the rest mass of the particle. When defining the global properties of an electron beam within this thesis, the nominal momentum of a single electron within the beam is given. This is in keeping with literature on the EMMA non-scaling FFAG (discussed in chapter 4), for which electrons are travelling at a speed greater than 0.99c. If the energy (E) of a particle is given by the rest energy (E_0) and kinetic energy (E_k) of the particle:

$$E = E_0 + E_k,$$

then the energy of the particle may be written as:

$$E = mc^2 = \gamma m_0 c^2,$$

where c is the speed of light. The kinetic energy of a particle can therefore be written:

$$E_k = (\gamma - 1)m_0 c^2.$$

When defining the global properties of a proton beam within this thesis, the nominal kinetic energy of a single proton within the beam is given. This follows the convention generally used in the literature on proton therapy. The energy, momentum and rest mass of a particle are related by:

$$E = \sqrt{(pc)^2 + (m_0 c^2)^2}.$$

The energy of particle beams are conventionally expressed in eV, which is the change in energy of an electron as it crosses a potential of 1 V. Particle momentum is given in units of eV/c and the mass in units of eV/c^2. For ultra-relativistic particles, where $pc \gg mc^2$, values for E_k and pc approach that of E. When dealing with the dynamics of particles within a bunch, it is often convenient to work with the momentum deviation of individual particles from a nominal momentum for the bunch. For both electrons and protons, the momentum of an individual particle within a beam is expressed as the fractional offset in its momentum (p) from the nominal momentum (p_0), $\delta = \frac{p-p_0}{p_0} = \frac{\Delta p}{p_0}$.

2.2 Fundamentals of motion within a circular accelerator

Using a Cartesian coordinate system, a particle given some velocity in the zx plane will remain in the same plane and follow a circular path if there is a magnetic field directed along the y axis [1]. The radius of this path is given by,

$$\rho = \frac{p}{qB_y}. \tag{2.2}$$

Figure 2.1 illustrates this case. A disc within the zx plane has been taken, within the disc there is a uniform magnetic field directed along the y axis. A charged particle given the appropriate starting conditions (position and velocity) will orbit around the centre of the disc (the path of such a particle is marked red in Fig. 2.1). Taking two more particles of the same energy, and giving one a slightly different starting velocity (Fig. 2.1a) and the second a different starting location (Fig. 2.1b), it is seen that the orbits of these two extra particles can be said to oscillate around the orbit of the original particle, but that in all cases each particle has the same location and velocity at the end of one orbit as it had at the start. The number of oscillations of a particle around the original path per orbit is called the betatron tune, which for pure dipole field is 1.

[1] In the absence of synchrotron radiation.

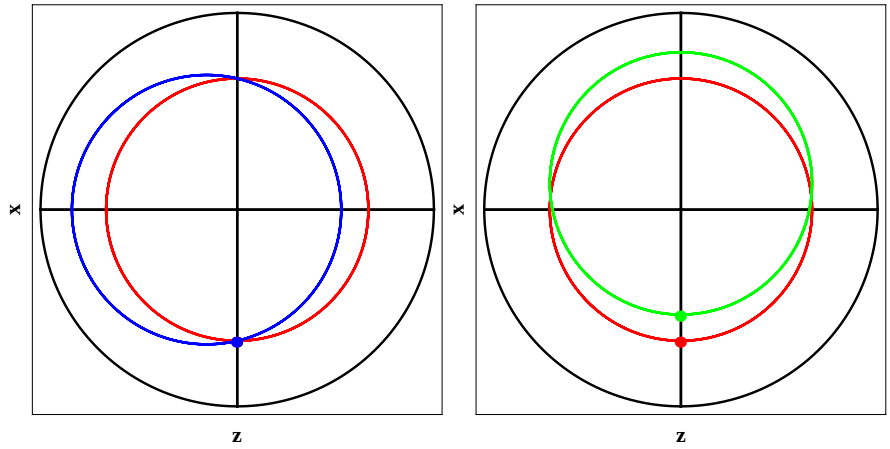

(a) Momentum offset from equilibrium orbit.

(b) Spatial offset from equilibrium orbit.

FIGURE 2.1: A particle with velocity in the zx plane travels on a disc that has a magnetic field directed along the y axis, the force experienced by the particle causes it to orbit around the centre of the plot (red path). If the initial velocity (blue path) or position (green path) of the particle is changed, then these new paths are seen to oscillate around the original path.

The disc is now divided into quarters, and the quarters of the disc are moved outwards (as illustrated in Fig. 2.2) so that they are separated by regions with no magnetic field: these field free regions are referred to as drift spaces. Figure 2.2a shows a path for which a particle will have the same location and velocity at the end of an orbit as at the start; this path is referred to as the closed (or equilibrium) orbit. Particles following any other path are seen to oscillate around the closed orbit. However, unlike in Fig. 2.1, the particles do not return to their starting conditions after one orbit, but instead do so after three. The betatron tune in the case shown in Fig. 2.2 is 4/3.

These two examples are crude representations of the arrangement of magnets within a cyclotron (Fig. 2.1) and synchrotron (Fig. 2.2). For both arrangements, the return of the particles to their starting conditions after a given number of orbits is evidence of the focusing properties of a uniform magnetic field, there is clearly a range of initial conditions within the zx plane for which a particle will orbit indefinitely. The particle sources used for accelerators emit particles which have a spread of initial locations, velocities and energies; focusing methods are important in ensuring that as many particles as possible survive to the end of the acceleration cycle without being lost to the walls of the accelerator.

Equation 2.2 tells us that the radius of an arc subtended by a charged particle that is travelling through a magnetic field increases with momentum, the result of which is a closed orbit that is dependent on momentum. The change in closed orbit with the fractional offset in momentum, δ, is referred to as dispersion. The momentum of the particle tracked through the split dipole is now increased, and the effects of dispersion can be seen in Fig. 2.3a. The betatron tune of the high momentum particle is not 4/3, but is instead either an irrational number or expressed by a fraction with a denominator much

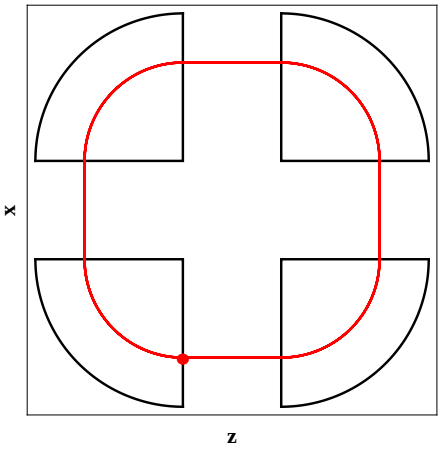

(a) Equilibrium orbit.

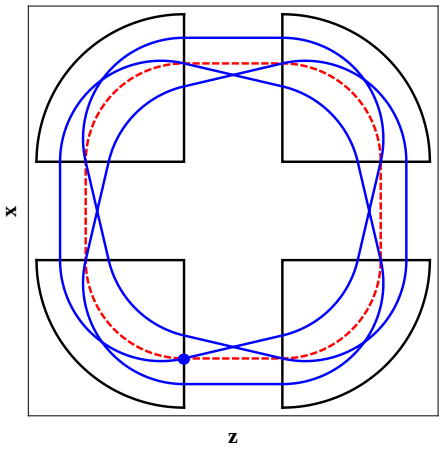

(b) Momentum offset from equilibrium orbit.

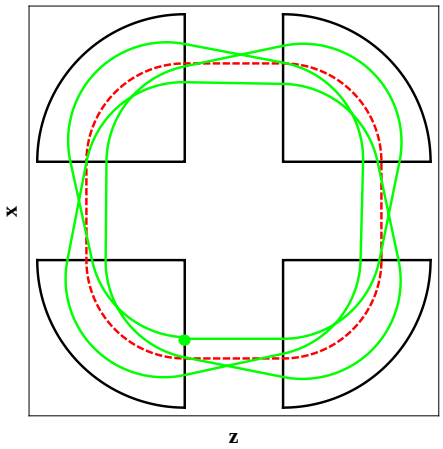

(c) Spatial offset from equilibrium orbit.

FIGURE 2.2: Separating the quadrants of the disc shown in Fig. 2.1 (with regions that have no fields) changes the particle dynamics. Figure (a) shows the equilibrium orbit for the new dipole configuration (this path is marked by the red dashed line in Figs. (b) and (c)); if particles have a momentum (Fig. (b)) or spatial (Fig. (c)) offset from the equilibrium orbit, then these particles will oscillate around the equilibrium orbit. In this example the betatron tune is 4/3.

larger than 3 (Fig. 2.3b). The change in tune with the fractional offset in momentum is called the chromaticity. We now select an entrance plane to one of the dipole quarters as an observation point, and track a particle through 200 orbits in the disc. Each time the particle passes an observation point, its position and momentum along the x axis are recorded. Figure 2.3 shows that a plot of the recorded momentum vs. position forms an ellipse, at the centre of which is the closed orbit. The orientation of the ellipse in a plot of the transverse dynamical variables relative to the closed orbit is dependent upon where along the closed orbit the measurement is made; however, the area enclosed by the ellipse will remain constant (a matrix describing the propagation of the motion is symplectic).

In a cyclotron, dispersion is the key to maintaining isochronicity: as particles are

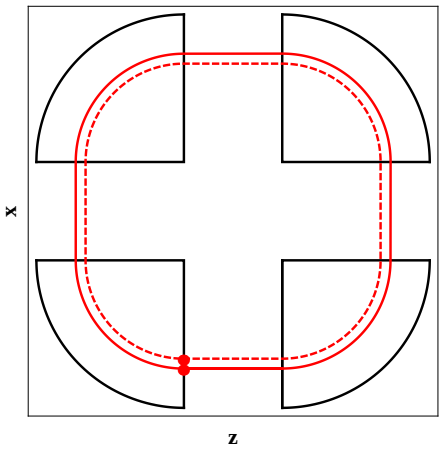

(a) Change in closed orbit with momentum.

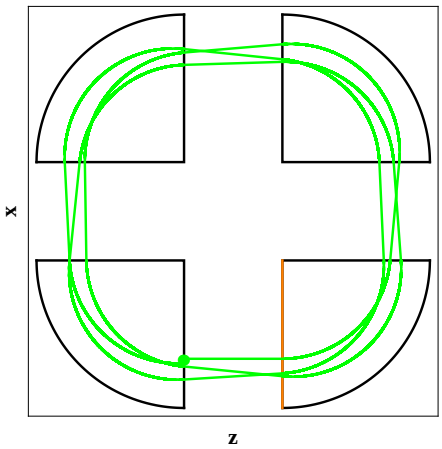

(b) Spatial offset from equilibrium orbit.

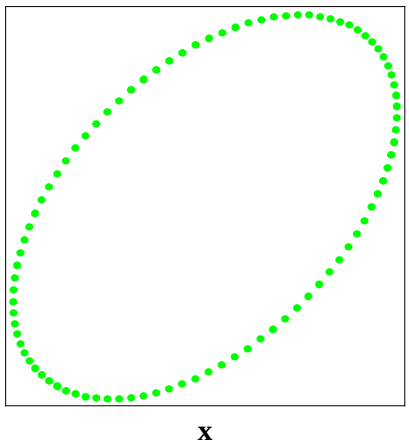

(c) Poincaré plot.

FIGURE 2.3: The momentum of the particle in these plots is greater than that in Fig. 2.2. Dispersion leads to a closed orbit that is found at a larger radius: the closed orbit for the new and previous momenta are shown by the solid and dashed red lines in Fig. (a) respectively. Figure (b) shows that the path of a particle that is offset from the equilibrium orbit no longer closes every three turns; the change in betatron tune with a fractional change in momentum is called the chromaticity. A particle is tracked through 200 turns, and the phase space variables, position and momentum, of the particle in the direction of the x axis are recorded at a boundary of the lower right quadrant (marked orange in Fig. (c)); the plot of p_x vs. x at the boundary traces out an ellipse. If the values of the phase space variables are recorded at successive points in the remainder of the quadrant (along an axis transverse to the closed orbit), then the ellipse will be seen to change shape, whilst the area enclosed by the ellipse remains constant. Given the periodicity of the system, the evolution of the ellipse is the same for each quadrant.

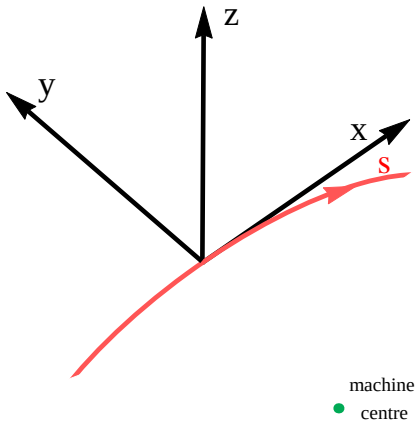

FIGURE 2.4: Coordinate system local to the equilibrium path of a particle beam. x is the longitudinal coordinate (tangential to the closed orbit). y and z are transverse horizontal and vertical respectively.

accelerated the increase in velocity is compensated by the increase in orbital circumference, so the orbital period remains constant. However, in a synchrotron, the fields are ramped up during acceleration so as to overcome the effect of dispersion and maintain a fixed closed orbit. In FFAGs, the magnetic fields increase with radius, which results in a reduction of the dispersion. For both synchrotrons and FFAGs the motion of particles can be described in terms of small oscillations (compared to the orbital radius) around closed orbits: for this reason it is convenient to use a coordinate system that has an origin at or close to the position of the closed orbit of a beam at any point around the circumference of the accelerator (Fig. 2.4).

The focusing provided by a uniform field directed along the y axis acts only in the zx plane. Eventually, any particle with some component of velocity in the y direction will be lost. A number of non-uniform field profiles (introduced in section 2.3) are commonly applied in accelerators to provide both horizontal and vertical focusing and to ensure that the beam dynamics within an accelerator are as required for any given application.

2.3 Magnets for accelerators

In section 2.2, we demonstrated that a uniform magnetic field could have a focusing effect of the motion of a charged particle. The focusing was limited to constraining particles along an axis that was perpendicular to both the magnetic field and the direction of motion of a particle travelling along the closed orbit (the transverse horizontal y axis in Fig. 2.4); if we want to be able to constrain motion to some region along the transverse vertical z axis, then we will need additional magnetic field profiles. Furthermore, we saw the momentum dependence of the focusing properties of a uniform field

when chromaticity was introduced; later in this chapter we will see why we may want to control the chromaticity of a lattice, and how this can be achieved using a variety of magnetic field profiles.

An understanding of the magnetic field profiles that can be produced for use within a particle accelerator can be gained by starting with two of Maxwell's equations,

$$\nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}, \tag{2.3}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.4}$$

where $\mathbf{B} = \mu_0 \mathbf{H}$, with μ_0 being the permeability of free space, $\mathbf{D}$ the electric displacement and $\mathbf{J}$ the current density. In the vacuum of a beamline that is surrounded by magnets and zero (or constant) electric fields, Eq. 2.3 may be rewritten as

$$\nabla \times \mathbf{H} = 0. \tag{2.5}$$

It follows that we can write the magnetic field in terms of a magnetic scalar potential, ψ:

$$\mathbf{B} = -\nabla \psi. \tag{2.6}$$

Combining Eq. 2.4 and Eq. 2.6 then gives Laplace's equation for the scalar potential, ψ: $\nabla^2 \psi = 0$. We consider the case where magnetic fields vary along the transverse axes of a magnet, but are constant along the longitudinal axis. This can be a good representation of the fields within a magnet when away from the entrance and exit, in which regions the fields do not vary with longitudinal position. Restricting this derivation to two dimensions, Laplace's equation in polar coordinates has the form:

$$\nabla^2 \psi = \frac{\partial^2 \psi}{\partial r^2} + \frac{1}{r}\frac{\partial \psi}{\partial r} + \frac{1}{r^2}\frac{\partial^2 \psi}{\partial \phi^2} = 0, \tag{2.7}$$

where, based on the local coordinate system of Fig. 2.4, $r = \sqrt{y^2 + z^2}$ and $\phi = \arctan\left(\frac{z}{y}\right)$. A solution to Eq. 2.7 is given by [32]:

$$\psi = \sum_{n=1}^{\infty} \left(a_n r^n \cos(n\phi) + b_n r^n \sin(n\phi)\right), \tag{2.8}$$

where, in the case of a pure multipole field, $2n$ is the number of magnet poles. The a_n and b_n coefficients determine the field strengths of the cosine and sine terms which are at an angle of $\pi/2n$ (skew multipole) and normal (upright multipole) to the xy plane respectively. Taking the normal quadrupole as an example, $n = 2$ and $a_2 = 0$, then Eq. 2.8 evaluates to:

$$\psi_2 = b_2 r^2 \sin(2\phi),$$

with a transformation to Cartesian coordinates giving:

$$\psi_2 = 2b_2yz. \tag{2.9}$$

The magnetic field is then found by using Eqs. 2.6 and 2.9:

$$\mathbf{B} = -2b_2(z, y, 0),$$

which may be written in terms of the field gradient along an axis in the quadrupole:

$$\mathbf{B} = \frac{\partial B_y}{\partial z}(z, y, 0).$$

The Lorentz force on a particle travelling perpendicular to the longitudinal axis of the quadrupole ($\mathbf{v} = (0, 0, v_x)$) is:

$$\mathbf{F} = qv_x\frac{\partial B_y}{\partial z}(z, -y, 0).$$

At this point we note that the force along the horizontal and vertical transverse axes of the normal quadrupole are dependent on z and $-y$ respectively. This property makes the quadrupole focusing for one transverse axis and defocusing for the other. In section 2.4 we see how pairs of quadrupole magnets are used to provide focusing along both transverse axes. The accelerators within this thesis are constructed with upright multipole magnets; however, in practice, there will be small skew multipole field components due to errors in the rotational alignment of magnets. Equation 2.8 has been evaluated for the first four values of n, with the results in Cartesian coordinates of the upright multipole fields given in table 2.1 and the field profiles shown in Fig 2.5. Figure 2.6 shows the force experienced by a particle as it passes through a quadrupole magnet and through a sextupole magnet, as a function of transverse position in the magnet. Throughout this thesis, we define multipole magnets by their normalised gradient (or strength), which is given by:

$$K_{n-1} = \frac{\partial^{n-1}B_y}{\partial z^{n-1}}\frac{1}{B\rho}, \tag{2.10}$$

where $B\rho$ is referred to as the particle rigidity, and relates the magnetic field (B) and bending radius (ρ) experienced by a particle to the particle momentum (p) and charge (q):

$$B\rho = \frac{p}{q}. \tag{2.11}$$

In Eq. 1.3 a field profile for the midplane of a scaling FFAG was given. A binomial expansion of this equation shows how the FFAG theoretically consists of an infinite

TABLE 2.1: Analytic expressions of multipole magnet field components.

Component	B_y	B_z
Dipole	0	$B\rho K_0$
Quadrupole	$B\rho K_1 z$	$B\rho K_1 y$
Sextupole	$B\rho K_2 yz$	$\frac{1}{2}B\rho K_2(y^2 - z^2)$
Octupole	$\frac{1}{6}B\rho K_3(3y^2 z - z^3)$	$\frac{1}{6}B\rho K_3(y^3 - 3yz^2)$

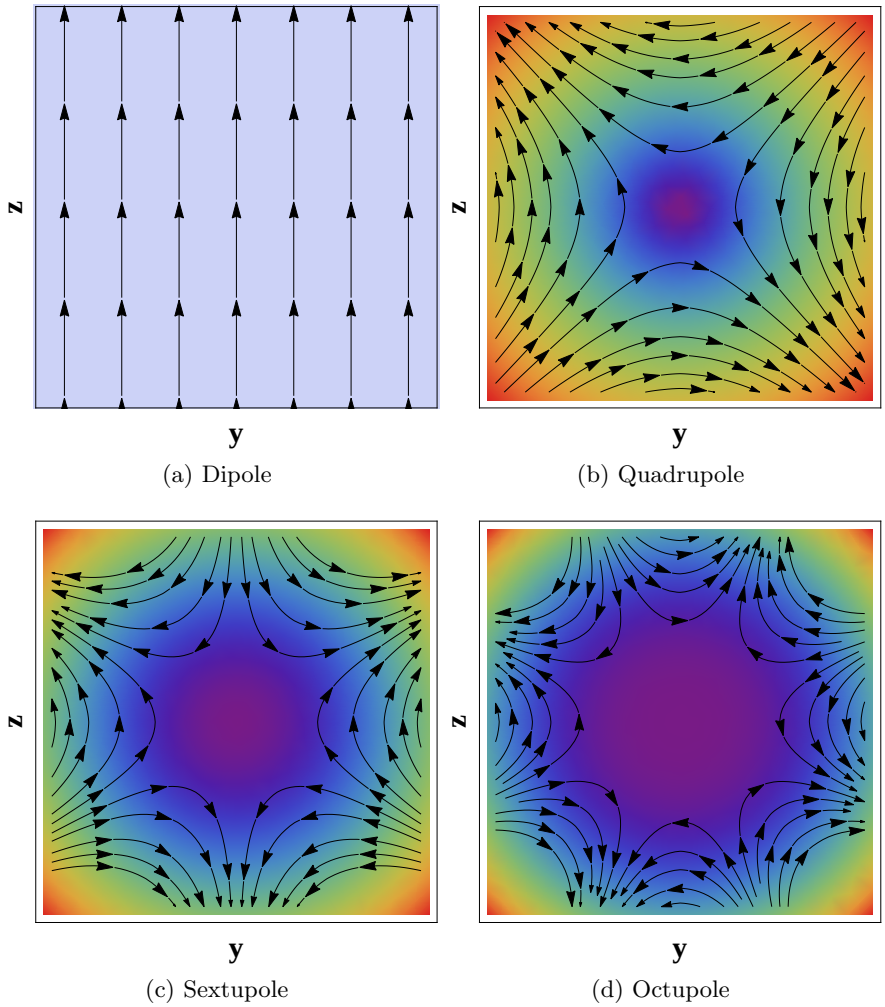

(a) Dipole (b) Quadrupole

(c) Sextupole (d) Octupole

FIGURE 2.5: Field profiles for the first four multipole magnet components. With the defined coordinate system, particles travel out of the page; the quadrupole field shown will be horizontally focusing for a negatively charged particle, and defocusing for a positively charged particle.

number of multipole magnet components:

$$\begin{aligned} B_z &= B_{z,0}\left(\frac{R}{R_0}\right)^k \\ &= B_{z,0}\left(1+\frac{k\Delta R}{R_0}+\frac{(k-1)k\Delta R^2}{2R_0^2}+\frac{(k-2)(k-1)k\Delta R^3}{6R_0^3}+...\right) \\ &= B_{z,0}\left(1+\sum_{n=1}^{\infty}\frac{k(k-1)(k-2)...(k-n+1)\Delta R^n}{n!R_0^n}\right), \end{aligned} \tag{2.12}$$

where $R = R_0 + \Delta R$. Figure 2.7 shows the field profile that satisfies the first 5 terms of Eq. 2.12, which, for the range shown, is a good approximation for the scaling FFAG magnet. Later in this chapter, two FFAG accelerator designs are discussed: one is built using scaling FFAG magnets; the second is designed with combined function magnets (where the fields within an accelerator magnet consist of two or more multipole components) that satisfy the first 5 terms of Eq. 2.12.

2.4 Linear particle dynamics

Understanding the motion of a particle with respect to the closed orbit can give key information about the performance of a particle accelerator. A popular approach is to describe each accelerator component in terms of a Hamiltonian that determines the dynamics of a charged particle within the component. From this, Hamiltonian maps that propagate the dynamical variables through the different accelerator components can be derived. The coefficients within these maps tell us about the behaviour of a particle beam at the location of the corresponding component within the accelerator. Maps for a sequence of components can be combined to produce a map for some section of accelerator, or even for an entire accelerator. The particle tracking studies within this thesis have been carried out using numerical methods (which are discussed in chapter 3), the results of tracking are often then used to calculate the map that describes the linear particle motion through the accelerator.

In Fig. 2.4 a coordinate system local to the closed orbit of a particle was defined. The transverse motion of a particle with respect to the closed orbit may be described by using the dynamical variables y, y', z and z', where the prime indicates the change in a variable with respect to the path length, s, of a particle following the closed orbit; the value of y', for example, can be calculated as:

$$y' = \arctan\left(\frac{p_y}{p_x}\right),$$

where p_y is transverse horizontal momentum with respect to the closed orbit and p_x is the momentum tangential to the closed orbit at a given location.

In section 2.2 we saw that a charged particle travelling with some component of velocity perpendicular to a uniform magnetic field oscillates around the closed orbit in the

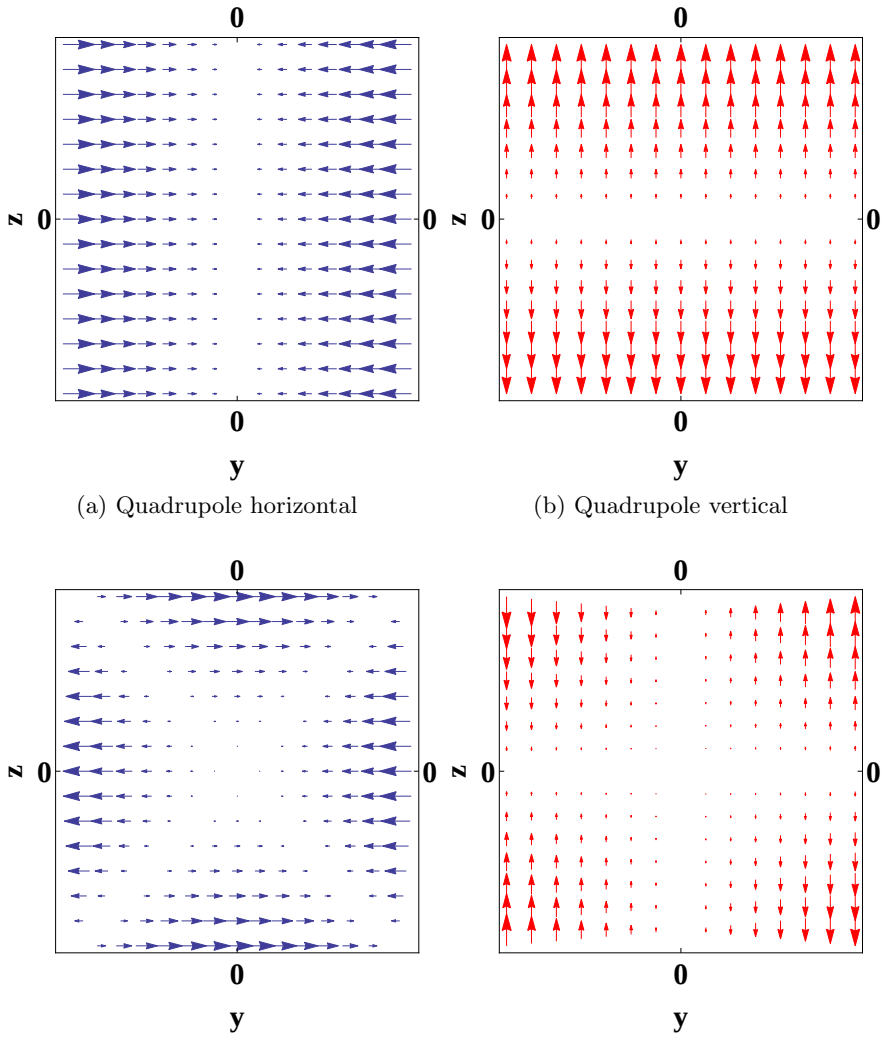

(a) Quadrupole horizontal (b) Quadrupole vertical

(c) Sextupole horizontal (d) Sextupole vertical

FIGURE 2.6: Horizontal and vertical forces experienced by a charged particle travelling through a magnetic multipole field.

plane perpendicular to the field. This oscillation indicated that the motion of the particle was stable along one of the transverse axes, however there was nothing constraining the particle along the second transverse axis. In section 2.3 we introduced several magnet field profiles, including that of the quadrupole magnet; the force experienced by a particle as it travels through a quadrupole increases linearly with the particle offset from the quadrupole axis. The field profile of the quadrupole focuses the particle towards the closed orbit along one transverse axis, and defocuses along the second. The small transverse oscillations of particles around a closed orbit as they travel through dipole and quadrupole fields may be represented by a pair of periodic second order differential

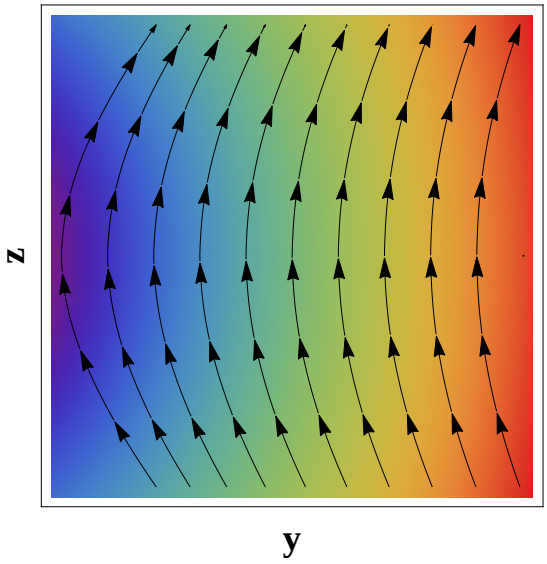

FIGURE 2.7: Field profile of a combined function magnet containing components from dipole up to decapole. This is approximately equivalent to a scaling FFAG magnet over the range shown.

equations, which are known as Hill's equations:

$$\begin{aligned} \frac{\partial^2 y}{\partial s^2} + K_y(s)y = 0, \\ \frac{\partial^2 z}{\partial s^2} + K_z(s)z = 0, \end{aligned} \tag{2.13}$$

where the functions $K_y(s)$ and $K_z(s)$ are the piecewise constants that give respectively the transverse horizontal and vertical focusing strength of a component at a position s. In a circular accelerator, a particle will encounter the different steering and focusing components once every turn in the accelerator, so that along the horizontal axis, for example, $K_y(s) = K_y(s + C)$, where C is the path length of a particle following the closed orbit for one turn. If steering occurs only in the xy plane, then the horizontal and vertical focusing functions at a given position, s, are written as:

$$K_y = \frac{1}{\rho(s)^2} + K_1(s), \quad K_z = -K_1(s), \tag{2.14}$$

where ρ is the bending radius of a dipole magnetic field component and K_1 is the quadrupole magnet focusing strength that was defined in section 2.3. A solution to Hill's equation for positive values of K is given by [33]:

$$y(s) = a\cos(\sqrt{K_y}s + b), \tag{2.15}$$

where a and b are determined by the initial phase space coordinates of a particle (y_0, y_0'). Through expansion of the cosine term of Eq. 2.15, we can write:

$$y(s) = A\cos(\sqrt{K_y}s) + B\sin(\sqrt{K_y}s),$$

and then by differentiating y with respect to s, we obtain an expression for y':

$$y'(s) = -\sqrt{K_y}A\sin\left(\sqrt{K_y}s\right) + \sqrt{K}B\cos\left(\sqrt{K_y}s\right).$$

If we take the limit of $s \to 0$, we find that $A = y_0$ and $B = y'_0/\sqrt{K_y}$, and we can now write the transfer map for propagating the dynamical variables through a combined function dipole (with horizontally focusing quadrupole terms, $K_y > 0$) of length, L:

$$\mathrm{M}_f = \begin{pmatrix} \cos\left(\sqrt{K_y}L\right) & \frac{1}{\sqrt{K_y}}\sin\left(\sqrt{K_y}L\right) & 0 & 0 \\ -\sqrt{K_y}\sin\left(\sqrt{K_y}L\right) & \cos\left(\sqrt{K_y}L\right) & 0 & 0 \\ 0 & 0 & \cosh\left(\sqrt{|K_z|}L\right) & \frac{1}{\sqrt{|K_z|}}\sinh\left(\sqrt{|K_z|}L\right) \\ 0 & 0 & \sqrt{|K_z|}\sinh\left(\sqrt{|K_z|}L\right) & \cosh\left(\sqrt{|K_z|}L\right) \end{pmatrix} \quad (2.16)$$

To find the dynamical variables after the component, we then multiply the map by the initial dynamical variables:

$$\begin{pmatrix} y_1 \\ y'_1 \\ z_1 \\ z'_1 \end{pmatrix} = M_f \begin{pmatrix} y_0 \\ y'_0 \\ z_0 \\ z'_0 \end{pmatrix}.$$

The transfer matrix, M_f, contains both ordinary and hyperbolic trigonometric terms; this provides a description of the focusing action of the quadrupole along the horizontal axis and defocusing along the vertical axis. In the case that $K_y < 0$, we refer to the quadrupole as being defocusing, and the map is given by:

$$\mathrm{M}_d = \begin{pmatrix} \cosh\left(\sqrt{|K_y|}L\right) & \frac{1}{\sqrt{|K_y|}}\sinh\left(\sqrt{|K_y|}L\right) & 0 & 0 \\ \sqrt{|K_y|}\sinh\left(\sqrt{|K_y|}L\right) & \cosh\left(\sqrt{|K_y|}L\right) & 0 & 0 \\ 0 & 0 & \cos\left(\sqrt{K_z}L\right) & \frac{1}{\sqrt{K_z}}\sin\left(\sqrt{K_z}L\right) \\ 0 & 0 & -\sqrt{K_z}\sin\left(\sqrt{K_z}L\right) & \cos\left(\sqrt{K_z}L\right) \end{pmatrix}. \quad (2.17)$$

It is possible to have a net beam focusing effect along both transverse axes by constructing accelerator lattices with sequences of focusing and defocusing quadrupoles that are separated by drift spaces. The FODO (focusing-drift-defocusing-drift) cell is a common example; we can calculate a map for transporting particles through a FODO cell by multiplying the maps for the individual components together in the following fashion:

$$M = M_o M_d M_o M_f, \quad (2.18)$$

where M_o is the map for a drift space. The quadrupole maps given describe motion through a quadrupole that is perfectly aligned with respect to the reference axis, meaning that motion along the transverse axes is uncoupled (i.e. the motion along the y axis is independent of the motion along the z axis and vice-versa). This results in a map for the FODO cell that is block diagonal; the upper-left and lower-right 2×2 matrices within the map individually describe motion along the horizontal and vertical axis, respectively.

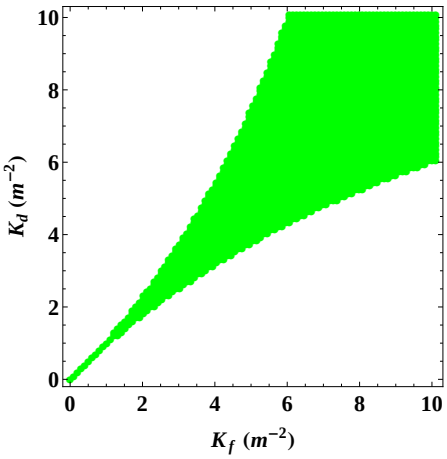

FIGURE 2.8: Stability diagram for a periodic FODO cell. Two 20 cm long quadrupoles are separated by 60 cm drift spaces, K_f and K_d give the strength of the horizontally focusing and defocusing quadrupoles respectively. A transfer matrix for the cell is calculated by multiplying the matrices for the individual components (as in Eq. 2.18); the stability of the cell is determined by calculating the trace of the 2×2 matrices that describe the motion along each transverse axis. Quadrupole strengths that give stability along both transverse axes (when $|\mathrm{Tr}(M)| \leq 2$) are marked in green. The choice of quadrupole strengths from those that give transverse stability is dictated by the beam dynamics required of the accelerator.

It can be shown [34] that for net focusing along a given transverse axis, the trace of the 2×2 matrix must meet the criterion:

$$|\mathrm{Tr}(M)| \leq 2.$$

For a single FODO cell configuration there can be a range of quadrupole strengths that give stability along both transverse axes (Fig. 2.8). From the quadrupole strengths that provide stability along both transverse axes, the selected strength is the one that gives the best match for a number of requirements of the beam dynamics (for example, in chapter 5 we choose quadrupole strengths that bring the transverse motion close to a resonant condition).

The trace of the matrix being less than 2 indicates a real (rather than complex or imaginary) phase advance of the transverse oscillation over one period (in this case the FODO cell).

A second approach to solving Hill's equations allows for a more general description of the linear transverse motion at a given position, s, to be developed. We start by taking a solution, y, that is written in terms of the amplitude and phase of the oscillation:

$$y(s) = a\sqrt{\beta_y(s)}\cos\left(\psi(s) + \phi_0\right), \tag{2.19}$$

where $a\sqrt{\beta(s)}$ gives the amplitude of the particle oscillation around the closed orbit at a position s, ϕ_0 is the initial phase of the oscillation and $\psi(s)$ is the phase advance of the particle since ϕ_0. The values of a and ϕ_0 are determined by the starting coordinates

(y, y') of a particle, and are independent of s. If we take Hill's equation (Eq. 2.4) as:

$$y'' = -K_y y, \tag{2.20}$$

and then substitute Eq. 2.19 into Eq. 2.20, then we find:

$$\begin{aligned}\frac{-a}{\sqrt{\beta(s)}} &\sin\left(\psi(s)+\phi_0\right)\left(\beta'(s)\psi'(s)+\beta(s)\psi''(s)\right) \\ &- \frac{a}{4\beta(s)^{3/2}}\cos\left(\psi(s)+\phi_0\right)\left(-2\beta(s)\beta''(s)+\beta'(s)^2+4\beta(s)^2\psi'(s)^2\right) \\ &= -K_y a\sqrt{\beta(s)}\cos\left(\psi(s)+\phi_0\right).\end{aligned} \tag{2.21}$$

After equating the sine and cosine terms on the left hand side of Eq. 2.21 to those on the right hand side, we note that the sine term must always be zero, which is true if:

$$\beta'(s)\psi'(s) + \beta(s)\psi''(s) = 0. \tag{2.22}$$

Equation 2.22 has the solution:

$$\psi'(s) = \frac{1}{\beta(s)}, \tag{2.23}$$

and it is noted that integrating this over the circumference of an accelerator and dividing by 2π gives the betatron tune:

$$Q = \frac{1}{2\pi}\oint_C \frac{1}{\beta(s)}ds. \tag{2.24}$$

To give a description of the motion of a particle in transverse phase space over many turns we start with Eq. 2.19 and its derivative:

$$y'(s) = -\frac{a}{\sqrt{\beta_y(s)}}(\alpha_y(s)\cos(\psi(s)+\phi_0) + \sin(\psi(s)+\phi_0)), \tag{2.25}$$

where $\alpha_y(s) = -\frac{\beta_y'(s)}{2}$. We then write:

$$\cos\left(\psi(s)+\phi_0\right) = \frac{y}{a\sqrt{\beta_y(s)}}, \tag{2.26}$$

and then substitute this into Eq. 2.25 to obtain:

$$\sin\left(\psi(s)+\phi_0\right) = -\left(\frac{\sqrt{\beta_y(s)}y'}{a} + \frac{\alpha_y(s)y}{a\sqrt{\beta_y(s)}}\right). \tag{2.27}$$

Using the identity $\sin^2\theta + \cos^2\theta = 1$, Eqs. 2.26 and 2.27 can be combined to eliminate $\psi(s)$ and give:

$$\frac{y^2}{\beta_y(s)} + \left(\frac{\alpha_y(s)}{\sqrt{\beta_y(s)}}y + \sqrt{\beta_y(s)}y'\right)^2 = a^2. \tag{2.28}$$

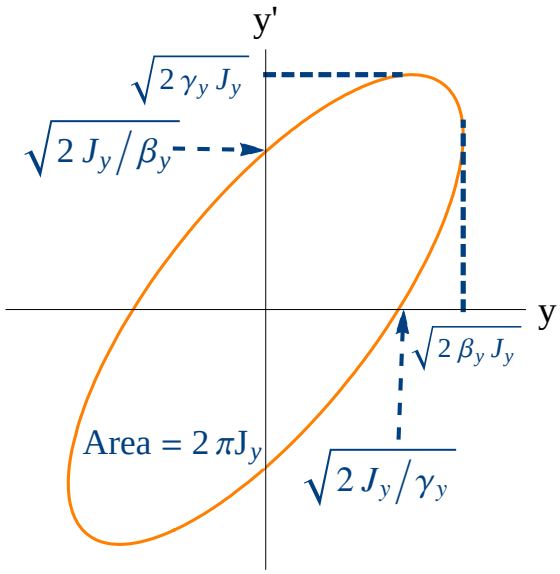

FIGURE 2.9: The motion of a particle following a stable trajectory within a linear system will form an ellipse in transverse phase space. The Courant-Snyder parameters (α, β and γ) define the orientation of the ellipse at a location s, whilst the area of the ellipse is constant for all s.

Finally, we introduce the term $\gamma_y(s) = \frac{1+\alpha_y^2(s)}{\beta_y(s)}$, which allows us to write Eq. 2.28 as:

$$\gamma_y(s)y^2(s) + 2\alpha_y(s)y(s)y'(s) + \beta_y(s)y'^2(s) = a^2. \tag{2.29}$$

Equation 2.29 describes an ellipse (as could be expected from Fig. 2.3c) of area πa^2. The variables $\alpha_y(s)$, $\beta_y(s)$ and $\gamma_y(s)$ are referred to the Courant-Snyder parameters and describe the orientation of the ellipse traced in phase space at a given position, s. Within our equations the value a is constant, which means that the area enclosed by a phase space ellipse remains the same irrespective of the ellipse orientation. If we are tracking a single particle, then a is related to the action, J_y, of the particle in phase space by $a = \sqrt{2J_y}$, and Eq. 2.29 becomes:

$$J_y = \frac{1}{2}\left(\gamma_y(s)y^2 + 2\alpha_y(s)yy' + \beta_y(s)y'^2\right). \tag{2.30}$$

An ellipse that is described by Eq. 2.30 is shown in Fig. 2.9. Alternatively, we may describe a bunch of many particles after first defining the emittance, ϵ_y, as the mean action of the particles within the bunch:

$$\epsilon_y = < J_y > .$$

In this case, $a = \sqrt{\epsilon_y}$ and Eq. 2.29 is then:

$$\epsilon_y = \gamma_y(s)y^2 + 2\alpha_y(s)yy' + \beta_y(s)y'^2. \tag{2.31}$$

We now have two metrics for describing the relative size of transverse oscillations that

are independent of the location of the particles longitudinally. It can be shown that the action and emittance will vary inverse proportionally to the particle momentum during acceleration [35] (leading to a reduced beam size at higher momentum); it is helpful to define normalised values of the action and emittance, which are independent of particle momentum. The normalised values are given by:

$$J_{N,y} = \beta\gamma J_y,$$
$$\epsilon_{N,y} = \beta\gamma\epsilon_y,$$

where β and γ are the relativistic parameters. Both non-normalised (or geometric) and normalised values of action and emittance are used within this thesis, with the preferred metric depending upon the situation at hand. For example, when investigating single particle dynamics within a lattice, we often use the geometric action, whereas when describing a bunch of particles during acceleration, it is the normalised emittance that is preferred.

A matrix that transports particles through one turn in the accelerator can be written in terms of the Courant-Snyder parameters and the phase advance of the particle. In later chapters, particle tracking is used to calculate the coefficients of the one turn map, and subsequently the Courant-Snyder parameters. If we take Eqs. 2.19 and 2.25, and then expand the trigonometric terms, we obtain:

$$y(s) = a\sqrt{\beta(s)}\left(\cos\left(\psi(s)\right)\cos\left(\phi_0\right) - \sin\left(\psi(s)\right)\sin\left(\phi_0\right)\right), \tag{2.32}$$

$$\begin{aligned} y'(s) = -\frac{a}{\sqrt{\beta(s)}}(\alpha(s)(&\cos\left(\psi(s)\right)\cos\left(\phi_0\right) - \sin\left(\psi(s)\right)\sin\left(\phi_0\right)) \\ &+ \sin\left(\psi(s)\right)\cos\left(\phi_0\right) + \cos\left(\psi(s)\right)\sin\left(\phi_0\right)). \end{aligned} \tag{2.33}$$

When $s = 0$, then $y(0) = y_0$, $y'(0) = y'_0$, $\alpha(0) = \alpha_0$, $\beta(0) = \beta_0$ and $\psi(0) = 0$, and we can show that:

$$\cos\left(\phi_0\right) = \frac{y_0}{\sqrt{2\beta_0 J}} \tag{2.34}$$

$$\sin\left(\phi_0\right) = -\left(\sqrt{\frac{\beta_0}{2J}}y'_0 + \frac{\alpha_0 y_0}{\sqrt{2\beta_0 J}}\right). \tag{2.35}$$

Inserting these into Eqs. 2.32 and 2.33, gives following equations for the betatron motion:

$$y(s) = \sqrt{\frac{\beta_y(s)}{\beta_{y,0}}}\left(\cos\left(\psi(s)\right) + \alpha_{y,0}\sin\left(\psi(s)\right)\right)y_0 + \sqrt{\beta_y(s)\beta_{y,0}}\sin\left(\psi(s)\right)y'_0, \tag{2.36}$$

$$\begin{aligned} y'(s) = \frac{1}{\sqrt{\beta_y(s)\beta_{y,0}}}&((\alpha_{y,0} - \alpha_y(s))\cos\left(\psi(s)\right) - (1 + \alpha_{y,0}\alpha_y(s))\sin\left(\psi(s)\right))y_0 \\ &+ \sqrt{\frac{\beta_{y,0}}{\beta_y(s)}}(\cos\left(\psi(s)\right) - \alpha_y(s)\sin\left(\psi(s)\right))y'_0. \end{aligned} \tag{2.37}$$

If we observe the betatron motion over a number of turns at a single point within the lattice (for example, at the location of a beam position monitor), then the matrix that transports particles over successive turns is given by:

$$M_y = \begin{pmatrix} \cos\mu_y + \alpha_{y,0}\sin\mu_y & \beta_{y,0}\sin\mu_y \\ -\gamma_{y,0}\sin\mu_x & \cos\mu_y - \alpha_{y,0}\sin\mu_y \end{pmatrix}, \tag{2.38}$$

where:

$$\begin{pmatrix} y_1 \\ y_1' \end{pmatrix} = M_y \begin{pmatrix} y_0 \\ y_0' \end{pmatrix},$$

and $\mu_y = 2\pi Q_y$.

Having a general description of the motion of particles (as given by the Courant-Snyder parameters) for the entire lattice allows for optimal positioning of components that control the beam dynamics (for example, sextupoles). Given a matrix consisting of the Courant-Snyder parameters along a single transverse axis at a longitudinal position, s_0:

$$A_y = \begin{pmatrix} \gamma_y & \alpha_y \\ \alpha_y & \beta_y \end{pmatrix}, \tag{2.39}$$

it can be shown [36] that, if the matrices that propagate the dynamical variables along the lattice are known, then the matrix of Courant-Snyder parameters can be transported from a position s_0 to s_1 by:

$$A_y(s_1)^{-1} = M_y(s_1, s_0) A_y(s_0)^{-1} M_y(s_1, s_0)^{\mathrm{T}}, \tag{2.40}$$

where T indicates the transpose of the matrix.

2.5 Resonant particle motion

In Fig. 2.8, we saw for the FODO cell example that there can be a range of quadrupole strengths that lead to stable motion along the horizontal and vertical transverse axes. The condition for stability was defined by the trace of a matrix that describes the motion along a transverse axis being less than 2, with the physical meaning of this being a real phase advance of the particle along the transverse axis. However, if the betatron tune is integer or some fraction of an integer, the effect on the beam dynamics of field errors may add coherently, and motion may again be unstable. Typically, lower order resonances are more likely to lead to beam loss, and accelerators are designed to operate away from these conditions. Alternatively, an accelerator may be brought close to a transverse resonance along one axis in order to extract beam in a controlled fashion; these extraction methods, when used for proton therapy, are a focus of this thesis. The resonance condition for betatron motion is:

$$aQ_y + bQ_z = nP, \tag{2.41}$$

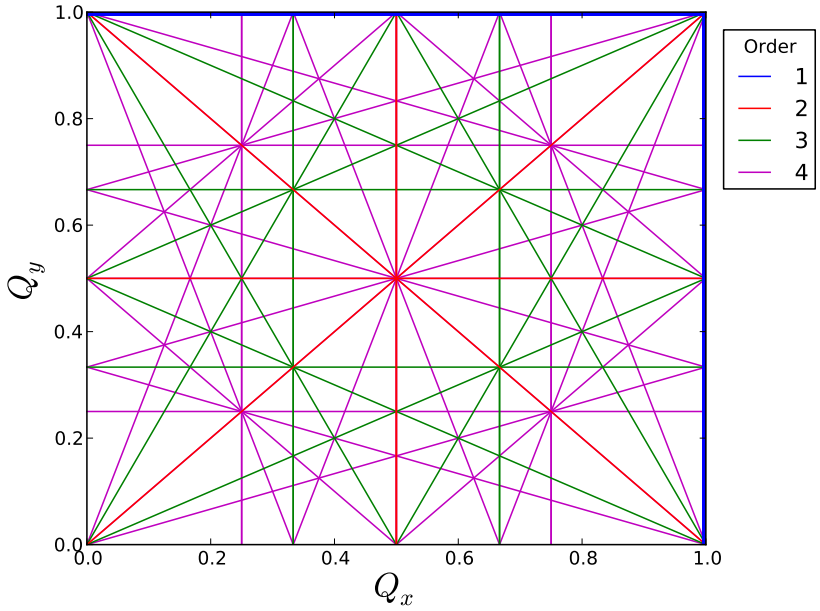

FIGURE 2.10: Tune resonance diagram for the first four orders of resonance for a lattice that is made up of a single period ($P = 1$ in Eq. 2.41). Accelerators usually have tune working points that are away from lower order resonances in order to minimise beam loss.

where a, b and P are integers, $|a|+|b|$ is the order of the resonance, n is an integer and P is the number of periodic cells that makes up the accelerator. Figure 2.10 shows a tune diagram for the first four orders of resonance in the case that $P = 1$. A resonance may involve one (a or b equal to zero) or both (a and b non-zero) transverse axes, with the latter referred to as a coupling resonance. Each order of resonance requires a field error of at least the same order to drive it, for example, a dipole field error may drive integer resonances, whilst a sextupole may drive integer, half-integer or third-integer resonances. The field components that can drive resonances may be introduced intentionally to control some aspect of the beam dynamics (e.g. chromaticity) or unintentionally, for example, through the misalignment of magnets. In the case of the field components included intentionally, increasing the number of periods that makes up a lattice reduces the density of the resonance lines within a tune diagram, and may allow the accelerator to operate further from resonant conditions.

2.6 Chromaticity control

In section 2.5, we saw that certain values of the betatron tune would lead to resonance conditions for the transverse motion, and that a resonance of a given order required a magnetic field term of at least the same order to drive the resonance. A working point of an accelerator is usually selected that leads to a betatron tune that is away from resonant conditions. In the discussion on linear beam dynamics (section 2.4), the beam was steered and focussed using dipole and quadrupole magnets, suggesting that it is

the integer and half-integer resonances that should be primarily avoided. However, the particles that are injected into an accelerator will have a range of momenta, which are distributed around the nominal momentum of the beam. Equations 2.10 and 2.14 tell us that the focusing strength of a magnetic field is dependent on the momentum of a particle, with higher momentum particles experiencing lower focusing strengths than low momentum particles; the result of this is a change in betatron tune with momentum, which is referred to as chromaticity. If the tune spread of a bunch is large enough, then off-momentum particles may encounter resonant conditions, and be lost from the accelerator.

By using the field profiles given in table 2.1, it can be shown that a sextupole magnet offers an additional quadrupole-type focusing term for off-momentum particles, provided that the sextupole is located in an area of non-zero dispersion. For a particle that has some momentum deviation, δ, the transverse positions of particle can be written so as to include the dispersion of the accelerator at the location of the sextupole. If the dispersion, $D_y(s)$, is non-zero along the y axis and zero along the z axis, then,

$$y = D_y(s)\delta + \Delta y,$$

$$z = \Delta z,$$

where Δy and Δz are the transverse offsets from the off-momentum closed orbit positions. Substituting this into the equation for the sextupole then gives:

$$\begin{aligned} B_z &= \frac{1}{2} B\rho K_2((D_y(s)\delta)^2 + 2D_y(s)\delta\Delta y + \Delta y^2 - \Delta z^2) \\ B_y &= -B\rho K_2(D_y(s)\delta\Delta z + \Delta y\Delta z). \end{aligned} \tag{2.43}$$

It can be seen that for the vertical field component there are dipole, quadrupole and sextupole terms, whilst for the horizontal component there is just a quadrupole and sextupole term. The chromaticity of a lattice is calculated by:

$$\xi_y = -\frac{1}{4\pi} \oint K_y \beta_y \mathrm{d}s. \tag{2.44}$$

In our lattice that is made up of dipoles and quadrupoles, alternating gradient focusing leads to β functions that are larger in the focusing quadrupoles than in the defocusing quadrupoles; evaluation of the integral in Eq. 2.44 will give a negative chromaticity (in keeping with our qualitative description). By including sextupole components, the chromaticity may be controlled (to the first order of δ), and the linear term of chromaticity set to a required value. Equation 2.44 with the quadrupole terms that arise from the

the inclusion of sextupole fields (Eq. 2.43) is written as:

$$\xi_y = -\frac{1}{4\pi}\oint \beta_y(K_1 - K_2 D_y)\mathrm{d}s,$$
$$\xi_z = \frac{1}{4\pi}\oint \beta_z(K_1 - K_2 D_y)\mathrm{d}s, \tag{2.45}$$

and it can be seen that by careful selection and positioning, sextupoles can bring the chromaticity of the lattice to zero.

2.7 Longitudinal dynamics

In the machines studied in this thesis, the acceleration of particles takes place within rf cavities that have length that is short compared to the total circumference of the accelerator. The cavities are designed to produce an oscillating electric field that is directed along the axis of the longitudinal motion of the beam. Starting with Eq. 2.1, then the rate of change of momentum of a particle as it travels through the electric field, $E_x(t)$, is given by

$$\frac{d\mathbf{p}}{dt} = qE_x(t), \tag{2.46}$$

where $E_x(t) = E_{x,0}\sin(\omega t + \phi_0)$, with ω the angular frequency of the field, t the time and ϕ_0 the phase of the oscillation at $t = 0$. Assuming that the velocity of a particle remains constant whilst being accelerated within a cavity (which is approximately true if the change in energy of a particle within a single cavity is small or if the beam is relativistic), then the change in kinetic energy of a particle can be found by integrating Eq. 2.46 over the length, L, of a cavity,

$$\Delta E_k = q\int_L E_x(t)dx.$$

This leads to the relationship

$$\Delta E_k = qV_0\sin(\phi)T, \tag{2.47}$$

where V_0 is the maximum voltage across the cavity, ϕ is the phase of the electric field as a particle passes the midpoint of the cavity and T is the transit time factor, which is given by:

$$T = \frac{\sin\left(\frac{L\pi}{\beta\lambda}\right)}{\frac{L\pi}{\beta\lambda}}. \tag{2.48}$$

The transit time factor scales the change in kinetic energy of a particle according to how the electric field varies as a particle passes through the cavity. For example, if the phase of the electric field changes by more than π whilst the particle is in the cavity, then the particle will experience both accelerating and decelerating forces and the efficiency of the cavity will be limited. Figure 2.11 shows the transit time factor vs. the fraction

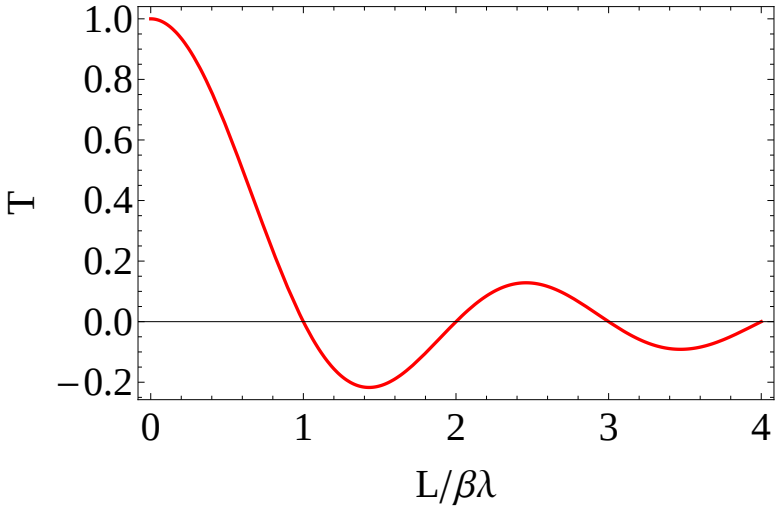

FIGURE 2.11: Transit time factor.

of the rf wavelength that a particle will encounter. It can be seen that the smaller the change in the electric field, the closer T is to unity. For the accelerators investigated in this thesis, T is taken to be equal to 1, and acceleration is modelled by an instantaneous kick to the energy of the particle as it reaches the location of the midpoint of an rf cavity. The magnitude of this kick, and whether the particle is accelerated or decelerated, is determined by the phase of the rf when the particle encounters the cavity. For net acceleration as a particle passes a single cavity over many turns, the kicks must add coherently; which requires careful selection of the rf frequency. At any given time, the momentum of a particle is equal to,

$$|\mathbf{p}|c = \sqrt{E_k(E_k + 2m_0c^2)} \tag{2.49}$$

Within the simulation codes used for the studies in this thesis, the kinetic energy of a particle is updated at an rf cavity according to $E_{k,f} = E_{k,i} + \Delta E_k$. The magnitude of the particle momentum after the kick is calculated using Eq. 2.49 and, given a change in momentum directed along the x axis, the individual components of the momentum are then found by setting:

$$p_{f,y} = p_{i,y}, \qquad p_{f,z} = p_{i,z},$$

followed by:

$$p_{f,x} = \sqrt{|\mathbf{p}_f|^2 - p_{f,y}^2 - p_{f,z}^2},$$

where $\mathbf{p}_i = (p_{i,y}, p_{i,z}, p_{i,x})$ is the initial momentum and $\mathbf{p}_f = (p_{f,y}, p_{f,z}, p_{f,x})$ is the momentum after the cavity. Selecting an rf frequency that is some harmonic of the revolution frequency:

$$f_{\mathrm{rf}} = \frac{h}{t_{\mathrm{rev}}}, \tag{2.50}$$

where h is an integer and t_{rev} is the time taken for an on-momentum particle to make a single revolution in the accelerator, means that a particle can arrive at the same phase (referred to as the synchronous phase) of rf voltage for numerous turns. Careful selection of the synchronous phase leads to focusing of the beam in longitudinal phase space, and maximises the number of particles that survive an acceleration cycle.

As an example, we consider the acceleration of low energy protons within a synchrotron accelerator. The particles that make up the beam are injected into the synchrotron with some range of momenta that are distributed around a design injection momentum. A particle with the design momentum will arrive at the first rf cavity when the rf field is at the synchronous phase (ϕ_s). However, dispersion in the lattice (which affects the path length of a particle over one turn in the accelerator) and a difference in particle speed means that a particle with some offset from the design momentum will arrive at the rf cavity either earlier or later than the design particle, and subsequently at a different phase of the rf field. For a low energy bunch, it is typically the difference in speed of an off-momentum particle that dominates over the change in path length in determining the arrival time of the particle at a cavity. A high-momentum particle within the bunch will arrive at a cavity earlier than a low-momentum particle.

For a high energy (relativistic) bunch, the change in particle speed with momentum will be small, and dispersion will instead determine the arrival time of the particle at a cavity. In this second case, it is the low-momentum particle that arrives at the cavity earlier. The energy at which dispersion comes to dominate over particle speed is a property of the accelerator, and is referred to as the transition energy.

The range of synchronous phase that will give longitudinal focusing and acceleration depends upon whether the design energy is below or above transition. In our low energy proton synchrotron, we assume that the bunch energy is below transition. Given Eq. 2.47, if the synchronous phase is above 0 and below $\pi/2$ radians, then the bunch will be accelerated, and individual particles within the bunch subject to longitudinal focusing.

During an acceleration cycle in a synchrotron, the revolution frequency of a bunch will increase; the rf frequency must also increase to ensure that a particle with the design momentum arrives at the synchronous phase of the rf every time the cavity is encountered. For further insight into the longitudinal dynamics of a particle, we can describe a system longitudinally with the Hamiltonian [37]:

$$H = h\left(\frac{\eta_0\delta^2}{2} + \frac{\eta_1\delta^3}{3} + ...\right) + \frac{eV_{\text{rf}}}{2\pi\beta^2 E_s}(\cos(\phi) - \cos(\phi_s) + (\phi - \phi_s)\sin(\phi_s)) \qquad (2.51)$$

where δ is the fractional offset in momentum of a particle from a reference momentum, V_{rf} is the rf accelerating voltage, β is the relativistic speed of a particle, E_s and ϕ_s are the synchronous energy and synchronous phase respectively and ϕ is the phase of a particle with respect to the rf voltage. The η terms are referred to as the phase slip factors, and describe the change in ϕ per turn of an off-momentum particle. Plotting contours

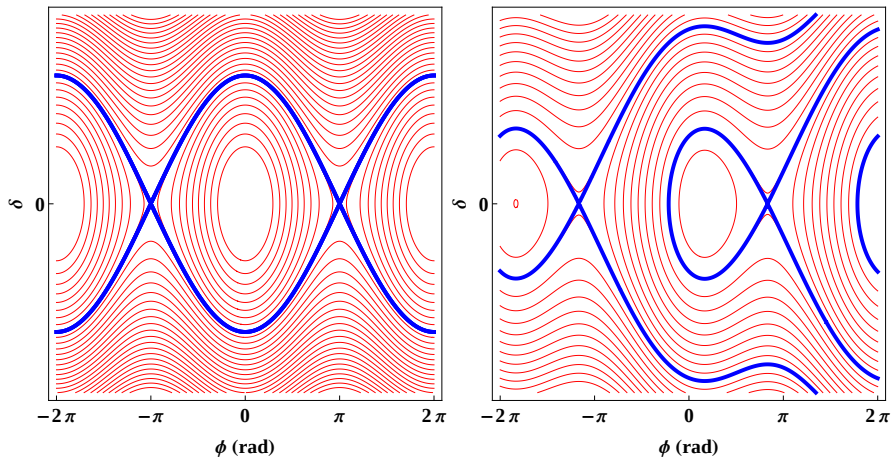

(a) $\phi_s = 0\,\text{rad}$ (b) $\phi_s = \pi/6\,\text{rad}$

FIGURE 2.12: Longitudinal phase space portraits for synchronous phases of 0 and $\pi/6\,\text{rad}$. Red lines show contours of constant Hamiltonian, and blue lines show the specific Hamiltonian values that define separatrices. Stable fixed points are observed when the phase of a particle is equal to the synchronous phase, and the closed blue contours that enclose the stable fixed points are referred to as rf buckets. In both plots, the phase slip factor is negative ($\eta < 0$), meaning that the lattice is below transition.

of constant Hamiltonian for Eq. 2.51 demonstrates how particles can be expected to behave in longitudinal phase space (Fig. 2.12). Stable fixed points are observed at ($\phi = \phi_s$, $\delta = 0$), whilst unstable fixed points are found at ($\phi = \pi - \phi_s$, $\delta = 0$). Plotting a contour for Eq. 2.51 for the Hamiltonian being equal to the value found at the unstable fixed point shows the separatrix in phase space, with the closed loop formed by the separatrix defining an rf bucket.

2.8 Beam dynamics in a synchrotron: HIT

The Heidelberg Ion Therapy (HIT) facility was first used for cancer treatment in November 2009, and by May 2013 had been used to deliver therapy to almost 2000 patients [38]. HIT features two horizontal beam lines and one gantry, which provides a beam that can rotate around a patient (with the tumour at the centre of rotation). Particles are delivered to the beamlines by a synchrotron that is routinely used to accelerate beams of protons and carbon ions for patient treatment, as well as helium and oxygen ions for research purposes. Treatment can be delivered through active scanning methods, with the extracted beam energy, size and intensity offered as variables. The time taken for one acceleration cycle (from particle injection to extraction) is dominated by the time needed to ramp the magnetic fields to provide a beam of the required energy. The ramp rate of the dipole magnets at HIT is $1.5\,\text{T/s}$, which leads to a repetition rate of $\approx 0.5\,\text{Hz}$ when extracting particles at the maximum energy of $430\,\text{MeV/u}$. To limit the amount of time taken to treat a tumour volume, a single acceleration cycle is used to treat many voxels. So that each of these voxels can be treated individually, the extraction of

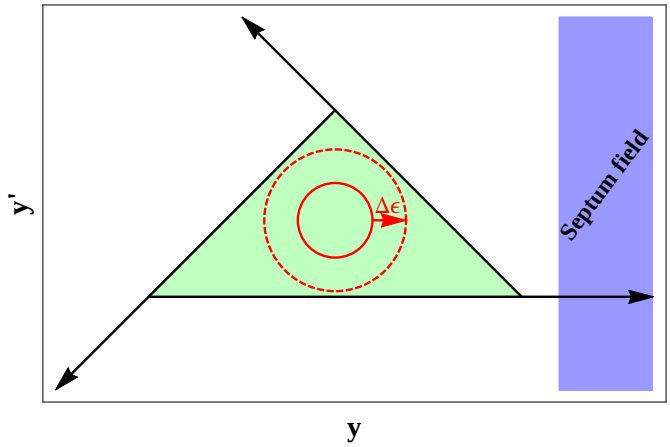

FIGURE 2.13: Emittance blow up of a beam (red), due to excitation by rf field, causes particles to cross the separatrix. Once a particle has left the stable region of phase space (green), its amplitude will increase until it encounters the field of a electrostatic or magnetic septum.

particles from the accelerator happens over many turns. This slow extraction has been possible at HIT through application of a method known as rf-knockout.

2.8.1 Slow extraction through third-integer resonance and rf knockout

An intensity controlled, continuous beam can be extracted from the Heidelberg synchrotron for up to 5 seconds. This is possible due to a method called rf knockout slow extraction. The method relies upon two separate perturbations to the beam motion, one from a sextupole field and a second from an rf cavity operating in the transverse electric mode.

third-integer resonance

Once a particle beam has been accelerated to the treatment energy, the strength of the quadrupole magnets is changed so that the horizontal tune is close to a third-integer (table 2.2), and a number of sextupoles are switched on in order to perturb the beam.

Figure 2.13 illustrates the situation in phase space once the quadrupoles and sextupoles have been adjusted. The boundary of the green triangle represents the separatrix that is created by the sextupole fields, when the tune is close to a third-integer tune. Particle motion within the separatrix is stable; particle motion outside the separatrix is unstable. The solid red circle represents the particle beam immediately after the magnets have been tuned on. In the stable region of phase space (coloured green), the beam will continue within the accelerator indefinitely. If a particle crosses the separatrix into the unstable region of phase space, then the amplitude of the motion of that will grow, and eventually the particle will encounter the field of a electrostatic or magnetic septum (Fig. 2.14) and be extracted. Control over the rate at which particles cross the separatrix means that the intensity of the extracted treatment beam can be varied. In earlier extraction systems, particles would become unstable (and be extracted) as the

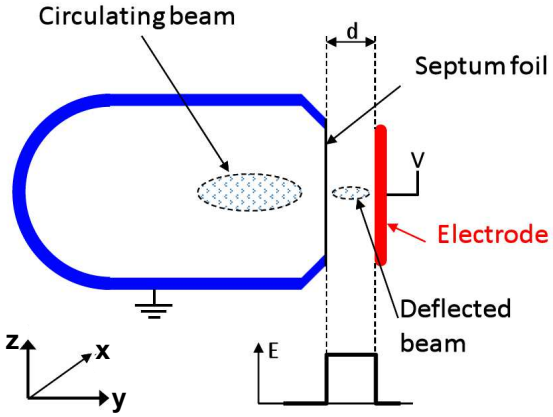

(a) Electrostatic septum

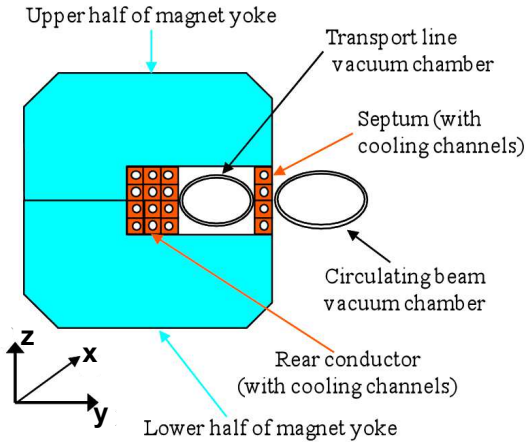

(b) Magnetic septum

FIGURE 2.14: Transverse cross section diagrams of electrostatic (Fig. (a)) and magnetic (Fig. (b)) septa (beam travels into, or out of, the page) [39]. A septum separates (along a transverse axis) dipole field and field-free regions of space: such components are often used to inject/extract beams into/out of rings. Taking extraction as an example, a circulating beam initially travels on the field free side of the septum. During resonant extraction, the betatron oscillation amplitude of a particle slowly increases, which eventually causes the particle to travel on the field side of the septum. A particle on the field side of the septum is steered towards further extraction components, which may include additional septa and a transport line for extracted beam.

quadrupole or sextupole strengths were altered in order to shrink the area contained by the separatrix. Varying the magnet ramping rate offered intensity control, however the variation in intensity was slow in comparison to that required by therapy.

RF knockout

For the rf knockout method, the area contained by the separatrix remains constant and to extract particles from the accelerator the transverse horizontal emittance of the beam is increased by using an rf cavity to give a transverse kick. In order for the transverse kicks to add coherently, the frequency of the rf cavity (f_{rf}) should resonate with the horizontal betatron motion,

TABLE 2.2: Summary of the HIT synchrotron parameters.

Lattice	
Circumference	64.986 m
Periods	6
Superperiodicity	2
Lattice type	Doublet
Injection	
Injection energy (kinetic energy per nucleon)	7 MeV/u
Injection momentum spread (δ)	$\pm 0.1\%$
Emittance	7 π mm mrad
Tune (Q_h/Q_v)	1.72/1.74
Extraction	
Extraction energy	48-430 MeV/u
Horizontal emittance	1-5 π mm mrad
Vertical Emittance	5 π mm mrad
Tune (Q_h/Q_v)	1.672/1.72

$$f_{\mathrm{rf}} = f_{\mathrm{rev}}(Q \pm m) \tag{2.52}$$

where m is an integer, and f_{rev} is the revolution frequency of the beam. Given the chromaticity of a lattice, as well as the amplitude dependence of tune and revolution frequency, it can be seen that a single rf frequency may be a driving term for a resonance with only a fraction of the total number of particles within a beam. To improve the extraction efficiency, the frequency of the rf cavity is modulated so that particles at the centre of a bunch diffuse out to amplitudes close to the separatrix. The intensity of the extracted beam is then controlled by varying the amplitude of the rf voltage. Given a non-uniform distribution of particles in phase space, a feedback loop between detectors that measure the intensity of the extracted beam and the rf amplitude allow for beams with only small fluctuations in intensity with time to be delivered to a patient.

2.9 FFAG based accelerators

2.9.1 Radial sector scaling FFAG: KEK

Following the successful construction and commissioning of a proof of principle radial sector scaling FFAG at the KEK laboratory, a second FFAG of this type was built to accelerate protons from 12 MeV to 150 MeV (Fig. 2.15); it is this second FFAG that is introduced here [40, 41], and the main parameters summarised in table 2.3. The accelerator, which has a a DFD cell structure, is the only one of the three FFAG examples given in this section to have been constructed. It was built as a prototype for a number of applications (including proton therapy) and has magnetic fields that vary in the midplane according to:

$$B_z = B_{z,0}\left(\frac{R}{R_0}\right)^k F(\theta), \tag{2.53}$$

FIGURE 2.15: The 150 MeV proton radial sector FFAG that was constructed/first operated at KEK, Japan. The accelerator has since been moved to Kyushu University, Japan [42]. The blue circular component located within the ring is a 12 MeV cyclotron that is used to accelerate protons up to the injection energy of the FFAG.

TABLE 2.3: Parameters of the radial sector scaling FFAG at KEK.

KEK lattice parameters	
Energy range	12-150 MeV
Number of cells	12
Cell type	DFD
k value	7.6
Maximum field (F/D)	1.63/0.78 T
Horizontal tune	3.69 - 3.8
Vertical tune	1.14 - 1.3

where R is the radius measured from the centre of the ring, R_0 is a reference radius, k is the scaling index, $B_{z,0}$ is the field at the reference radius, $F(\theta)$ is a periodic function and θ is the azimuthal angle. We have already shown, by the Taylor expansion of the scaling FFAG field profile (Eq. 2.12), that the field can be decomposed into multipole components. Whether a magnet is focusing or defocusing is set by the sign of the dipole field term, $B_{z,0}$; alternating gradient focusing within a radial sector FFAG therefore requires bending particles both into and against the curvature of the ring (with net bending into the ring). As with the synchrotron, the focusing and defocusing magnets are separated by drift spaces. For the KEK accelerator, the magnet boundaries lie along radii drawn from the ring centre, and the θ dependence of $B_{z,0}$ in Eq. 2.53 shows the azimuthal variation of the field. A consequence of the need for bending in two directions is that radial sector FFAGs have a larger circumference than a cyclotron or spiral sector FFAG (section 2.9.3) that is made with magnets of a comparable field strength and operates over the same momentum range.

The label scaling refers to how the closed orbit changes with momentum: as the momentum increases, the mean radius of the orbit will increase, but the shape of the orbit will remain identical to the low momentum orbit (i.e. the closed orbit scales). Given the drift spaces and reverse bending elements, the bending radius, ρ, of a particle within the magnet must be smaller than the radius of the closed orbit measured from the machine centre, R. A condition of the scaling law is that the ratio of the bending radius to the orbit radius is fixed for all momenta along any azimuthal angle:

$$\frac{\partial}{\partial p}\frac{\rho}{R}\bigg|_{\theta=\text{const}} = 0. \tag{2.54}$$

By using magnets that have the field profile of Eq. 2.53 the focusing strength of the lattice can be maintained with momentum:

$$\frac{\partial K}{\partial p}\bigg|_{\theta=\text{const}} = 0, \tag{2.55}$$

which gives a constant betatron tune with momentum, and is used to limit the effects of resonances on the betatron motion.

Given Eqs. 2.11, 2.53 and 2.54, it can be shown that the relationship between momentum and the average orbit radius is:

$$p = p_0 \frac{R}{R_0}^{k+1}. \tag{2.56}$$

A scaling index, k, of 7.6 for the KEK accelerator leads to a radial aperture that extends from 4.3 m to 5.47 m. Increasing the scaling index of the lattice would lead to a reduced aperture, but k is also related to the betatron tune and issues relating to the stability of the lattice must be considered (this is discussed further when we introduce the PAMELA FFAG).

The FFAG at KEK has been characterised both experimentally [41] and through tracking simulations [43]. It was found that imperfections in the scaling of the dynamics with momentum leads to a small tune shift ($\Delta Q \approx 0.1$). Extraction is carried out using a pulsed kicker and septum magnet [44]: the kicker magnet pushes particles towards the outside of the ring, where they can encounter the septum magnet. The septum magnet has a thin iron blade (with its length parallel to the direction of the beam): on the inner side of the blade there are negligible additional magnetic fields, whilst on the outside of the blade there is a field that steers the beam out of the accelerator. Using the kicker to give adequate radial separation between the penultimate and final turn of the beam in the accelerator is key in obtaining good extraction efficiency (minimising beam loss to the septum blade). The relative positioning of the kicker and septum is chosen to maximise the effect of the kicker, given:

$$\Delta y = \theta_{\text{kick}} \sqrt{\beta_{\text{kicker}} \beta_{\text{septum}}} \sin\psi, \tag{2.57}$$

Table 2.4: PAMELA proton lattice parameters.

PAMELA proton lattice parameters	
Energy range	30.95-250 MeV
Number of cells	12
Cell type	FDF
k value	36.721
r_0	6.251 m
B_0 (F/D)	1.7282/-2.0 T
Horizontal tune	8.64
Vertical tune	3.24
Orbit excursion	0.176 m

where Δy is the change in orbit position at the septum, θ_{kick} is the deflection angle of the beam by the kicker, β_{kicker} and β_{septum} are the β functions at the kicker and septum respectively and ψ is the phase advance of the beam from the kicker to the septum. The largest change in the orbit position is attained when the phase advance between the kicker and septum is given by $\psi = \left(n + \frac{1}{2}\right)$, where n is an integer. The extraction process occurs over a couple of turns, and is referred to as fast extraction, and results in a pulsed extracted beam with a frequency that is determined by the repetition rate of the accelerator (which is 100 Hz). No options for varying the extraction energy of the KEK FFAG are discussed within the literature reviewed.

2.9.2 Non-scaling FFAG: PAMELA

The PAMELA lattice [29] is a non-scaling FFAG design concept which consists of two concentric rings (Fig. 2.16): An inner ring accelerates protons from 31 MeV to 250 MeV, and carbon ions from 8 MeV/u to 68 MeV/u; the outer ring is then used to accelerate the carbon ions extracted from the first ring to 400 MeV/u. Here we will focus on the inner ring only.

The inner ring consists of 12 FDF triplet cells, which, as with the radial sector scaling FFAG, provide alternating gradient focusing whilst bending the beam both forwards and backwards. However, there are several key differences between PAMELA and the FFAG at KEK, these are:

1. In PAMELA, the field profiles are described by a Taylor series expansion of the scaling law (Eq. 2.12) that has been truncated at the decapole term. PAMELA is a non-scaling FFAG, however the non-linear field terms that are present lead to a low chromaticity.

2. The horizontal phase advance per cell of the betatron motion within PAMELA is more than π radians. This allows for a large scaling index to be used ($k = 36.721$), which reduces the orbit excursion during acceleration. The first two regions of stability for Hill's equation are demonstrated for a FDF cell in Fig. 2.18 [45]. For the purpose of the plot, we represent the focusing term in Hill's equation by $g(\theta)$,

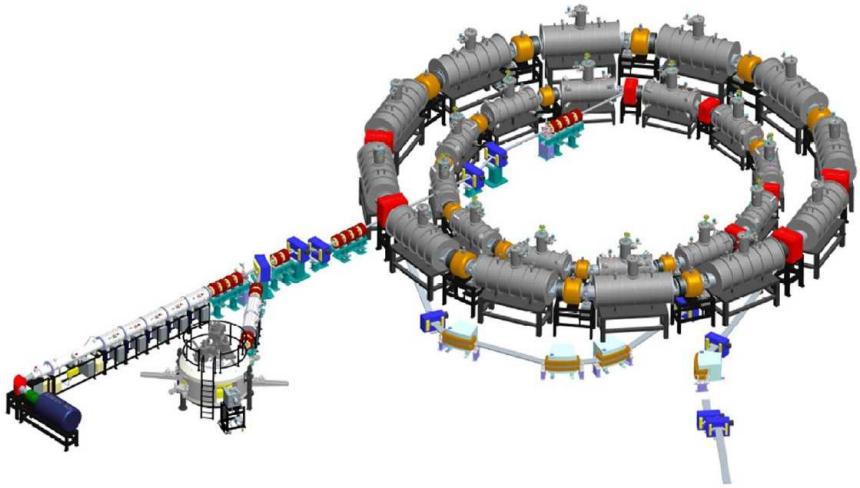

FIGURE 2.16: The PAMELA design [29]. The inner ring would accelerate through the full range of treatment energies for protons and up to 68 MeV/u for carbon ions. The outer ring would accelerate carbon ions from 68 to 400 MeV/u.

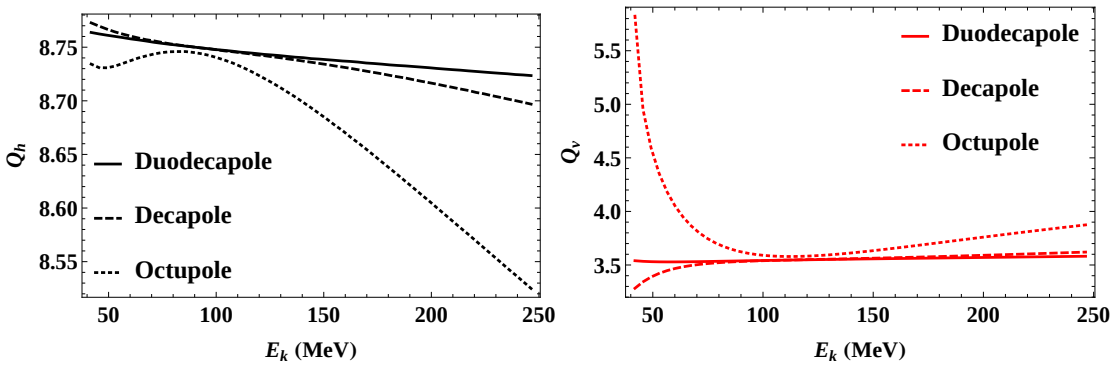

(a) Horizontal (b) Vertical

FIGURE 2.17: The PAMELA lattice has a low chromaticity due to the inclusion of higher order field terms in Eq. 2.12. The above shows the horizontal and vertical tune profiles when the series expansion is truncated at the octupole, decapole and duodecapole term.

where θ indicates the variation of the focusing strength with the azimuthal angle around the ring centre:

$$y'' + g(\theta)y = 0,$$
$$z'' - g(\theta)z = 0.$$

We give each of the magnets within the cell an aperture of one third of the total aperture of the cell (θ_{cell}), and $g(\theta)$ is given by:

$$g(\theta) = \begin{cases} g_0 + g_1 & \text{for } 0 < \theta < \frac{1}{3}\theta_{\text{cell}}, \\ g_0 - g_1 & \text{for } \frac{1}{3}\theta_{\text{cell}} < \theta < \frac{2}{3}\theta_{\text{cell}}, \\ g_0 + g_1 & \text{for } \frac{2}{3}\theta_{\text{cell}} < \theta < \theta_{\text{cell}}, \end{cases}$$

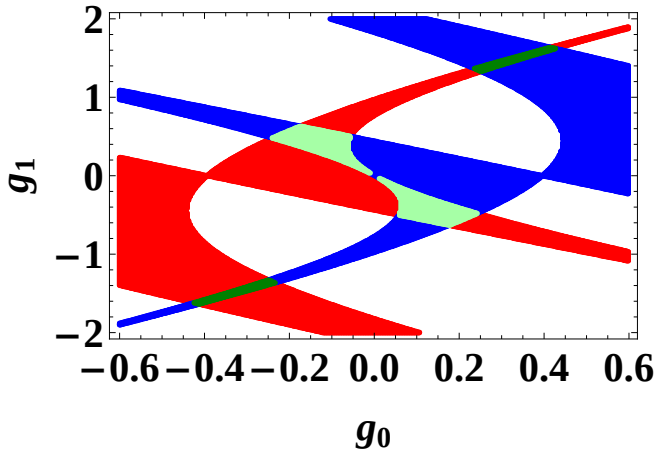

FIGURE 2.18: Stability plot for a FDF triplet. Conventionally the phase advance of a beam per period is between 0 and π radians; the focusing strengths for which this criteria is met for both the horizontal and vertical axes are marked in lighter green in the plot, and this region is referred to as the first stability region of Hill's equation. PAMELA is designed to operate in the second stability region of Hill's equation, which is marked in darker green; the phase advance per period in this region is between π and 2π radians. The regions marked blue and red give stability respectively for the horizontal and vertical axis only.

where g_0 is the average focusing strength of the entire cell for the horizontal axis ($-g_0$ for the vertical axis) and g_1 gives the difference for each magnet between the average and local focusing strength. The issue with using a larger scaling index and constructing a lattice that operates with a period horizontal phase advance greater than π radians is that the increased non-linearity of the magnetic field with radius will significantly limit the dynamic aperture of the ring. A tracking study for the PAMELA lattice [46] found a normalised dynamic aperture of approximately $30\,\pi$ mm mrad for the error free lattice, whilst a previous study [45] found that the inclusion of an rms alignment error of $50\,\mu$m on the magnets within a lattice operating in a similar regime as PAMELA led to a reduction in the dynamic aperture of approximately two thirds. In either case, the studies indicate that the dynamic aperture of PAMELA would be large enough for the proton therapy application.

3. The magnets are rectangular (rather than sector shaped), and are placed along a straight line within in a cell (rather than along an arc). It is intended that this will help to reduce construction costs and simplify the alignment process during machine commissioning.

A requirement assumed for the PAMELA design is that it should be possible to vary the extraction energy by 40 MeV in less than 1 second; one of the obstacles to realising this within a FFAG is the shift in the horizontal closed orbit with momentum. Extracting the beam horizontally from the lattice using a kicker magnet and septum was investigated, however it was found that the requirements of the beam dynamics could not be met. In particular, extraction of the low energy treatment protons (70 MeV)

needs the kicker to produce an orbit separation of 10 cm, the non-linearity of the field leads to significant amplitude detuning of the bunch at the required amplitudes and ultimately to particle loss. If it were possible to obtain the required orbit separation, then a second problem of the angular dispersion at the extraction septum poses an additional challenge [47]. For these reasons, vertical extraction was considered to be the best solution for the PAMELA lattice.

Both kicker and resonance-based extraction methods have been investigated for vertical extraction from PAMELA [48]. In comparison with the kicker-based horizontal extraction scheme, the kicker strength required for vertical extraction is much lower and it is simpler to match the post-extraction transport line to the dynamics of the bunch. The resonant extraction method investigated is based on particles crossing a half-integer vertical tune, with the half-integer tune set to coincide with the required extraction energy by varying the F/D ratio and the resonance driven by introducing quadrupole field errors. Upwards and downwards crossing of the resonance have been investigated, and it was found that downward crossing gives the output that is closest to being monoenergetic. This is explained after first considering a positive tune shift with increasing particle amplitude; for upward crossing of a half-integer tune, the exact energy at which a particle crosses the half-integer tune is determined by its amplitude, with large amplitude particles crossing the resonance and being extracted first. For downward crossing, it is the particles at the centre of the bunch that encounter the resonance first, once the amplitude of a particle that is initially at the centre of the bunch has grown to match that of a higher amplitude particle, then both particles will have approximately the same tune and their amplitudes will grow at a similar rate to one another. Yokoi [48] highlights that these extraction methods are difficult to control, and that although extraction happens over a number of turns, it is not necessarily slow extraction (when compared to synchrotron slow extraction that can take place over a few seconds).

2.9.3 Spiral sector scaling FFAG: RACCAM

The Research in ACCelerators and Applications in Medicine (RACCAM) design concept is a spiral sector scaling FFAG (Fig. 2.19), for which the midplane field varies according to:

$$B_z = B_{z,0}\left(\frac{R}{R_0}\right)^k \left\{1 + f\sin\left[N\theta - N\tan(\zeta)\ln\left(\frac{R}{R_0}\right)\right]\right\}, \tag{2.58}$$

where R is the radius measured from the centre of the ring, R_0 is the reference radius, f describes the field flutter (the field modulation with θ), N is the number of periods and θ is the azimuthal angle. The periodic cell of the RACCAM lattice (table 2.5) contains just one magnet type; however, alternating gradient focusing is achieved through use of edge focusing, where the beam receives a vertical kick that is dependent upon the entrance angle of the beam measured with respect to the normal of the magnet face. In order to keep constant vertical focusing, the profile of the magnet edge is shaped so that the nominal beam entrance angle remains constant with momentum. The term ζ

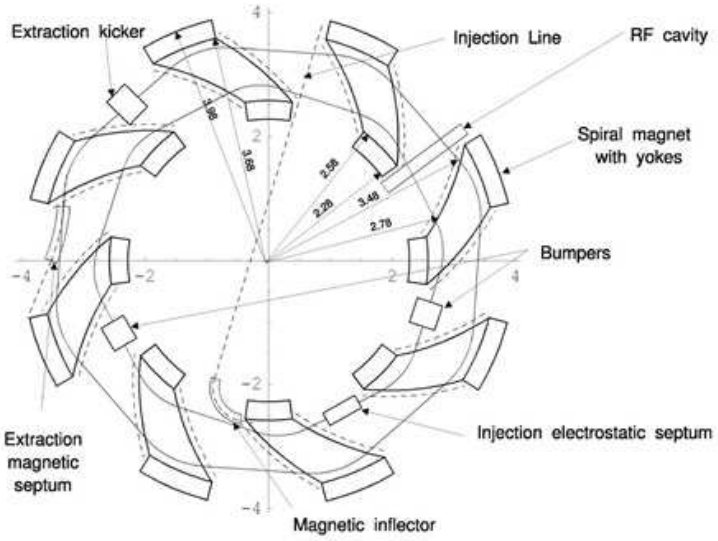

FIGURE 2.19: Diagram showing the RACCAM spiral sector FFAG design [49]. The use of edge focusing means that alternating gradient focusing could be achieved within RACCAM without having to include reverse bending.

TABLE 2.5: Parameters for the RACCAM design.

RACCAM lattice parameters	
Injection energy range	5.549-15 MeV
Extraction energy range	70-180 MeV
Number of cells	10
Cell type	Spiral
k value	5
Spiral angle, ζ	53.7°
$r_{min/max}$	2.79/3.46 m
Orbit excursion	0.67 m
B_{max}	1.7 T
Horizontal tune (inj/ext)	2.758/2.761
Vertical tune (inj/ext)	1.549/1.603

in Eq. 2.58, which is referred to as the spiral angle, gives the constant angle between a line drawn from the ring centre to a point in the midplane on the magnet edge and the normal to the magnet edge at that same point. The equation that describes the shape of a magnet edge is:

$$R = R_0 \exp\left(-\frac{\theta}{\tan\zeta}\right).$$

For the RACCAM design, extraction occurs during one turn by using a kicker to deflect particles out of the ring. A scheme that allows for variable extraction energy whilst maintaining the injection/extraction radius and betatron tune has been investigated [28]; to vary the extraction energy, both the injection energy and magnetic flux density have to be controlled in order to keep the field strength constant. With the field strength fixed, the momentum gain between the injection and extraction radius is also fixed, to a factor of ≈ 3.6, and as a result it is possible to tune the extraction

energy to anything between 70 and 180 MeV. The fixed dynamics at all output energies should make the extraction of particle bunches to a matched transport line relatively straightforward. However, the non-linearity of the magnet excitation poses an additional challenge; field map modelling of the RACCAM accelerator has shown that for the variable extraction energy scheme considered, the field profile of the magnets does not scale linearly with changing current due to the saturation of the magnet steel. As a result, there is a variation of betatron tune with extraction energy, which can limit the dynamic aperture at some energies (due to resonances). For proton therapy it is important that the extraction energy can be varied quickly and accurately: it is not clear to what extent the extraction method presented for the spiral sector FFAG meets these requirements.

Chapter 3

Tracking methods for accelerator design

Numerous computer codes have been developed for modelling the motion of particles through an accelerator. Often, assumptions made in the tracking methods mean that a particular code is suitable only for accelerators operating in a certain regime, for example, a code may only be relevant for ultra-relativistic particles or for motion that does not deviate far from a reference axis. There are two main approaches to particle tracking: the first is based on the numerical integration of the Lorentz force experienced by a particle as it passes through the different accelerator components, whilst the second involves developing transfer maps that can be used to propagate the phase space coordinates through discrete elements (which often cover complete components). Numerical integrators offer good flexibility, it is straightforward to introduce new elements and to track through field maps that represent real components (e.g. based on magnetic measurements). However, care has to be taken to ensure that the integration step length is small enough for the model to offer a good approximation of the path of a real particle. For this reason, numerical integration can often be more CPU intensive than using transfer maps. The coefficients within a transfer map can give important information about the optics of an accelerator, and a map can be constructed in such a way as to maintain symplecticity. However, implementing a transfer map for a new component so that it may be included in a tracking code is not necessarily straightforward.

Most of the tracking studies within this thesis have been carried out using the Zgoubi tracking code [50]. A problem with using a code that has undergone extensive development (such as Zgoubi) is that the underlying physics may not be comprehensively described by the user manual. Before deciding upon Zgoubi, a number of options were considered, including the implementation of a code based upon a new numerical integration method. The aim of this chapter is to further discuss several approaches to particle tracking, including the methods used in Zgoubi, and to justify the choice of Zgoubi as the preferred tracking code.

3.1 Transfer map tracking

A transfer map propagates a set of phase space coordinates through a discrete element. Tracking codes based on transfer maps generally implement maps for standard accelerator components (such as dipoles, quadrupoles and rf cavities). Codes also exist that can construct transfer maps based on some appropriate specification of the fields within a given component. Transfer maps can transport particles from the start of a component to the end in a single step, which can be less CPU intensive than some alternative methods (such as numerical integration), or the component can be subdivided into a number of smaller steps if more information about the particle trajectory is needed in a particular case.

As an example, we consider the construction of a transfer map in the form of a Taylor series for a quadrupole. The Hamiltonian for a relativistic particle in an electromagnetic field is given by:

$$H = \sqrt{(\mathbf{p} - q\mathbf{A})^2 c^2 + m_0^2 c^2} + q\phi, \tag{3.1}$$

where $\mathbf{p}$ is the canonical momentum ($\mathbf{p} = \boldsymbol{\beta}\gamma m_0 c + q\mathbf{A}$), $\mathbf{A}$ is the vector potential and ϕ is the scalar potential. The electric and magnetic fields are given by $\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t}$ and $\mathbf{B} = \nabla \times \mathbf{A}$. The equations of motion for a particle in an electromagnetic field then follow from Hamilton's equations, which are given by:

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}, \quad \frac{dq_i}{dt} = \frac{\partial H}{\partial p_i}, \tag{3.2}$$

where (q_i, p_i) are pairs of canonical variables (coordinate q_i and conjugate momentum p_i), and the independent variable is the time, t. An advantage of using a Hamiltonian to describe the system is that Hamilton's equations allow for the distance along a beam line, s, to be used as an independent variable rather than time. Using s as the independent variable means that the change in particle dynamics over an element of length L (which generally corresponds to the length of a component within the lattice, for example a dipole or quadrupole), may be calculated easily. It is convenient to work with the canonical variables (y, p_y), (z, p_z) and (x, δ). Here the definition of δ differs from that given earlier, and is the energy deviation $\delta = \frac{E}{p_0 c} - \frac{1}{\beta_0}$, where E is the energy of a particle and p_0 and β_0 are respectively the momentum and normalised speed of a reference particle (this definition of δ is valid for this section only). For a magnetic quadrupole, the vector potential, $\mathbf{A}$, is given by $A_y = A_z = 0$ and $A_x = -\frac{1}{2}K_1(y^2 - z^2)$. The Hamiltonian for a particle within quadrupole field is therefore:

$$H = \frac{\delta}{\beta_0} - \sqrt{\left(\delta + \frac{1}{\beta_0}\right)^2 - p_y^2 - p_z^2 - \frac{1}{\beta_0^2\gamma_0^2}} + \frac{1}{2}K_1(y^2 - z^2), \tag{3.3}$$

and the equations of motion, with the path length, s, as the independent variable are:

$$\begin{aligned} \frac{dy}{ds} &= \frac{\partial H}{\partial p_y}, \quad \frac{dp_y}{ds} = -\frac{\partial H}{\partial y}, \\ \frac{dz}{ds} &= \frac{\partial H}{\partial p_z}, \quad \frac{dp_z}{ds} = -\frac{\partial H}{\partial z}, \\ \frac{dx}{ds} &= \frac{\partial H}{\partial \delta}, \quad \frac{d\delta}{ds} = -\frac{\partial H}{\partial x}. \end{aligned} \tag{3.4}$$

At this stage, the order to which a Taylor series expansion is performed for the Hamiltonian determines the range of applications to which a transfer map will be appropriate. A second order expansion is a common choice, and is referred to as the paraxial approximation given that the resulting field maps are only accurate in describing motion that is almost parallel to the reference axis. For the quadrupole, taking only second order terms of the dynamical variables gives:

$$H_2 = \frac{1}{2}p_y^2 + \frac{1}{2}p_z^2 + \frac{1}{2}K_1 y^2 - \frac{1}{2}K_1 z^2 + \frac{1}{2\beta_0^2\gamma_0^2}\delta^2 \tag{3.5}$$

By applying Hamilton's equations to the second order Hamiltonian for the dynamical variables for the transverse axes, we then obtain equations of motion that are linear in the dynamical variables:

$$\begin{aligned} \frac{\mathrm{d}y}{\mathrm{d}s} &= p_y, \quad \frac{\mathrm{d}p_y}{\mathrm{d}s} = -K_1 y, \\ \frac{\mathrm{d}z}{\mathrm{d}s} &= p_z, \quad \frac{\mathrm{d}p_z}{\mathrm{d}s} = K_1 z, \end{aligned}$$

whilst for the longitudinal variables, we obtain:

$$\frac{\mathrm{d}x}{\mathrm{d}s} = \frac{\delta}{\beta_0^2\gamma_0^2}, \quad \frac{\mathrm{d}\delta}{\mathrm{d}s} = 0. \tag{3.6}$$

The equations for the transverse axes are identical to those already presented as Hill's equations in chapter 2 (Eq. 2.4). By solving the equations of motion, we can obtain the transfer map:

$$M = \begin{pmatrix} \cos(\omega L) & \frac{\sin(\omega L)}{\omega} & 0 & 0 & 0 & 0 \\ -\omega\sin(\omega L) & \cos(\omega L) & 0 & 0 & 0 & 0 \\ 0 & 0 & \cosh(\omega L) & \frac{\sinh(\omega L)}{\omega} & 0 & 0 \\ 0 & 0 & -\omega\sin(\omega L) & \cosh(\omega L) & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \frac{L}{\beta_0^2\gamma_0^2} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{3.7}$$

where $\omega = \sqrt{|k_1|}$. In addition to the transverse motion, we also have terms that describe the longitudinal position of a particle relative to the design particle. It is noted that there are features of the motion of a particle that will be missing when tracking a particle

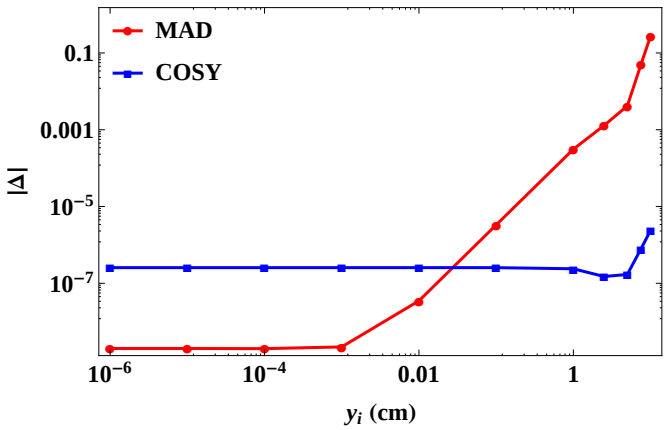

FIGURE 3.1: A 100 FODO cell line has been modelled using MAD, COSY (11^{th} order map) and Zgoubi. For a given offset at the start of the line (y_i), the absolute fractional difference at the end of the line (relative to the Zgoubi result, so that $\Delta = (y_f - y_{f,zgoubi})/y_{f,zgoubi}$) is given for MAD and COSY. The higher order COSY map is able to more accurately calculate the particle trajectory when the initial transverse offset is large, however calculating the coefficients of such a map can take as much CPU time as the numerical integration method (Zgoubi).

through a quadrupole with Eq. 3.7. For example, there are no terms in Eq. 3.7 that will affect the betatron tune if a particle deviates from the reference momentum or has a large transverse amplitude. These aspects of the particle motion were lost when we made the paraxial approximation. Further, were we to follow the same method for a sextupole magnet, with $A_x = -\frac{1}{6}K_2(y^3 - 3yz^2)$, then the linear transfer map obtained is equal to that of a drift space. The appropriate order of the series expansion is dependent upon the requirements of a given tracking simulation.

MAD (Methodical Accelerator Design) [51] is a code that can use maps that are calculated after truncating the Taylor series expansion of the Hamiltonian to third order in the dynamical variables (giving a transfer map with terms up to second order in the dynamical variables). The effects of octupole and higher order components are modelled by a kick to the transverse variables p_y and p_z made under the thin lens approximation. There are numerous examples of MAD being used to calculate the optics of non-scaling FFAGs [52–54], however the non-paraxial motion in such machines may mean that the MAD results are of limited accuracy for tracking over a large number of turns.

COSY is a tracking code that allows for the calculation and application of transfer maps of arbitrary order. The maps are calculated following the specification of an accelerator using either COSY's library of analytical expressions for electromagnetic fields or by inputting measured or simulated field data. COSY incorporates a symplectic numerical integrator that is used to track particles through fields, the results from which are then used to form maps. The exact methods used by COSY are not discussed in detail in this thesis, but can be found in [55–57]. COSY offers advantages for tracking through FFAG, including being able to accurately track non-paraxial particles.

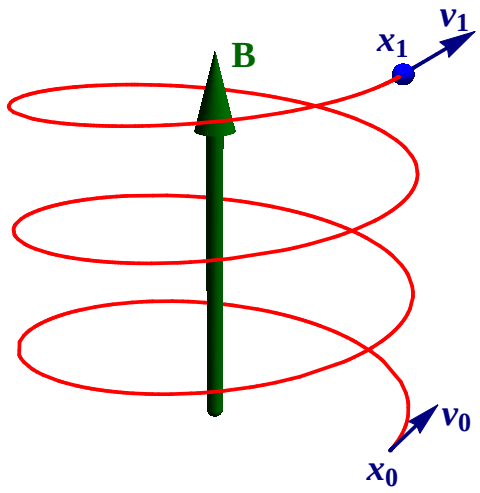

FIGURE 3.2: Helical motion of a charged particle in a uniform magnetic field. The arc integration method for particle tracking makes the approximation that the magnetic field is uniform over an integration step length, and that the motion is described by a helix.

3.2 The arc method for particle tracking

The arc method is a simple method for tracking particles through magnetic fields that was devised as an option for the tracking work carried out for this thesis. A code based upon the method was written in the C programming language; an advantage of using a self-developed code is that the methods for tracking are completely understood, however there is the drawback that significant further development would be required in order to attain a level of functionality comparable with the other codes tested. The arc method was therefore used only to benchmark more complete tracking codes, and to gain a better understanding of the tracking methods they employ. In this section we will summarise the arc method of particle tracking.

In the absence of synchrotron radiation, the exact path of a charged particle that has some component of velocity perpendicular to a uniform magnetic field is helical and can be calculated readily. Consider a particle of velocity $\mathbf{v_0}$, at a location $\mathbf{x_0}$ as it travels through a uniform magnetic field, $\mathbf{B}$. By taking the unit vector along the direction of the magnetic field:

$$\hat{\mathbf{B}} = \frac{\mathbf{B}}{B},$$

the components of velocity that are parallel and perpendicular to the direction of the magnetic field can be calculated as:

$$\mathbf{v}_{\parallel} = (\mathbf{v_0}.\hat{\mathbf{B}})\hat{\mathbf{B}},$$
$$\mathbf{v}_{\perp} = \mathbf{v_0} - \mathbf{v}_{\parallel}.$$

Equating the Lorentz force to the centrifugal force allows for the radius of the helical path to be calculated as:

$$\rho = \frac{mv_\perp}{qB},$$

where m is the relativistic mass. To obtain the angle subtended by a particle over one integration step length, the length of the arc projected onto the plane perpendicular to the magnetic field is used, $s_\perp = v_\perp \delta t$, where δt is the integration time interval:

$$\delta\theta = \frac{v_\perp \delta t}{\rho} = \omega \delta t.$$

The change in position of the particle in the plane perpendicular to the magnetic field is then based on circular motion:

$$\Delta \mathbf{x}_\perp = (\rho - \rho\cos\delta\theta)\hat{\mathbf{F}} + (\rho \sin\delta\theta)\hat{\mathbf{v}}_\perp,$$

where $\hat{\mathbf{F}}$ is the unit vector in the direction of the Lorentz force. The position of the particle after the time interval δt is:

$$\mathbf{x_1} = \mathbf{x_0} + (\rho - \rho\cos\delta\theta)\hat{\mathbf{F}} + (\rho\sin\delta\theta)\hat{\mathbf{v}}_\perp + \mathbf{v}_\parallel \delta t. \tag{3.8}$$

The velocity of the particle is found by using the derivative of Eq. 3.8 with respect to δt:

$$\mathbf{v_1} = (\rho\omega\sin\delta\theta)\hat{\mathbf{F}} + (\rho\omega\cos\delta\theta)\hat{\mathbf{v}}_\perp + \mathbf{v}_\parallel.$$

Unlike the map based tracking methods that were introduced in the previous section, the independent variable when tracking with the arc method is time (s is no longer the distance along the beam line, but is instead the path length of an individual particle). The time taken for a particle to reach the boundary of a component is dependent upon the dynamics of that particle, and iterative methods have to be used to find the exact time at which the particle reaches the boundary; we used the Newton-Raphson method when implementing the arc method.

The arc method has been used to track particles through a number of lattices, including a synchrotron and a spiral sector FFAG, for which the tracking results agreed qualitatively with the expectation of the dynamics within the systems studied [58]. Later in this chapter we will present a comparison of results obtained using the arc method to track particles through an arbitrary multipole field with those obtained when using the Zgoubi tracking code.

3.3 The Zgoubi tracking code

Zgoubi calculates particle trajectories based upon the numerical integration of the Lorentz force equation. Tracking can take place through regions that have electromagnetic fields defined analytically or by a field map. As with the arc method, the independent variable

is time, and iterative methods are again used to find the beam dynamics at the boundary of a hard edge magnet.

Zgoubi has an extensive library of analytical representations of accelerator components, including dipole magnets (based on Cartesian and polar coordinate systems), multipoles and rf cavities. A number of components have been developed specially for FFAG modelling, including radial and spiral sector scaling FFAG elements.

3.3.1 The Zgoubi tracking method

Starting by writing the particle momentum in terms of the particle rigidity and a unit vector in the direction of the velocity of the particle:

$$m\mathbf{v} = mv\hat{\mathbf{v}} = qB\rho\hat{\mathbf{v}}, \tag{3.9}$$

then the Lorentz force equation (Eq. 2.1) can be written as:

$$\frac{d(B\rho)}{ds}\hat{\mathbf{v}} + B\rho\frac{d\hat{\mathbf{v}}}{ds} = \frac{\mathbf{E}}{v} + \hat{\mathbf{v}} \times \mathbf{B}. \tag{3.10}$$

The propagation of the particle position and velocity are given by:

$$\mathbf{x_1} = \mathbf{x_0} + \int_s \hat{\mathbf{v}}_\mathbf{0}\,\mathrm{d}s, \tag{3.11}$$

$$\hat{\mathbf{v}}_\mathbf{1} = \hat{\mathbf{v}}_\mathbf{0} + \int_s \hat{\mathbf{v}}'_\mathbf{0}\,\mathrm{d}s, \tag{3.12}$$

where $\mathrm{d}s = \mathbf{v}\,\mathrm{d}t$ is the distance along the path length of the particle (rather than a reference trajectory). The new particle position and velocity are then found approximately, following a Taylor series expansion of Eq. 3.11 and 3.12, by:

$$\mathbf{x_1} \approx \mathbf{x_0} + \hat{\mathbf{v}}_\mathbf{0}\Delta s + \frac{d\hat{\mathbf{v}}_\mathbf{0}}{ds}\frac{\Delta s^2}{2!} + ... + \frac{d^5\hat{\mathbf{v}}_\mathbf{0}}{ds^5}\frac{\Delta s^6}{6!}, \tag{3.13}$$

$$\hat{\mathbf{v}}_\mathbf{1} \approx \hat{\mathbf{v}}_\mathbf{0} + \frac{d\hat{\mathbf{v}}_\mathbf{0}}{ds}\Delta s + \frac{d^2\hat{\mathbf{v}}_\mathbf{0}}{ds^2}\frac{\Delta s^2}{2!} + ... + \frac{d^5\hat{\mathbf{v}}_\mathbf{0}}{ds^5}\frac{\Delta s^5}{5!}. \tag{3.14}$$

The methods are considered for tracking through magnetic fields only ($\mathbf{E} = 0$), in which case the rigidity is fixed and:

$$\frac{\mathrm{d}\hat{\mathbf{v}}}{\mathrm{d}s} = \hat{\mathbf{v}} \times \frac{\mathbf{B}}{B\rho}. \tag{3.15}$$

Differentiating Eq. 3.15 gives the appropriate order of $d^n\hat{\mathbf{v}}/ds^n$ for Eq. 3.13, where, for example, the product rule gives:

$$\frac{d^2\hat{\mathbf{v}}}{ds^2} = \frac{d\hat{\mathbf{v}}}{ds} \times \frac{\mathbf{B}}{B\rho} + \hat{\mathbf{v}} \times \frac{1}{B\rho}\frac{d\mathbf{B}}{ds}. \tag{3.16}$$

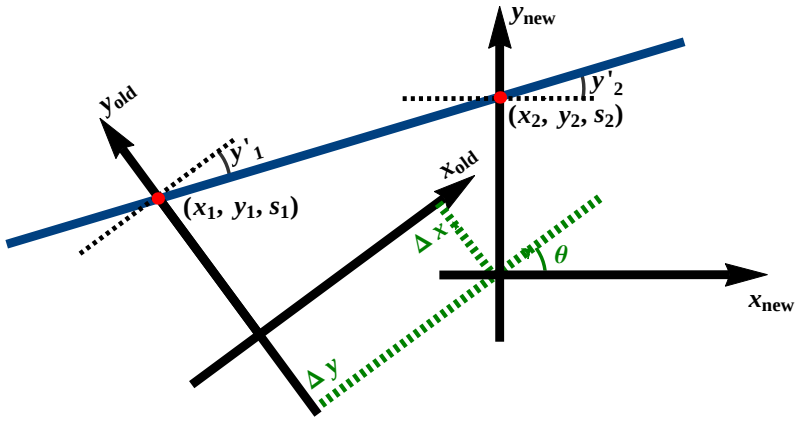

FIGURE 3.3: A particle at the coordinates (x_1, y_1, s_1) is transformed to a reference frame that is offset from the original frame by Δx, Δy and a rotation of θ. The transformation involves transporting the particle through a drift space to the point (x_2, y_2, s_2). The blue line shows the path of the particle and the red points show the location of the particle before and after the transformation.

The derivatives of $\mathbf{B}$ and $\hat{\mathbf{v}}$ are found alternately by:

$$\frac{d\mathbf{B}}{ds} = \sum_i \frac{\partial \mathbf{B}}{\partial x_i} \hat{v}_i, \tag{3.17}$$

$$\frac{d^2\mathbf{B}}{ds^2} = \sum_{ij} \frac{\partial^2 \mathbf{B}}{\partial x_i \partial x_j} \hat{v}_i \hat{v}_j + \sum_i \frac{\partial \mathbf{B}}{\partial x_i} \hat{v}_i', \tag{3.18}$$

and so on. The Zgoubi implementation of the above methods allows for tracking to be carried out following truncation of the Taylor series expansion (Eq. 3.13) to between the second and sixth order terms of Δs, with fourth order being the default.

3.3.2 Reference axis transformations

Zgoubi tracks particles using local coordinates (Fig. 2.4), and has components based on both polar and Cartesian coordinate systems. In the case of a polar system, the reference axis changes smoothly as a function of the azimuth of the component. For a Cartesian coordinate system, the reference axis is fixed, meaning that additional transformations must be made in order to describe a circular accelerator using the Cartesian based components only. The changes in the reference axis shown in Fig. 3.3 lead to the following

transformations in the dynamical variables of a particle:

$$\begin{aligned} y_2' &= y_1' - \theta, \\ y_2 &= \frac{(y_1 - \Delta y)\cos(y_1') + \Delta x \sin(y_1)}{\cos(y_2')}, \\ L^2 &= (\Delta x - y_2 \sin(\theta))^2 + (\Delta y - y_1 + y_2 \cos(\theta))^2, \\ z_2 &= z_1 + L \tan(z_1'), \\ z_2' &= z_1', \\ s_2 &= s_1 + \frac{L}{\cos(z_1')}. \end{aligned} \tag{3.19}$$

where Δy and Δz offset the start of the new reference axis from the end of the initial reference axis, and θ gives the rotation of the new reference axis with respect to the initial reference axis.

The above transformations are valid when the transformation is made in a region with negligible electromagnetic fields and L is the length of drift space that a particle is transported through during the transformation from the first reference frame to the second.

3.3.3 Field map tracking

As well as providing a range of analytical descriptions of electromagnetic fields, the Zgoubi code allows for particles to be tracked through field maps. In chapter 4, particles are tracked through a field map for the ns-FFAG EMMA. The maps used are defined in the yx plane (referred to as midplane field maps), and consist of a grid of equally spaced Cartesian coordinates, with the magnetic field components defined for each set of coordinates. The symmetry of the cell that is studied in chapter 4 means that a magnetic field will be perpendicular to and continuous across the midplane; in this case, $B_y = B_x = 0$, and the field map only has non-zero values for B_z. From the midplane field map, Zgoubi uses interpolation methods to calculate an estimated value for the magnetic field at any point on the midplane, and extrapolation based on Maxwell's equations to estimate the magnetic field at a point above or below the midplane.

3.4 Comparing numerical methods

Neither of the integrators described above force the tracking results to be symplectic. Attention is paid to the convergence of tracking simulations (when reducing the integration step length) to a solution that should be exactly symplectic, in which case, a deviation from symplecticity in the tracking results can be a measure of the accuracy of the solution. Figure 3.4 demonstrates the convergence of the Zgoubi integration (with Eq. 3.13 up to second order) and the arc integration method to a correct solution. In Fig. 3.4a an integration step length of 1 cm is used: the effects of non-symplecticity can

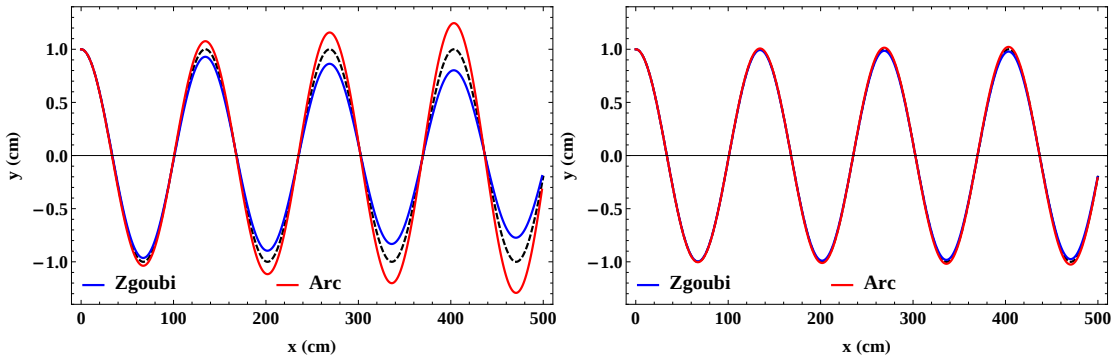

(a) 1 cm (b) 0.1 cm

FIGURE 3.4: A 10 MeV proton is tracked through a 5 m long quadrupole field (of gradient 10 Tm^{-1}) using the Zgoubi (truncated at second order) and arc integrators. The particle initially has a 1 cm horizontal offset from the quadrupole axis, and it's motion is directed along the length of the magnet. There is no initial vertical offset, meaning motion is limited to the horizontal plane. In Fig. (a) a 1 cm integration step length is applied, the non-symplecticity of both integrators is seen as the betatron oscillation amplitude dependence on longitudinal position (the correct solution is marked by the black dashed line). Reducing the step length (to 0.1 cm in Fig. (b)) leads to tracking results that are a better approximation of the true particle motion.

be seen for both methods as a change of betatron oscillation amplitude with longitudinal position. Reducing the integration step length to 0.1 cm (Fig. 3.4b) improves the accuracy of the solution, however the tracking results still visibly deviate from the correct solution before the end of the simulation. The appropriate integration step length for a particular tracking study is found on an ad hoc basis: the lattice configuration, the length being tracked through and the required degree of accuracy are all influencing factors. The default setting in Zgoubi is to track particles with Eq. 3.13 using terms up to fourth order (second to sixth order are available as an option). As expected, the inclusion of higher order terms gives greater accuracy for any given integration step length. Figure 3.5 shows the convergence of a tracking simulation to a solution for the different orders of the Zgoubi integrator and the arc method, it can be seen that the arc method converges at the same rate as the Zgoubi second order integrator.

3.5 Conclusions from tracking studies

We have investigated a number of methods for tracking particles through electromagnetic fields; with each method based on either forming transfer maps or on numerical integration. We found that, for the test scenarios considered, there was good agreement between the two types of tracking method, provided that the approximations made when tracking were valid for the given test configuration: for the transfer map based tracking, this meant having a map which described the motion of the particles to a sufficiently high order in the dynamical variables, whilst for numerical integration we had to choose an integration step length that makes a good compromise between the accuracy of tracking and the CPU time taken for the simulation. An advantage of the transfer map based

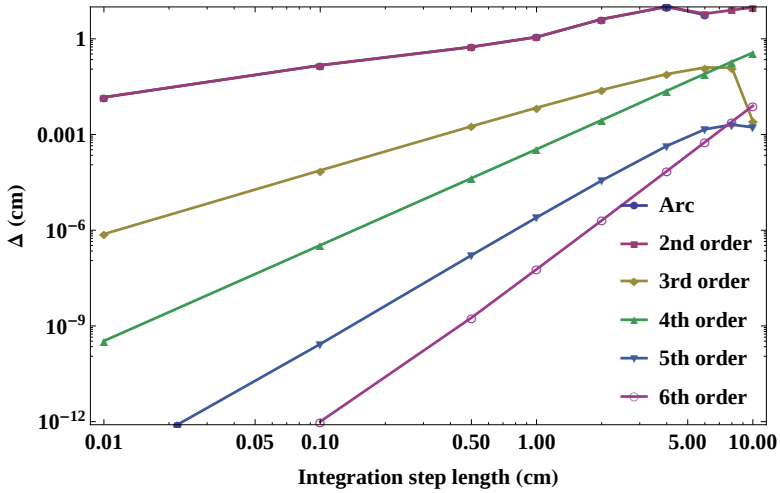

FIGURE 3.5: Convergence of tracking results to a solution for the different orders of Zgoubi integrator and arc method. For this example, a 10 MeV proton is tracked through a multipole field of 5 m in length. The field consists of quadrupole ($10\,\mathrm{Tm}^{-1}$), sextupole ($500\,\mathrm{Tm}^{-2}$) and octupole ($5000\,\mathrm{Tm}^{-3}$) components. Again, the particle initially has a 1 cm horizontal offset from the multipole axis, and its motion is directed along the length of the magnet. Δ is the difference between the particle position at the exit face of the magnet using a given integration regime and that found using the 6^{th} order Zgoubi integrator and 0.01 cm integration step length. The results for the second order Zgoubi integrator and arc method overlap.

methods is that the coefficients within the maps provide general information about the motion of particles at a given point within the lattice (rather than just of a single particle), however, creating maps for new field profiles is not necessarily straightforward. Results obtained using numerical integration methods can be used to create transfer maps, and it can be relatively straightforward to incorporate new field profiles into numerical integration codes; COSY and Zgoubi are examples of such codes. An advantage of COSY is that it can produce transfer maps of an arbitrary order, and then use these maps to track particles over many turns, potentially using less CPU time than using numerical integration alone. Advantages of using Zgoubi include the ability to describe a large variety of field types analytically and a number of options for combining and tracking through field maps. Additionally, the Zgoubi code had been used locally in the design of the EMMA FFAG, meaning that there was ready access to experienced users of the code.

A drawback to using any of the pre-existing tracking codes is that it can be difficult to fully understand all of the methods used. We addressed this concern (in part) by benchmarking the Zgoubi tracking code against the arc method of particle tracking (for which the tracking methods are understood). In the next chapter, we will compare the results of tracking particles through the ns-FFAG EMMA using Zgoubi, to results obtained experimentally.

Chapter 4

Verification of design studies

In investigating the relative merits of synchrotrons and FFAGs as machines for delivering proton therapy, tracking simulations based upon hard edge models of accelerator designs have been carried out. It is important to understand how the particle dynamics predicted by a hard edge model may differ from the particle dynamics observed in a real machine. The Electron Machine with Many Applications (EMMA) linear non-scaling FFAG is considered as a test lattice. The EMMA ring consists of 42 DOFO cells located around a 16 m circumference. The high density of magnets, combined with the short lengths and large apertures of the quadrupoles makes fringe field effects potentially significant. The situation is made more complex by the fact that the fringe fields can overlap in the short distances between magnets. These features make the EMMA lattice a rigorous testbed for hard edge magnet modelling.

The main aims of this chapter are:

1. to compare models based on different descriptions of the magnets (hard edge and field map), and to evaluate the extent to which the models are consistent with experiment;

2. to determine, based on the results of the comparison of the models with experimental data, whether the simpler hard edge magnet model is sufficient for the early stages of design of an FFAG, and when it is necessary to use a model that includes fringe fields.

4.1 The EMMA lattice

EMMA was the first (and is currently the only) non-scaling FFAG to have been built. The aims of the EMMA project were to show that a non-scaling FFAG could operate successfully, and identify whether such machines could be suitable for applications including proton therapy and muon acceleration. The lattice design for EMMA is based on the truncation of the scaling FFAG law after the two linear terms, so that the vertical field varies radially as:

$$B_z = B_{z,0}\left(1 + \frac{k\Delta R}{R_0}\right), \tag{4.1}$$

TABLE 4.1: EMMA DOFO cell parameters. These are the nominal values and include a 2 mm shift outwards (positive change) of both quadrupole offsets with respect to the original EMMA design parameters. The integrated field gradient is given by the quadrupole gradient multiplied by the quadrupole effective length.

Configuration	D/F doublet
Number of cells	42
D quad effective length	75.698 mm
D quad offset from reference axis	36.048 mm
D quad integrated field gradient	0.367 T
Short drift length	50.000 mm
F quad effective length	58.782 mm
F quad offset from reference axis	9.514 mm
F quad integrated field gradient	-0.402 T
Long drift length	210.000 mm

where $B_{z,0}$ gives the dipole field strength, k is the scaling index and R_0 is a reference radius. The resulting absence of higher order field components (e.g. sextupole) to correct chromaticity leads to the description of the machine as a linear non-scaling FFAG. The basis of the EMMA lattice is 42 DOFO cells (table 4.1), with the dipole fields that guide particles round the ring provided by introducing offsets to the transverse horizontal position of the quadrupoles that make up the cells.

The linear field design of the EMMA accelerator allows for a potentially large acceptance (meaning that a high emittance bunch may be successfully injected onto a stable orbit), however there is also a large betatron tune range encountered during acceleration from 10 MeV/c to 20 MeV/c, and the beam will cross a number of integer tune values. If a beam crosses integer tune values quickly (i.e. through rapid acceleration of the beam), then the effects of field errors on the beam only add coherently for a relatively small number of turns, and the beam may be accelerated to the extraction momentum successfully [59]. Short acceleration cycles in a machine such as EMMA are also of interest when considering the acceleration of beams consisting of short lifetime particles such as muons.

Rapid acceleration from the injection momentum to the extraction momentum in EMMA is achieved by applying a novel acceleration method [60]. For reference, we first describe the acceleration of low energy protons within a synchrotron accelerator. The particles that make up the beam are injected into the synchrotron with some range of momenta that are distributed around a design injection momentum. A particle of the design momentum will arrive at the first rf cavity when the rf field is at a design phase. However, dispersion of the lattice (which affects the path length of a particle per turn in the accelerator) and a difference in particle speed means that a particle with some offset from the design momentum will arrive at the rf cavity either earlier or later than a design particle, and subsequently at a different phase of the rf field. Meanwhile, the frequency of the rf changes to ensure that a particle with the design momentum (referred to as a synchronous particle) arrives at the design phase (referred to as the synchronous

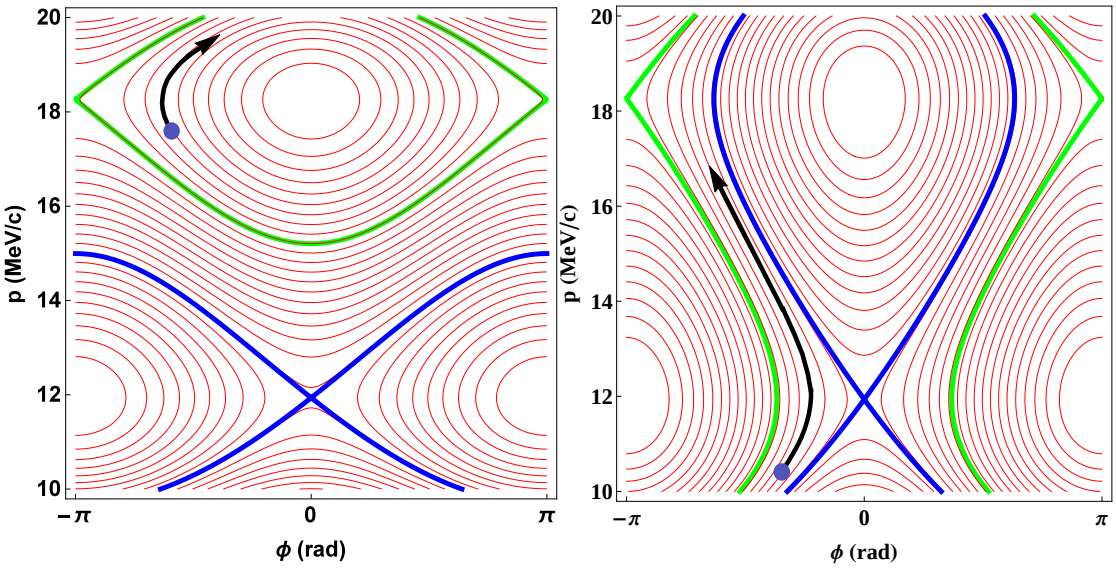

(a) 0.5 MV per turn.

(b) 1.5 MV per turn.

FIGURE 4.1: During acceleration, it is usually desirable to constrain particles within a stable area of longitudinal phase space, which is referred to as an rf bucket. The rf bucket is centred around a stable fixed point that corresponds to a particle having an orbital frequency that is a harmonic of the rf frequency. For EMMA, we see a low and high momentum rf bucket due to a parabolic orbital period with momentum, with the separatrices marked blue and green in Fig. 4.1a (a sample phase space trajectory is shown by a blue spot and black arrow). Given a sufficiently high accelerating voltage, the separatrices from the two rf buckets will cross and the serpentine channel opens (Fig. (b)); in this case a particle will travel around both stable fixed points.

phase) of the rf every time the cavity is encountered. The synchronous phase is selected so that off-momentum particles are focused towards the synchronous momentum during acceleration, as a result particles can be confined to a region in longitudinal phase space that is referred to as an rf bucket (demonstrated in Fig. 4.1a). The size of the rf bucket is determined by the rf voltage and by the rate at which the rf phase changes per turn for off-momentum particles (the phase slip factor).

In EMMA, electrons have ultra-relativistic momenta, meaning that the change in particle velocity during acceleration is small. The key contribution to the change of orbital period with particle momentum is therefore the change in path length with momentum. In a radial sector FFAG, particles move to greater radii orbit with increasing momentum, but if the FFAG is also non-scaling, then the shape of particle orbit will change with momentum too. The change in orbit shape with momentum has allowed for the EMMA lattice to be optimised so that a plot of orbital period vs. momentum is parabolic with a small change in the period over the EMMA momentum range. Having a design that is near isochronous during acceleration reduces the phase slippage per turn of a particle with respect to the rf, which, for a given rf voltage, leads to an increased rf bucket size. A longitudinal phase space portrait is shown for the EMMA lattice in Fig. 4.1a for an accelerating voltage of 0.5 MV (summed over the 19 cavities within the lattice) and an rf frequency of 1.301 GHz. There are two sets of separatrices marked on

the portrait (green and blue lines), which show the boundaries of two separate rf buckets. The red lines demonstrate contours of constant values of the longitudinal Hamiltonian, and we see that inside an rf bucket the contours form closed loops around a stable fixed point. A stable fixed point occurs when the rf frequency is some harmonic of the beam orbit frequency and the beam arrives at the cavity at an rf phase where the beam will see no voltage. In Fig. 4.1a we see stable fixed points at two different momenta, this is due to the parabolic orbital period, which leads to the rf frequency being at the 72^{nd} harmonic of the orbital period for momenta on either side of the minimum of the parabola. An increase in the accelerating voltage leads to an increase in the size of the rf bucket around each stable fixed point; eventually the separatrix from the low momentum rf bucket crosses the separatrix from the high momentum rf bucket (Fig. 4.1b). In the region of overlap of the two rf buckets we see a new type of behaviour, with a particle in this region following a path in longitudinal phase space that takes the particle around both the low and high momentum stable fixed points. Acceleration along this phase space path is known as serpentine acceleration, and allows for acceleration of a beam that does not have an exactly isochronous orbital period through a large momentum range given a fixed rf frequency. In order for serpentine acceleration to be implemented successfully, it is necessary for the EMMA lattice to have a parabolic orbital period close to that obtained during optimisation of the design.

In addition to the quadrupole magnets, other components for controlling the beam include 17 vertical correctors and 19 rf cavities. Injection is achieved using a septum and two kicker magnets; the beam can be extracted using a similar set of components. Beam diagnostics tools within the EMMA ring include 82 beam position monitors (BPMs).

4.2 Defining EMMA models

In making the comparison between hard edge and field map models, a simplified representation of the EMMA lattice is considered. The simplified model consists only of the DOFO cell components (i.e. quadrupoles and drift lengths), offering a representation of the ideal scenario for which the additional magnets (e.g. kickers) do not affect the multi-turn dynamics of the particle beam. This ideal lattice is shown in Fig. 4.2.

4.2.1 Hard edge model

The EMMA lattice is a 42 sided regular polygon, with each side (formed by a section of the reference trajectory) consisting of a DOFO cell. Each cell begins after a $\pi/21$ change in the reference axis following the previous cell. Simulations have been carried out for the hard edge model using the Zgoubi tracking code. In this instance it was convenient to use a periodic cell beginning at the entrance face of the defocusing quadrupole, and to use the same reference axis that is defined in the EMMA design drawings as a reference axis [61]. The initial hard edge model is based upon the parameters in table 4.1; other models, based upon variations of these parameters are considered at a later stage.

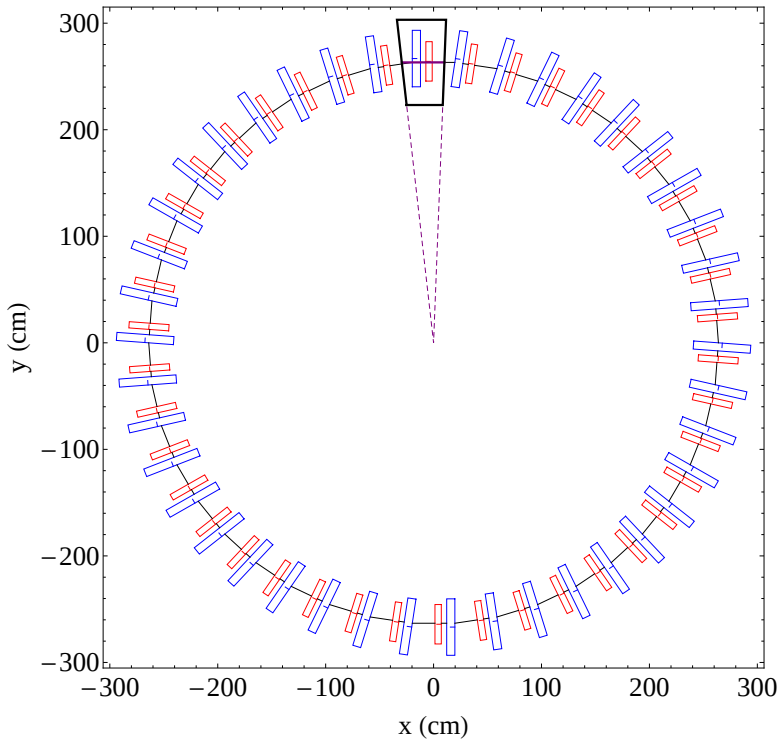

FIGURE 4.2: A basic representation of the EMMA lattice, consisting of defocusing (blue) and focusing (red) quadrupoles and drift sections only. The coloured lines through the axes of the quadrupoles show the offset of the quadrupoles with respect to the centre of the beam pipe (black line). The black box marks the periodic cell that is used for field map tracking.

4.2.2 Field map model

A field map for a single period of EMMA was produced using the Opera 3d design suite based on engineering designs for the magnets [62]. The records of amendments made to the magnets during the manufacturing process were found to be incomplete; however, the magnet descriptions used within the model are believed to be close enough to the manufactured magnets to give a fair comparison between field map simulations and experimental measurements. A full description of the magnet models (including yoke geometries and steel properties) is given in appendix B. The field map describes the magnetic fields on a 2D rectangular Cartesian grid that is located in the horizontal midplane (Fig. 4.3). In the following paragraphs the main features of the Opera model are discussed, including:

- the modification of the periodic cell definition;
- the enforcement of boundary conditions;
- the setting of current densities within conductors.

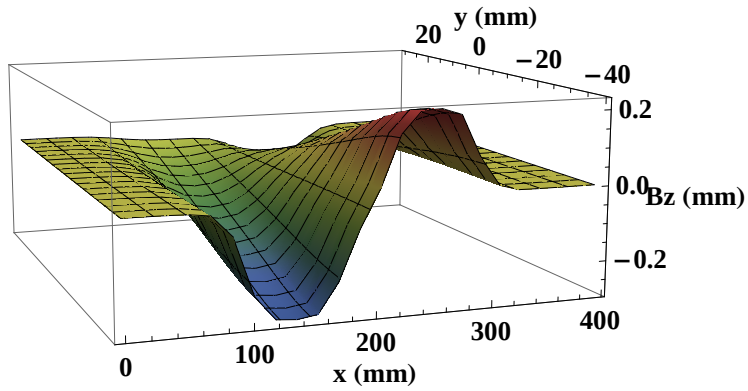

FIGURE 4.3: Plot to show the calculated midplane field map in one periodic cell of the EMMA lattice. The symmetry of an ideal period means that the magnetic field will cross the midplane normal to the midplane (i.e. $B_x = B_y = 0$).

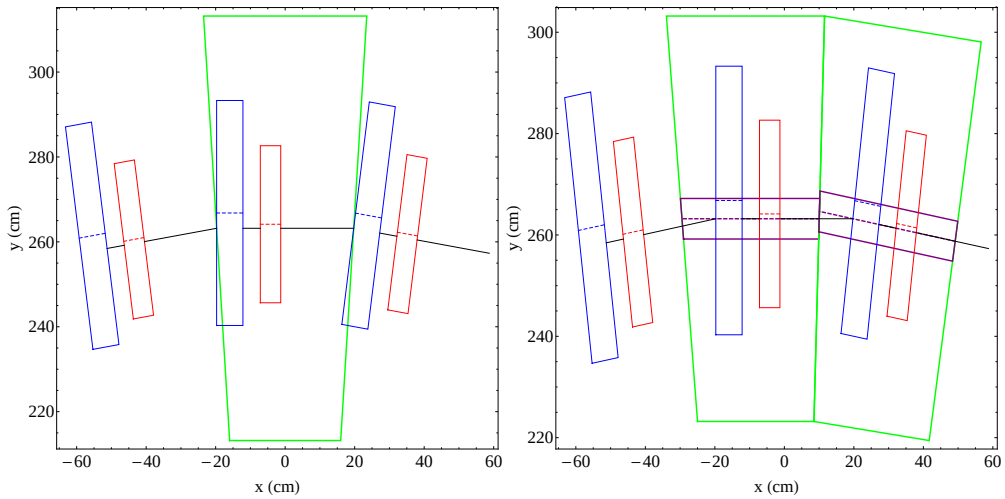

(a) Hard edge period.

(b) Field map period.

FIGURE 4.4: If the DOFO cell is taken as the design period of EMMA, then the boundaries of the periodic cell (marked green) will intersect the steel of the D quadrupole (Fig. (a)). For the purpose of modelling the cell within Opera a different periodic cell has been defined, which places the quadrupoles at the centre of the periodic cell, and minimizes the magnetic fields at the period cell boundaries; two of these periods are marked in Fig. (b). For tracking simulations, the field map is ideally defined for the areas enclosed by the solid purple lines (a rectangular map is used for compatibility with Zgoubi), and the reference axis is shown by the purple dashed line. It can be seen that, for the new period, the reference axis no longer runs continuously over the boundaries between periodic cells.

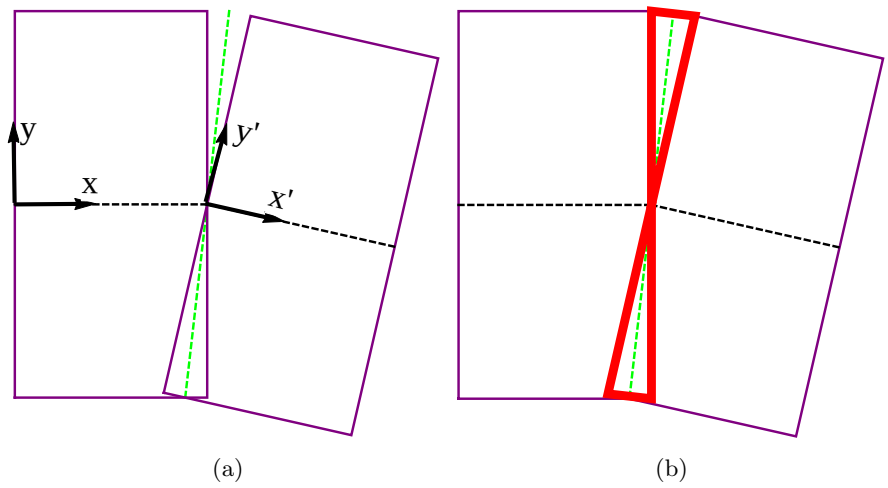

(a) (b)

FIGURE 4.5: If non-zero fields at the boundaries of a periodic cell are not taken into account, then use of a rectangular field map can lead to tracking errors when the transformation of the reference frame is applied (Fig. (a)). Field map boundaries are marked in purple, the reference axis is black, and the periodic boundary is shown in green. Potentially problematic regions are within the red triangles (Fig. (b)).

Modification of the periodic cell

In order to simplify the field calculation, it was helpful to consider a different period for the lattice than that used for the hard edge model. This is because the boundary of a single periodic cell would intersect the steel of the defocusing quadrupole within the Opera model (Fig. 4.4a). Instead, a periodic cell is chosen in which the D and F quadrupoles are located at the centre of the period (Fig. 4.4b), which ensures that the magnetic fields at the boundaries of the period are small. Apart from simplifying the modelling of the cell in Opera, this new period has an advantage when tracking with Zgoubi: for simplicity, the reference axis transformation made at the end of each period assumes that the particle is being transported through a field free region (Fig. 3.3). Choosing the start of the periodic cell at the entrance of the defocusing quadrupole would lead to an error from the non-negligible fields at this point. However, if the start of the periodic cell is near to the centre of the long drift, then the reference axis transformation can be applied where the fields are small.

Figure 4.5 illustrates a further aspect of the problem: the green dashed line shows the location of the boundary of a periodic cell. The field maps that we track through with Zgoubi are in the form of a rectangular Cartesian grid; it can be seen that the field map extends beyond the period boundary on one side of the reference axis and falls short of the period boundary on the other side.

For a particle travelling on the outer side of the reference axis, the field map starts after and ends before the periodic boundary: information about the fields close to the boundary has been lost. Upon transforming the particle to a new reference frame, it is assumed that the particle will travel through a region of zero field drift space (upper red triangle in Fig. 4.5b), though this may not be accurate. For a particle travelling

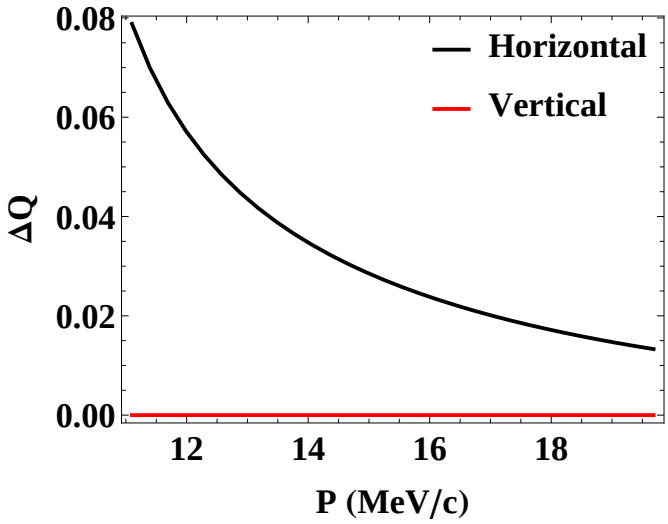

FIGURE 4.6: Change in lattice tune vs. momentum. The new period definition, with the magnets located away from the period boundary, leads to a small change in the horizontal tune. No differences were measured for the vertical tune or orbital period.

on the inner side of the reference axis, the field map starts before and ends after the periodic boundary. If, in these overlapping regions, the field map contains field data from the adjacent cells, then the particle will be tracked through the fields of the lower red triangle in Fig. 4.5b twice when the change in reference frame is made. Ensuring that the magnetic fields are negligible at the periodic boundaries means that the transformations already available within Zgoubi (Eq. 3.19) can be used without introducing significant errors to the beam dynamics.

As the newly defined period maintains a reference axis that is parallel with the axes of the quadrupoles, additional transformations must be made within Zgoubi at the end of the period. This is illustrated in Fig. 4.4b, where it can be seen that a particle with zero transverse horizontal offset from the reference axis at the end of the first period should have a non-zero offset at the start of the second period (given that the reference axis is not continuous across the boundary). The transformations used are $\Delta y = 1.45095\,\text{cm}$ and $\Delta x = 0.0553603\,\text{cm}$. To check that the new definition of the periodic cell leads to the same dynamics as the original periodic cell, both were modelled using a hard edge representation of the magnets; a small change in the horizontal ring tune (up to $\approx 0.6\%$, Fig. 4.6) was observed, however, the vertical tune and orbital period remained constant. The change in the horizontal tune needs further investigation to be understood.

Boundary conditions

The iron in an EMMA cell has midplane (xy) symmetry, and the magnetic field should cross this plane of symmetry normal to the midplane. When solving the finite element mesh, setting the boundary condition that the field is normal to the midplane means that a field map can be calculated after solving the finite element mesh for either the top or bottom half of the cell only. Reducing the volume of the mesh in this fashion

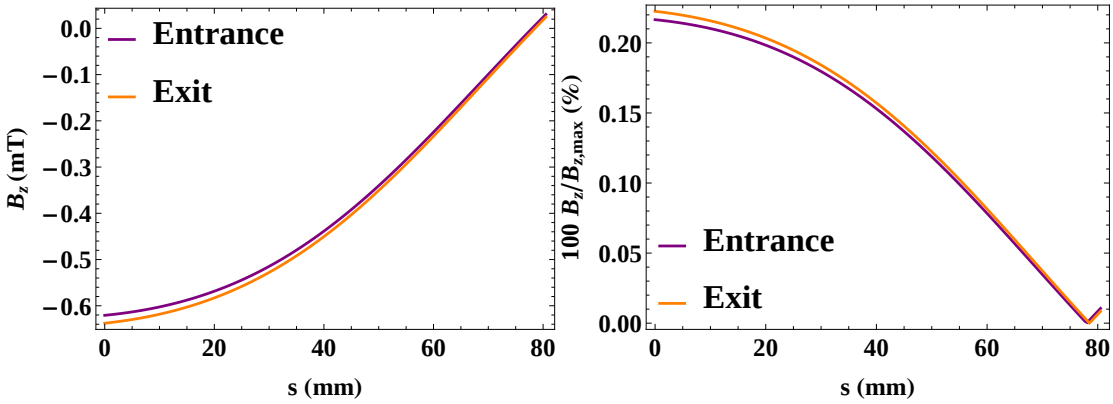

(a) (b)

FIGURE 4.7: An Opera model that takes into account the periodicity of the EMMA lattice (by including three consecutive cells) shows that the magnets in one period will contribute to fields in the adjacent cells. We take the central period of the three periods, and then find the magnetic field along the entrance and exit boundary (as seen by the beam) on the midplane of this period. Fig (a) is a plot of the boundary magnetic field vs. the distance along the boundary, s, with s increasing from a point on the boundary that is closer to the machine centre than the low momentum closed orbit position to just beyond high momentum closed orbit. The difference between the plot for the entrance and exit boundary indicates a small error in the FEM solution; this may be solved using a finer mesh. Fig. (b), which shows the fields from Fig. (a) as a percentage of the maximum field found within the field map, tells us that the fields along the boundaries (including the difference between the entrance and exit boundary) are relatively small.

allows for a significant reduction in the time taken for a solution to be found.
The boundary conditions at the remaining mesh volume boundaries have been set so that the magnetic fields are tangential to the boundary (implying that no magnetic flux leaves the volume of the model). In reality, the fields around a magnet extend out infinitely in free space. The periodicity of the EMMA lattice and the distance of the radial and uppermost boundaries from the quadrupole iron have been considered in defining the optimal volume for which the finite element mesh will be solved (balancing the accuracy of the solution with the time taken to solve the mesh).

The periodicity of the lattice was thought to be potentially of concern for EMMA modelling due to the close proximity of the quadrupoles to the boundaries, the effect of which could be that the fields within one period could contain a significant contribution from the magnets in the adjacent periods. To investigate this an Opera model containing three consecutive EMMA periods was produced and a field map calculated for the central period. The fields obtained at the boundaries of this single period are shown in Fig. 4.7. The differences in the beam dynamics that are obtained by tracking through two field maps, one where the periodicity condition has been enforced and one where it has not, were found to be small (shown in Fig. 4.8 and Fig. 4.9). As a result, field maps calculated after modelling a single cell have been used for the remainder of this chapter.

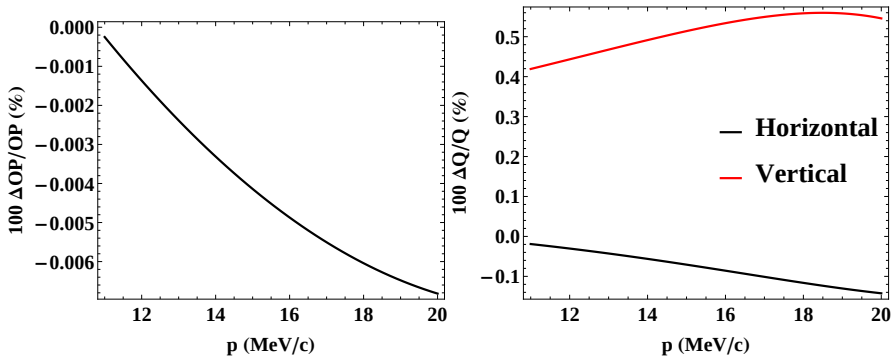

(a) Orbital period (b) Betatron tune

FIGURE 4.8: Making the approximation that a periodic cell is not influenced by the adjacent cells leads to small changes in the beam dynamics. Figures (a) and (b) show the percentage change in orbital period and tune when an Opera model based on three periodic cells is used instead of a model with one periodic cell.

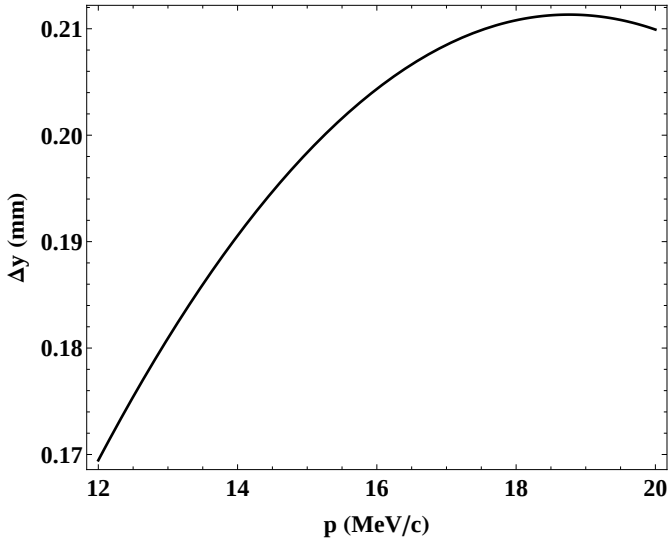

FIGURE 4.9: Change in the closed orbit position (Δy) at a single reference point within a cell found when the approximation that a periodic cell is not influenced by the adjacent cells is made.

Quadrupole currents

The magnetic field gradient integrated along the length of each individual EMMA quadrupole magnet has been measured by Tesla Engineering using a rotating coil of 35 mm radius. Each quadrupole has a field clamp attached to the side of the magnet that will face the long drift that separates neighbouring EMMA cells (Fig. 4.10). The field clamp has a high magnetic permeability relative to that of air, meaning that the flux density within the clamp is much higher than it would be for a corresponding volume of air at the same location. The quadrupole fringe field contribution to the magnetic field in the long drift, between the clamps for the focusing and defocusing magnets, is significantly reduced by the presence of the clamps. During the measurements of integrated field gradient, the field clamp position of each quadrupole was adjusted in an attempt

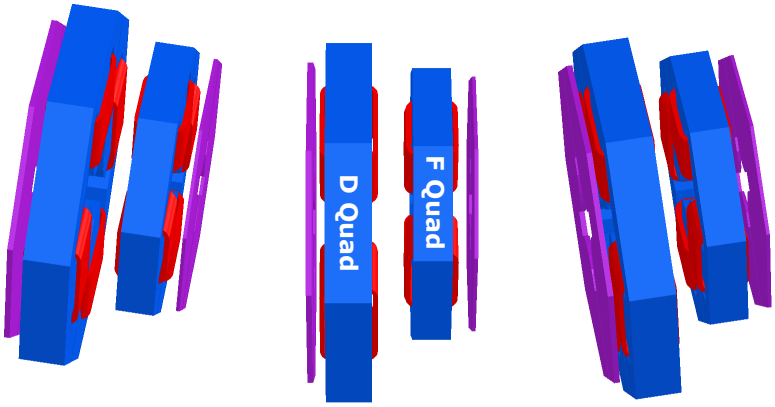

(a) Three cell top view

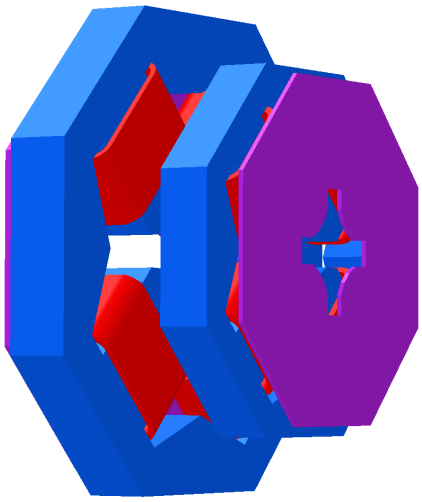

(b) Single cell view

FIGURE 4.10: Images to show the positions of the field clamps (purple) relative to the EMMA quadrupoles (blue). Conductors are marked in red. Figure (a) shows three consecutive cells, through which the beam would travel from left to right. The positioning of the field clamps, after the focusing quadrupole and before the defocusing quadrupole (in the direction of the beams travel), significantly reduces the magnetic field in the long drifts that separate cells. A single cell is shown from a different perspective for reference (Fig. (b)).

to reduce the variation in measurements between different magnets of the same type (i.e. focusing or defocusing) for a specified coil current.

The positions of the field clamps following the calibration procedure can be found in appendix B. With these new field clamp positions, the mean integrated field gradient was measured to be (0.55078 ± 0.00004) T for the focusing quadrupoles (coil current of 364 A) and (0.47240 ± 0.00003) T for the defocusing quadrupoles (coil current of 376 A).

Figure 4.11 shows that the mean measured integrated field gradients (labelled 'F calibration' and 'D calibration') are approximately 98% and 98.6% of the integrated field gradients calculated with Opera for the focusing and defocusing quadrupole respectively. It is assumed that these differences arise from errors in the conductor volume specified

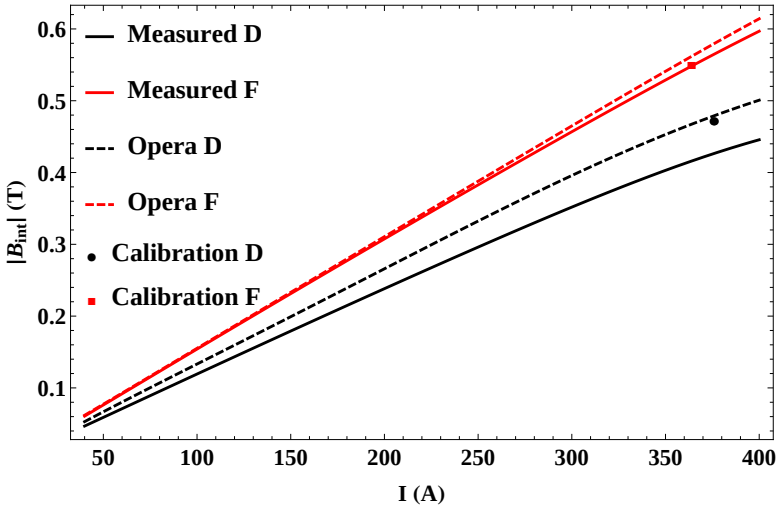

FIGURE 4.11: Integrated field gradient vs. quadrupole current. The dashed lines represent the excitation curves calculated with Opera. The solid lines and individual points were measured on separate occasions by Tesla Engineering using a rotating coil. The difference between the Opera curve and calibration points is assumed to be due to the approximate description of the conductor within the Opera model; before carrying out tracking studies, we scale the current used within the Opera model so that the integrated field gradient for the model matches that of the calibration points. The measured excitation curve (for which we assume a scaling error) is used as a reference for the nonlinear relationship of the integrated field gradient with quadrupole current.

within the Opera model; this has been compensated by scaling the quadrupole currents in the model so that the calculated integrated field gradient agrees with the mean measured integrated field gradient. The scaling factors are 0.97959 for the focusing quadrupole and 0.98173 for the defocusing quadrupole.

Tesla Engineering have measured excitation curves for the EMMA quadrupoles [1]. These measurements show a non-linear relationship between magnetic field and coil current (particularly for the defocusing magnet) due to the saturation of the quadrupole steel: this is considered later when we look at the experimental data measured using equivalent momenta (section 4.4.2). The field map produced at this stage is based upon the baseline EMMA quadrupole current settings, which were found following the optimisation of EMMA experimentally. These settings are currents of 303.8 A and 270.5 A for the defocusing and focusing quadrupoles respectively.

4.3 Comparing hard edge and field map models

Initially a hard edge model based upon the lattice definition given in table 4.1 is compared to a model using a field map calculated for the experimental baseline quadrupole currents. Qualitatively, the basic features shown in Fig. 4.12 are as expected, with the

[1] These were also measured using a rotating coil. It can be seen in Fig. 4.11 that the integrated field gradient measured during calibration does not sit on the measured excitation curve. Tesla Engineering could not offer a definitive explanation for this. It is assumed that whilst the excitation measurement is subject to a systematic error, the shape of the curve is accurate.

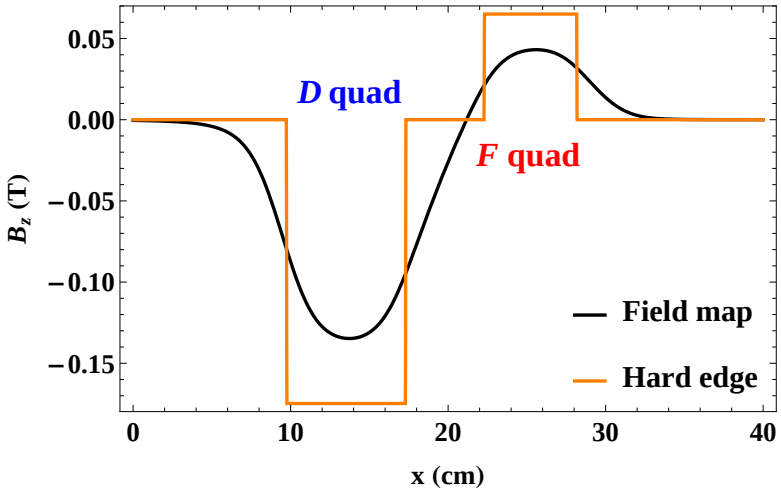

FIGURE 4.12: Vertical field component, B_z, vs. longitudinal position along the reference axis. The hard edge model is based upon the component lengths and integrated field gradients given in table 4.1. The field for the field map model uses the baseline experimental quadrupole currents. The instantaneous changes in B_z along the x axis seen for the hard edge model mark the boundaries of quadrupoles.

fields of the field map model extending beyond the quadrupole yoke. Quantitatively, a difference is found between the integrated field gradients measured along the reference axes of the field map and hard edge models. For the field map, we calculate the integrated field gradient for the defocusing quadrupole by integrating the field along the reference axis from the start of the cell until the polarity of the field changes, we then divide this integrated gradient by the offset of the defocusing quadrupole from the reference axis; to find the integrated field gradient for the focusing quadrupole, we instead integrate the field from the point at which the field polarity changes until the end of the cell, and use the offset of the focusing quadrupole from the reference axis. The integrated field gradients calculated from the field map for the defocusing and focusing quadrupoles are found to be 0.340 T and -0.309 T respectively. The shapes of the fields within the field map are complicated by the close proximity of the two magnet yokes and by the breaking of rotational symmetry around the reference axis because of the different transverse offsets of the defocusing and focusing quadrupoles from the reference axis. This is illustrated in Fig. 4.13 where it can be seen that the drift space that separates the two quadrupoles actually contains a combination of the fringe fields of both magnets.

One of the advantages of using a hard edge model is that it is straightforward to attribute random alignment errors to each magnet of the lattice within tracking studies. The same may be true for a field map model if the steel of one quadrupole does not influence the field of the second, and if the net field at a point within the map is simply a superposition of the fields from the individual magnets; if this is not the case, then it may be necessary to solve a finite element model for any combination of alignment errors used in a tracking study. Two field maps were produced to check whether these conditions are met within the EMMA cell. Both maps were calculated using an Opera

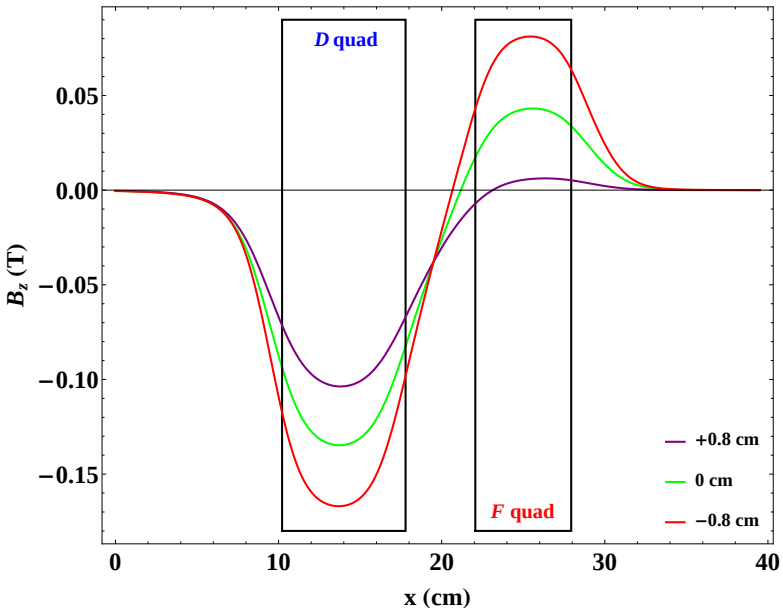

FIGURE 4.13: Vertical field component, B_z, vs. longitudinal period position along three straight reference lines. The 0 cm line represents the reference axis specified in the EMMA design, the +0.8 cm line shows the fields along an axis shifted towards the outside of the ring and the -0.8 cm line an axis shifted towards the inside. In this example, the intersection of the three plots within the drift space indicates the point at which the quadrupole term of the field becomes horizontally focusing rather than defocusing. The field being negative at this location means that the dipole term originating at the defocusing magnet is still dominant over the dipole term from the focusing magnet.

model that has the EMMA magnets at their design positions. For the first map the defocusing magnet is powered whilst the focusing magnet is not, and for the second map it is only the focusing magnet that is powered. In Fig. 4.14 it can be seen that the steel of the first quadrupole does affect the field profile of the second (and vice versa), and that the field map for a complete cell is not simply a superposition of the field from the individual magnets. Consequently it is expected that tracking through superposed fields will introduce errors to the beam dynamics. Figure 4.15 shows that, for the simple case where we do not introduce magnet position errors when superposing maps, the change in the particle dynamics as a result of superposing two maps is small compared to the likely errors in determining the dynamical variables experimentally. Although these differences are small, they suggest that further investigation is needed to determine how the errors change when we introduce magnet misalignments; for the remainder of this chapter we use field maps based on Opera models that have been solved with both magnets powered. Figures 4.16 to 4.18 show tracking results for the field map (with simultaneously powered quadrupoles) and hard edge models. In keeping with the significant differences between the two field profiles, there are noticeable differences in the particle dynamics. For example, in Fig. 4.16, it can be seen that the closed orbit positions for the hard edge model are shifted outwards with respect to the closed orbits of the field map model (by up to ≈ 0.6 cm), which in turn leads to the longer orbit

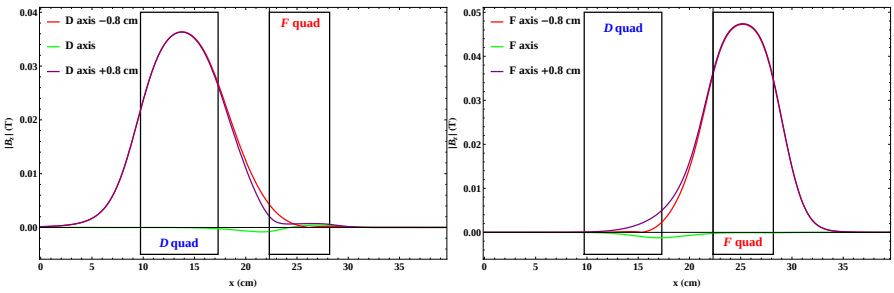

(a) D powered only

(b) F powered only

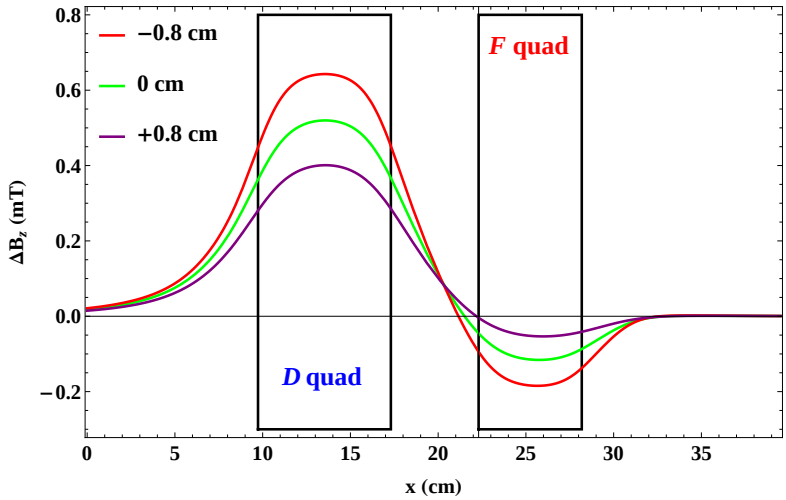

(c) Difference in magnetic field, ΔB, between the superposed maps of Figs. (a) and (b) and a map calculated with both quadrupoles powered.

FIGURE 4.14: Figures (a) and (b) show fields along lines parallel to the reference axis in a single EMMA cell in the situation where both quadrupoles are present within the Opera design, but only one is powered. For each case, the magnetic field is calculated along the magnetic axis of the powered quadrupole and then along lines ± 0.8 cm horizontally displaced parallel to the magnetic axis. It can be seen that the presence of the steel of the second quadrupole breaks the symmetry of the powered quadrupole (by the non-zero field along the magnetic axes and the differences between the magnitude of fields observed at $+0.8$ cm and -0.8 cm). Figure (c) shows the differences in the magnetic field found by subtracting the superposed fields of (a) and (b) from the field map calculated with a model where both quadrupoles are powered. The differences are believed to arise due to the non-linear saturation curve of the quadrupole steel.

periods for the hard edge model (Fig. 4.17). Despite this, some of the key features of the EMMA lattice are preserved, including the parabolic orbital period curve and the approximate number of integer tune values crossed during acceleration (Fig. 4.18).

One of the aims of the EMMA investigation presented in this thesis is to find the degree to which the dynamics obtained from a hard edge model are consistent with the real machine. At this stage, we limit the comparison to a comparison between the hard edge model and a field map model; in principle, the field map model provides a more realistic representation of the machine. We have already given some of the qualitative and quantitative differences between the hard edge and field map field profiles. Some

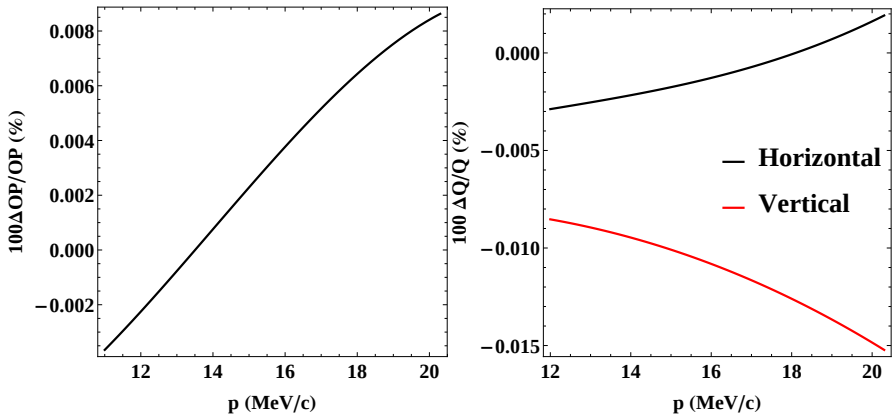

(a) Orbital period (b) Betatron tune

FIGURE 4.15: Difference in the orbital period (Fig. (a)) and betatron tune (Fig. (b)) found when tracking particles through a field map calculated by superposing the field maps for individually powered quadrupoles and when tracking particles through a field map calculated for both quadrupoles powered simultaneously.

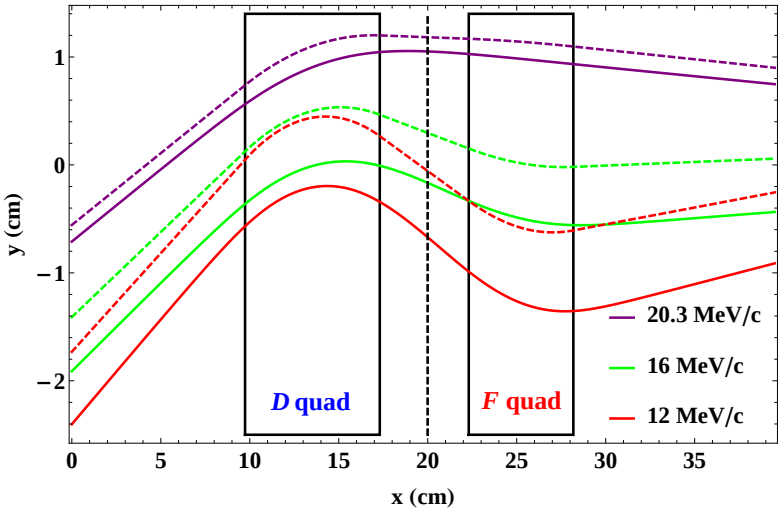

FIGURE 4.16: Simulated closed orbit positions along the length of an EMMA cell. Solid coloured lines are the closed orbits found for the field map model, whilst dashed coloured lines are as found for the hard edge model. The dashed black line marks the BPM position at which the closed orbits will be measured experimentally.

of these differences result from the field map being based on the machine settings used experimentally, which in turn were found after some optimisation of machine operation; the hard edge model has so far been based on design parameters. Given the time taken to calculate field maps, the degree to which a hard edge model may produce the same particle dynamics as a field map is sought by means of a parameter search. Assuming a machine that has already been constructed, the geometry of the machine within the model is fixed, excluding alignment errors, and is as described in table 4.1. The quadrupole strengths are left as free parameters within a search that has the objectives of reducing the differences between the betatron tune (horizontal and vertical), orbital period and closed orbit position (at the location of one of the BPMs) calculated across

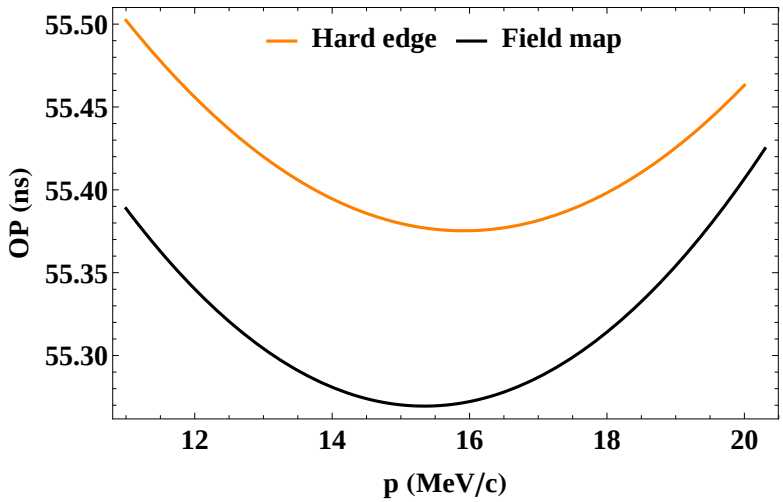

FIGURE 4.17: Orbital period (OP) vs. momentum for the hard edge model based on the design parameters for the magnetic fields (table 4.1) and for a field map with the quadrupole currents based on the EMMA experimental settings.

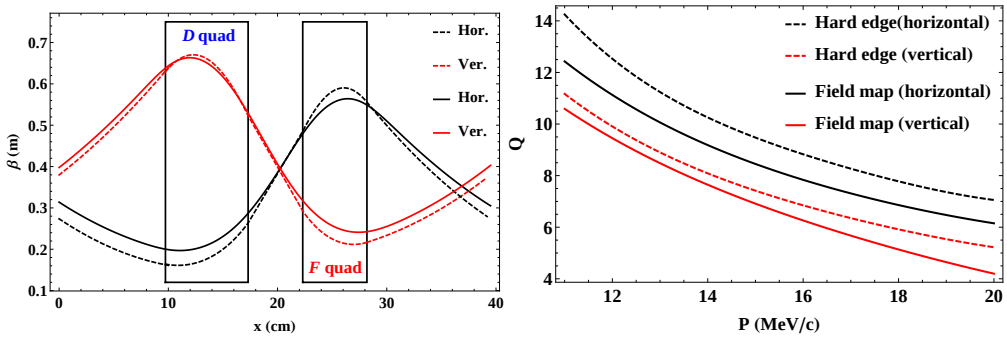

(a) β functions at 15 MeV/c (b) Betatron tune

FIGURE 4.18: Focusing properties of the hard edge (dashed lines) and field map (solid lines) model, with the β functions for 15 MeV/c shown as an example. Whilst there are differences between hard edge and field map tune values, the number of integer tune values crossed during acceleration is similar.

the full momentum range with each model. The results of this search can be seen in Fig. 4.19, where four regions of lower residuals can be identified (each representing one of the four objectives), implying that good agreement between hard edge and field map models for one of the objectives leads to worse agreement for the other objectives. This implication is shown more clearly in Fig. 4.20, where we can see how optimising for one of the metrics impacts on the agreement between models for the remaining metrics.

A hard edge model which gives particle dynamics that more closely match those of the real machine (or field map) can be found by varying the geometry of the cell within the model [63], however, for both the optimisation based on quadrupole strengths and cell geometry, such models can only be found retrospectively. For the EMMA accelerator, hard edge modelling plays an important role in the early stages of development, allowing for general features of dynamics of interest to be studied (for example, integer resonance crossing and serpentine acceleration). The extent of fringe field regions and the effects

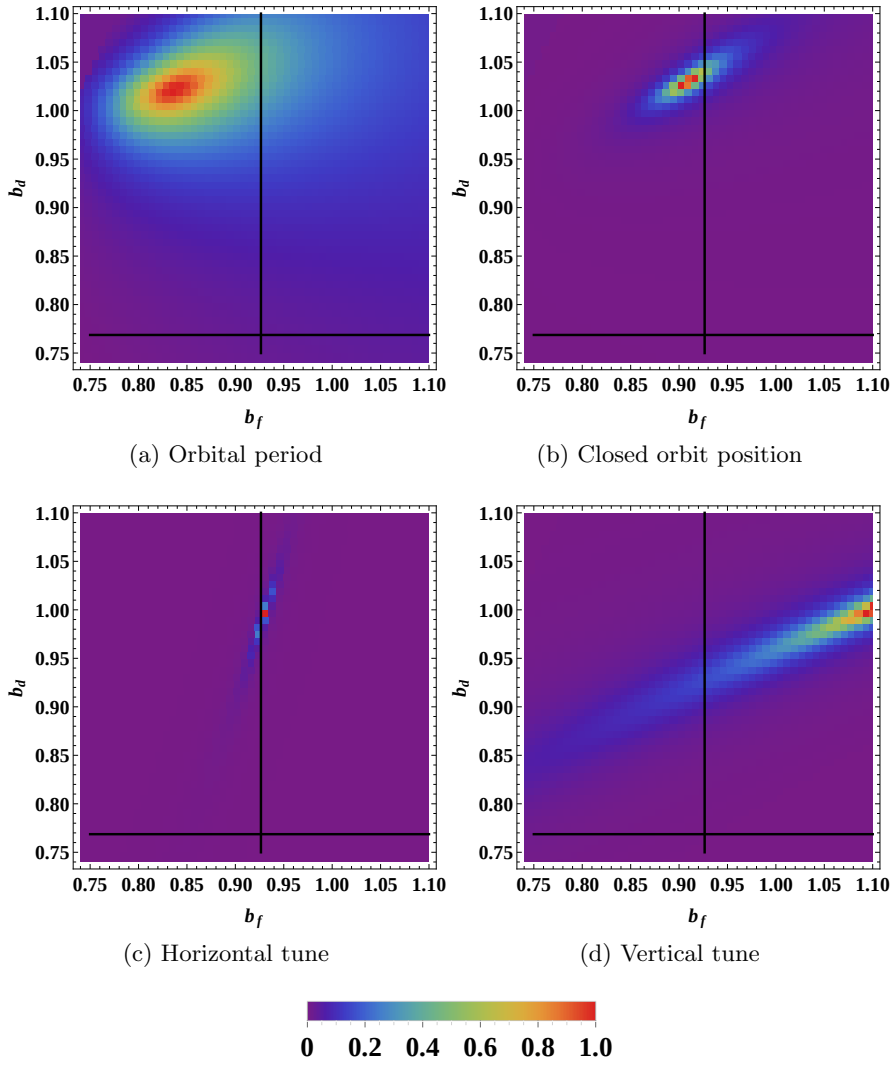

(a) Orbital period

(b) Closed orbit position

(c) Horizontal tune

(d) Vertical tune

FIGURE 4.19: Results from a simple parameter search aimed at finding the optimum hard edge representation of the field map. The cell set up is as described in table 4.1, but with the quadrupole integrated field gradients left as free parameters. The b_d and b_f axis represent a linear scaling coefficient of the design defocusing and focusing fields, respectively, whilst the black lines show the coefficients required to give the integrated field equal to that calculated for the field map. The optimisation is based on finding the scaling coefficients that reduce the differences in orbital period, betatron tune and closed orbit position between the hard edge and field map model. For each configuration, the residuals between the hard edge and field map model have been calculated; the reciprocals of the residuals for each metric have been scaled to give 1 for the best agreement between models (red end of the spectrum) and 0 for the worst agreement. The four red islands above represent the minimal residuals for each of the optimisation goals. The impact of optimising for one of the metrics upon the agreement between models for the remaining metrics is demonstrated in Fig. 4.20.

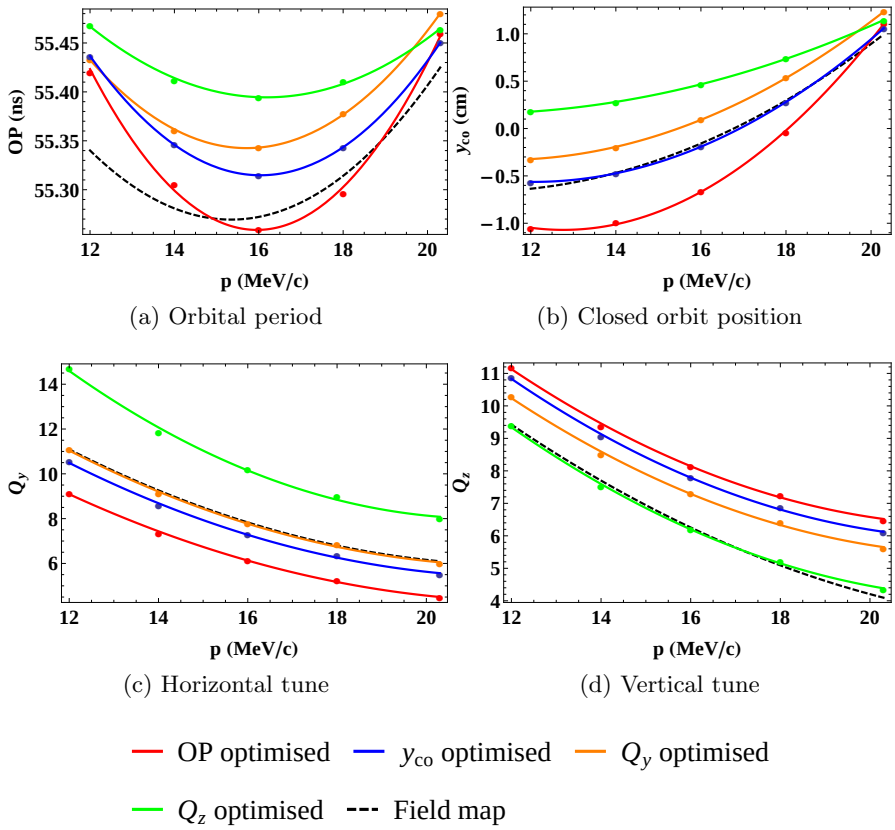

(a) Orbital period (b) Closed orbit position

(c) Horizontal tune (d) Vertical tune

FIGURE 4.20: Tracking results for the four sets of quadrupole strengths that minimise the residuals in Figs. 4.19a – 4.19d. The red curves, for example, show the tracking results for the hard edge model quadrupole strengths that minimise the residuals for orbital period data when comparing hard edge and field map models (Fig. 4.19a). We see that optimising the agreement between hard edge and field map for one of the metrics does not give optimal agreement for the remaining metrics.

on the field profile due to the close proximity of magnets and the breaking of rotational symmetry, mean that significant differences between the dynamics predicted by hard edge and field map models can be expected.

4.4 Principal experimental investigation

4.4.1 The ALICE/EMMA accelerator facility

Particle bunches are injected into EMMA from the energy recovery linac based ALICE (Accelerators and Light In Combined Experiments) accelerator (Fig. 4.21). For a more typical ALICE set up, trains of electron bunches (up to $\approx 100\,\mathrm{pC}$ bunch charge) are accelerated up to $30\,\mathrm{MeV/c}$ within a super conducting rf cavity [64]. These bunches are then steered through a chicane, that can be used to produce intense pulses of THz light, and then through the undulator magnets of a free electron laser (FEL) that can produce infra-red light. Following the FEL, bunches are steered round to the cavity that was

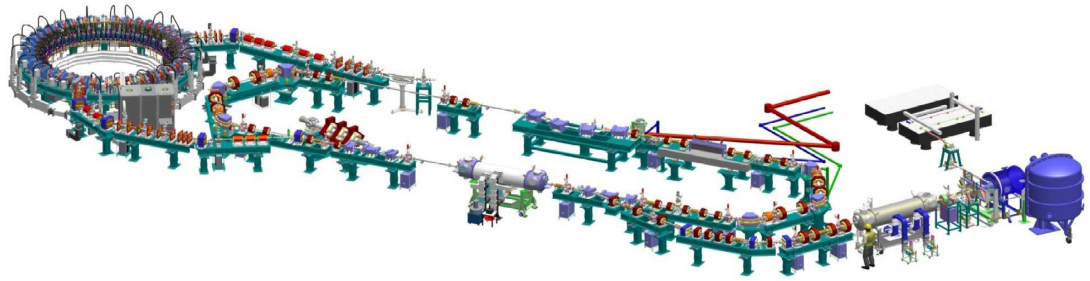

FIGURE 4.21: The ALICE/EMMA accelerator facility. ALICE acts as an injector when operating EMMA (the ring structure furthermost left in the image).

encountered earlier, where the energy that had previously been imparted to a bunch is now recovered.

For EMMA operation, ALICE accelerates single bunches ($\approx 40\,\mathrm{pC}$ bunch charge) of electrons up to 12.5 MeV/c. A dipole magnet, located shortly after the rf cavity, is used to steer particles out of ALICE and into the EMMA injection line. At the end of the injection line, a septum magnet guides particles into the EMMA ring, where two kicker magnets then push the particle bunch closer to the closed orbit position. The bunch injected from ALICE has an emittance that is much smaller than the EMMA design acceptance; to test the acceptance of the EMMA accelerator [65], the EMMA phase space is painted by varying the strength of a vertical corrector magnet in the injection line and of the injection septum.

4.4.2 Equivalent momentum

The dynamics of a particle in a quadrupole are determined by the normalised quadrupole strength, which is given by:

$$K_1 = \frac{1}{B\rho}\frac{\partial B_z}{\partial y}. \tag{4.2}$$

The dynamics of particles with different momenta may be studied by keeping the injection momentum constant and instead varying the quadrupole gradients. For the EMMA experimental work carried out for this thesis, the ALICE accelerator was set up to provide particles of 12.5 MeV/c; the quadrupole currents were then varied so that the quadrupole strengths experienced by a particle would be consistent with the strengths experienced by a particle of a different momentum as it travels through the nominally powered lattice. From here on, the momentum quoted in any experimental results refers to the equivalent momentum corresponding to the appropriate scaling of the quadrupole strengths. The quadrupole currents used for the experiments in EMMA were based on modelling and optimisation carried out early in the commissioning process, and are calculated by:

$$I_{Q_d} = 245\frac{15.5}{p_{\mathrm{equiv}}}\frac{p_{\mathrm{real}}}{12.5}, \tag{4.3}$$

$$I_{Q_f} = \frac{I_{Q_d}}{1.122926}, \tag{4.4}$$

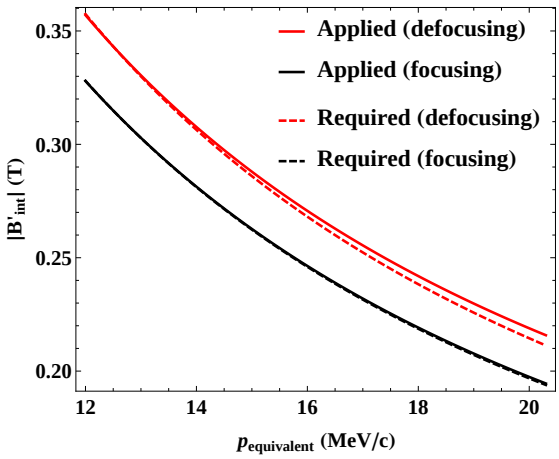

FIGURE 4.22: The integrated field gradients required to give a particular equivalent momentum (based on field map calculations), are compared with the integrated field gradients that were applied experimentally.

where $I_{\mathrm{Q_d}}$ and $I_{\mathrm{Q_f}}$ are the currents at the defocusing and focusing quadrupoles respectively, p_{real} is the real momentum of the particle and p_{equiv} is the equivalent momentum.

The accuracy of applying the excitation curves shown in Fig. 4.22 (solid lines) has been considered. Initially, quadrupole currents are calculated from Eq. 4.3 and Eq. 4.4, and the Opera model is solved for an equivalent and real momentum of 12.5 MeV/c; this is the baseline current setting (i.e. a 12.5 MeV/c particle has the dynamics expected of a 12.5 MeV/c particle). For the Opera field map, we use the method described at the beginning of section 4.3 to calculate the integrated field gradient of each quadrupole. The integrated field gradient should then scale according to Eq. 4.2 in order to give an equivalent momentum (shown by the dashed lines in Fig. 4.22). We then calculate five further maps, in these maps we set the quadrupole currents according to the experimental settings used for the equivalent momenta of 12, 14, 16, 18 and 20.3 MeV/c (but still with a real momentum of 12.5 MeV/c). The integrated field gradients are then calculated for each of the five maps (the field profile along the reference axis for these maps is shown in Fig. 4.23), with the results shown by the solid lines in Fig. 4.22. A small difference is seen in the plot of integrated field gradient vs. equivalent momentum between the lines produced by scaling the fields of the 12.5 MeV/c map and the lines given by the five equivalent momenta maps; which is potential evidence of an error in the coefficients of Eqs. 4.3 and 4.4. The difference in the required and applied integrated field gradients for any given equivalent momentum is largest for the defocusing quadrupole, and the effects of this difference can be seen when comparing the results of tracking simulations. The dynamics have been calculated by tracking particles of different momenta through the 12.5 MeV/c map and by tracking a particle of 12.5 MeV/c through the five maps corresponding to equivalent momenta. The differences in the dynamics are as expected from the excitation curves, with better agreement between the methods at lower momentum (Fig. 4.24).

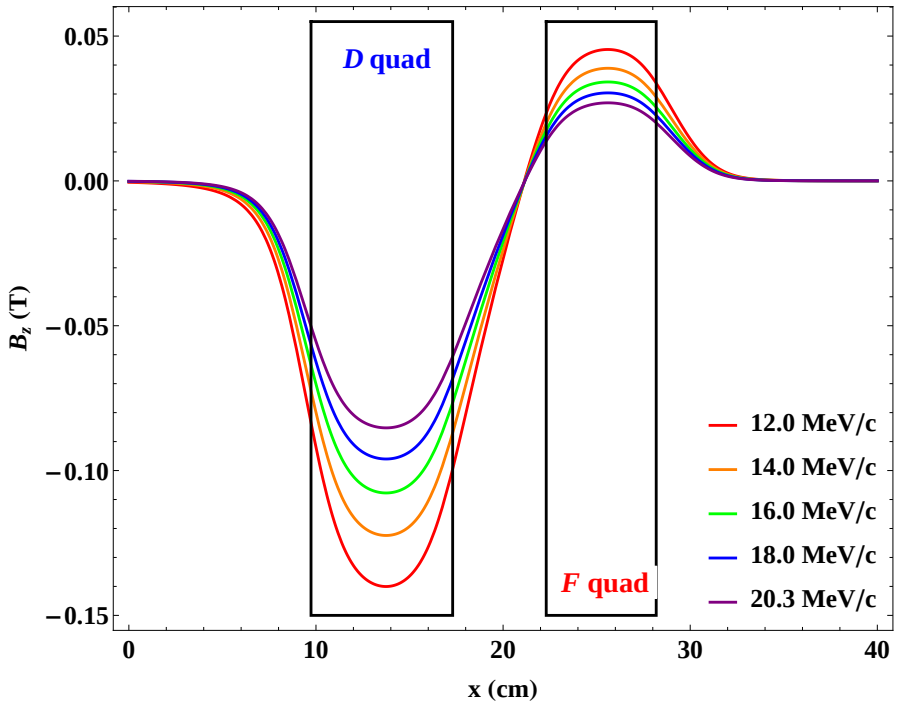

FIGURE 4.23: The experimental measurements of orbital period, betatron tune and closed orbit have been made using the method of equivalent momenta. Five field maps have been produced (with the quadrupole currents adjust according to Eq. 4.3) for the purpose of making a comparison between these experimental measurements and field map tracking. The above plot shows the magnetic field along the cell reference axis for the five equivalent momenta used.

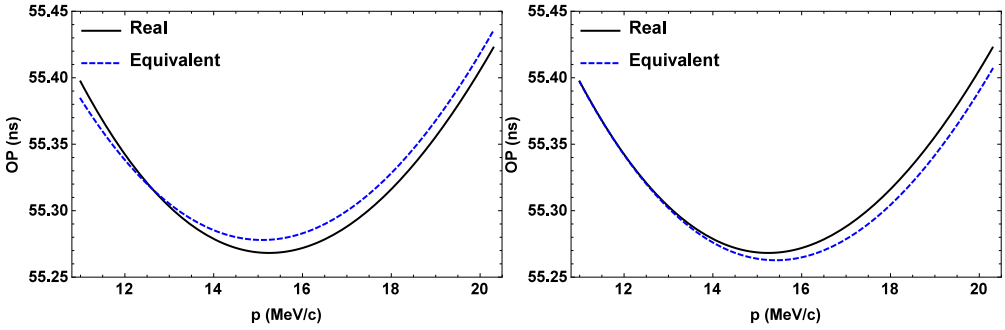

(a) Uncorrected orbital period (b) Corrected orbital period

FIGURE 4.24: Comparison of the field map orbital period curves calculated by using real and equivalent momenta. Figure (a) gives the orbital period of a 12.5 MeV/c particle for the equivalent momentum maps (blue dashed line). In Fig. (b) we account for change in speed of the particle, and multiply the equivalent momentum curve of Fig. (a) by β_0/β, where β_0 and β are,respectively, the relativistic speed of a particle at 12.5 MeV/c and a particle with a real momentum equal to the equivalent momentum. As expected, the two curves intersect at 12.5 MeV/c (Fig. (a)) and better agreement is seen between the two methods at lower momentum (Fig. (b)).

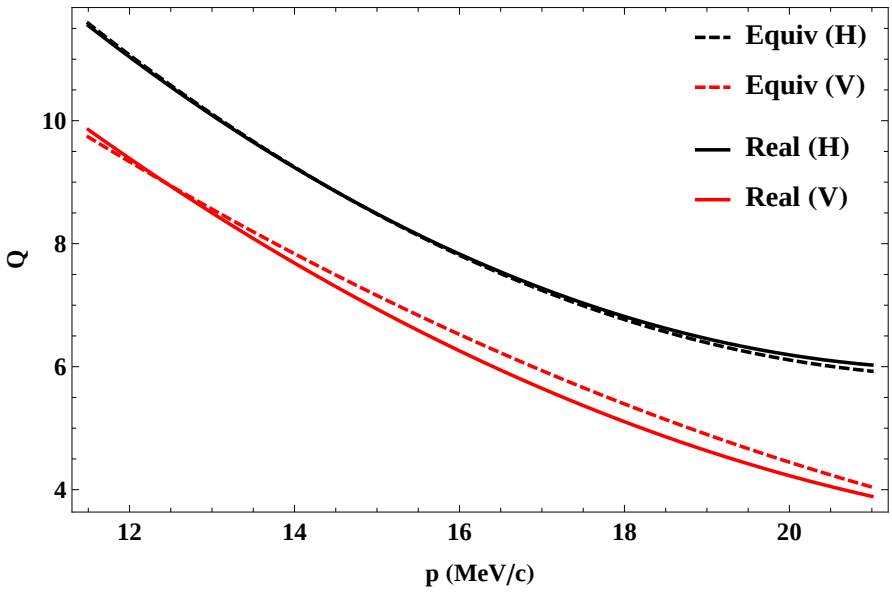

FIGURE 4.25: Comparison of the tune curves calculated by using real and equivalent momentum. The greatest difference between the curves is seen for the vertical tune, which is consistent with the larger field error being observed for the defocusing quadrupole in Fig. 4.22.

The implication of these results is that the experimental equivalent momentum procedure may not reproduce the dynamics of the baseline lattice exactly, and that different lattice configurations are being applied for each current setting of the quadrupoles. Figures 4.24 and 4.25 give an indication as to how large the errors on the beam dynamics may be as a result of the quadrupole excitation curves used.

4.4.3 Betatron tune and chromaticity

Betatron tune

The tune excursion of the EMMA lattice has already been measured experimentally using equivalent momenta, with data collected at 12, 14, 16, 18 and 20.3 MeV/c. For each momentum setting, 10 turns worth of data was collected for each of the BPMs that are sited between the defocusing and focusing quadrupoles, giving 420 data points per momentum (data collected for an equivalent momentum of 14 MeV/c is presented in Fig. 4.26 as an example). The initial calculation of the cell tunes (as presented in [60]) was based on splitting the data set for each momentum into 400 subsets of 21 points (e.g. data points 1 to 21 make the first subset, 2 to 22 the second, etc.), and applying the numerical analysis of fundamental frequency (NAFF) correlator [66] to each of the subsets. The mean tune value and standard deviation of all 400 subsets is then calculated. This approach for calculating the tune had been developed for application to an accelerating beam, for which the tune can change rapidly (with the advantage of the NAFF correlator being that the error on the calculated frequency should fall at a rate of $1/N^4$, where N is the number of data points). On reviewing the data for the original calculations of tune, it was seen that the calculation method used had led to a large

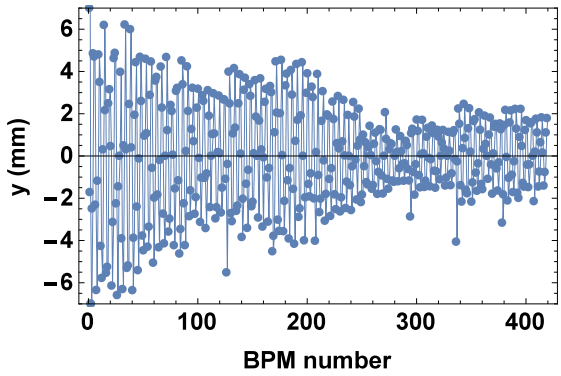

(a) Horizontal axis

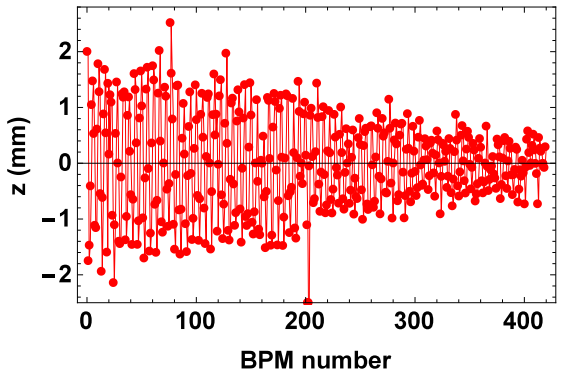

(b) Vertical axis

FIGURE 4.26: Plots showing the BPM data that is used to calculate the cell tune for 14 MeV/c particles. The vertical axes in the plots give the betatron oscillation of the bunch centroid along the horizontal (a) and vertical (b) transverse axis of EMMA. The horizontal axes of the plots are the order in which the BPMs are encountered by the bunch; the data has been collected using the BPMs located between the defocusing and focusing magnets in each of EMMA's 42 cells, meaning that the above plots show 10 turns worth of data. The betatron oscillation has been obtained for each axis by subtracting the average closed orbit position (the mean of the transverse position over all BPM measurements) from the individual BPM measurements.

standard deviation on the mean of the tune for some momenta. The cell tune values for the results presented in [60] have been recalculated for this thesis, with a discrete Fourier transform (DFT) applied to the entire 420 points of data for each set (table 4.2). Before presenting these results, we will further discuss the methods for tune calculation that have been applied.

For a DFT, a set of N equally spaced BPM readings are transformed from the time domain to N equally spaced coefficients in the frequency domain. Each coefficient is given by:

$$Y_\nu = \sum_{n=0}^{N-1} y_n e^{-2\pi i n \nu}, \tag{4.5}$$

where y_n is the transverse position of a particle bunch centroid with respect to the closed orbit at the n^{th} measurement, and $2\pi\nu$ is a phase advance between two successive BPMs. The coefficients are calculated for N values (between 0 and 1) of ν, with the individual

values of ν separated by intervals of $\frac{1}{N-1}$. The coefficients will be largest when ν is close to some harmonic of the betatron tune, and a mono-energetic particle bunch will give two peaks in the DFT coefficients in the frequency range 0 to 1 that correspond to Q and $1 - Q$ [2]. As the BPMs used for the EMMA tune calculation are separated by a single cell, and the phase advance through one cell is less than π radians for all momenta by design, the tune is given by the first peak of the DFT. The tune value is taken as the value of ν that has the largest DFT coefficient, and the error on the tune is taken as $1/N$.

If we consider Eq. 4.5 as a continuous function of ν, then we will see a sinc-like plot in the frequency domain (blue plot in Fig. 4.27a). This occurs as we truncated the oscillation in the time domain at N turns, which results in the plot in the frequency domain being a convolution of the frequency domain representation of the betatron oscillation and a rectangular pulse. Instead of finding the largest DFT coefficient for discrete values of ν, we can instead search numerically for the maximum of the main lobe of the sinc-like function (we use the NMaximize routine in Mathematica [67]); this is the main principle in the application of the NAFF correlator, which is given as:

$$Y_\nu = \sum_{n=0}^{N-1} \chi_n y_n e^{-2\pi i n \nu}. \tag{4.6}$$

The NAFF correlator has the same form as the DFT, apart from the inclusion of a Hann window:

$$\chi_n = \frac{1}{2}\left(1 - \cos\left(\frac{2\pi n}{N-1}\right)\right). \tag{4.7}$$

The Hann window reduces the effects of the truncation of the BPM signal in the time domain by scaling the signal to zero at the beginning and end of the signal, which is seen as a reduction of the side lobes in Fig. 4.27b. The combination of assuming a continuous range of ν and the application of the Hann window leads to an error on the calculated tune as low as $1/N^4$. In Fig. 4.28, we show the convergence of the tune calculation methods to the tune in the case that the betatron oscillation is given by a single particle (with $Q = 0.32$) and there is no noise in the measurement of the oscillation. In reality a particle bunch is made up of particles with a range of momenta (giving a range of frequency content within the oscillation measurement) and there will be some error on the measurement of the bunch centroid position. The momentum distribution ($\Phi(\delta)$) of the particles within the bunch leads to the decoherence of the bunch in transverse phase space (this is discussed further in section 4.5), the effect of which is that a measurement of the betatron oscillation appears dampened. A more appropriate model for the measurement of the transverse position of the bunch centroid

[2] If the DFT (Eq. 4.5) is applied to a pair of transverse variables that are based on the position (y) and momentum (p_y) that are written as a complex number, then just a single peak is found (corresponding to the tune, Q). This is discussed in greater detail later in this chapter.

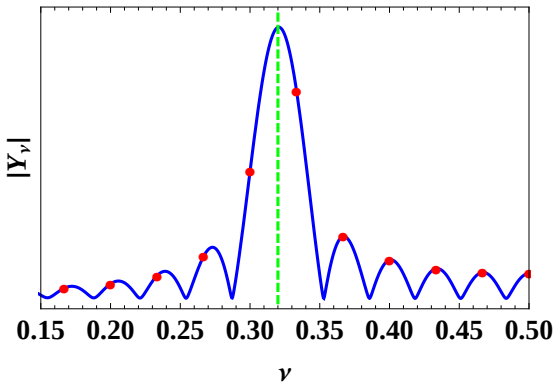

(a) DFT

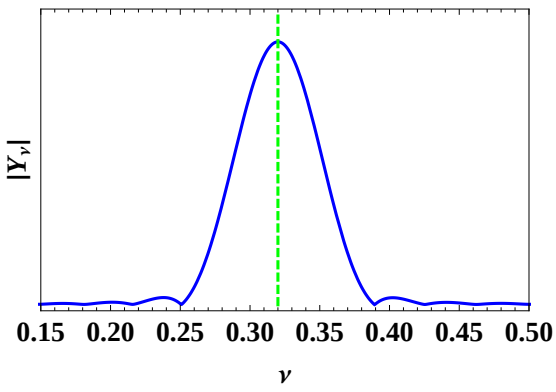

(b) NAFF correlator

FIGURE 4.27: We consider the betatron oscillation of a particle that has a tune of $Q = 0.32$ (indicated by the green dashed lines), the transverse position of the particle with respect to the closed orbit is measured over N turns at a reference point within the ring. A DFT can be performed on the turn by turn data to obtain N equally spaced coefficients in the frequency domain (red points in Fig. (a)); an estimate of the tune is found with an uncertainty of $1/N$ at the tune corresponding to the largest coefficient. We know that the signal from the betatron oscillation has only a single frequency, and the true value of the tune must be somewhere between the location of the two largest DFT coefficients; if we consider the DFT for continuous Q, then we plot a sinc-like function (blue plot in Fig. (a)), and we see that the peak of this plot corresponds to the true value of the tune. A numerical search for this peak is the basis of the application of the NAFF correlator. The sinc-like function results from the truncation of the signal at N points in the time domain: the NAFF correlator includes a Hann window which eliminates the discontinuity of the signal (due to truncation) in the time domain, and as a result, the side lobes are less prominent in Fig. (b). The combination of searching for the maximum of the NAFF correlator and windowing leads to uncertainties as low as $1/N^4$ in the tune calculation.

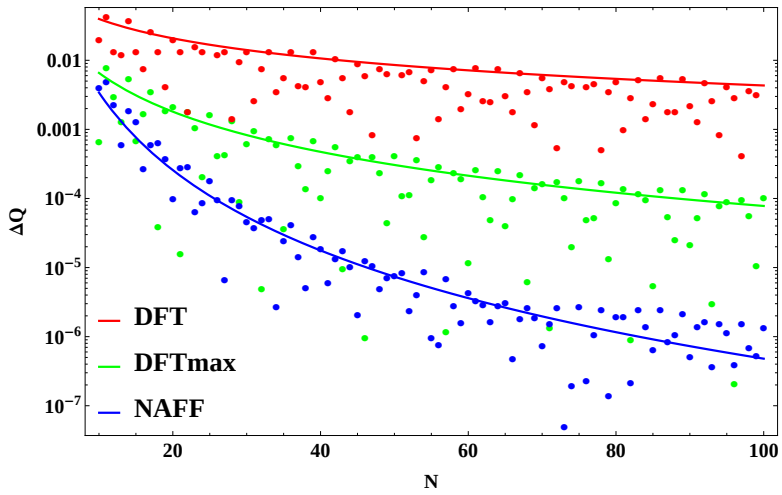

FIGURE 4.28: Plot to show the convergence of tune calculation methods to the actual tune value when increasing numbers of data points, N, are used. ΔQ is the absolute difference between the calculated and actual tune. We present the individual results for a tune calculation based on N data points, and a line drawn from the maximum value of ΔQ that falls in the expected fashion for a given method. Following on from Fig. 4.27, we consider application of the methods to the oscillation of a single particle with $Q = 0.32$: for DFT, the tune is estimated as being at the location of the largest DFT coefficient, the line changes proportional to $1/N$; for DFTmax, we assume a continuous range of Q in the frequency domain, and then search for the maximum numerically, the line changes proportional to $1/N^2$; for NAFF, we apply a Hann window to the oscillation data, and then search for the maximum in the frequency domain, the line changes proportional to $1/N^4$.

at the n^{th} turn is given by:

$$y_n = \sqrt{2\beta_y J_y} \int_{\delta_{\min}}^{\delta_{\max}} \cos\left(2\pi n(Q_y + \xi_y \delta)\right) \Phi(\delta) \mathrm{d}\delta, \tag{4.8}$$

where $\sqrt{2\beta_y J_y}$ gives the amplitude of the oscillation, ξ_y is the chromaticity of the lattice, δ is the fractional offset in momentum and $\delta_{\min}$ and $\delta_{\max}$ are respectively the minimum and maximum values of δ for the particle bunch. An example of such an oscillation, based on chromaticity and momentum spread values that are typical of the EMMA lattice, is given in Fig. 4.29.

Initially we consider the effects of decoherence and measurement errors on the calculation of the betatron tune independently. Starting with decoherence, we consider three values for the standard deviation of the momentum spread of the bunch: $\sigma_\delta = 0$, 5×10^{-4} and 1×10^{-3}. The convergence of the different calculation methods to a value of the tune given with increasing BPM measurements is shown in Fig. 4.30; it is seen that decoherence impacts the convergence of both the numerical search for the maximum of the DFT and the NAFF methods, whilst there is no notable difference in the convergence of the DFT calculation over the range of time domain sample sizes used. The NAFF correlator initially converges to a tune value at a greater rate when the momentum spread of the bunch (and therefore the rate of decoherence of the beam) is larger, however in Fig. 4.30c we note that the method appears stop converging after some number of turns

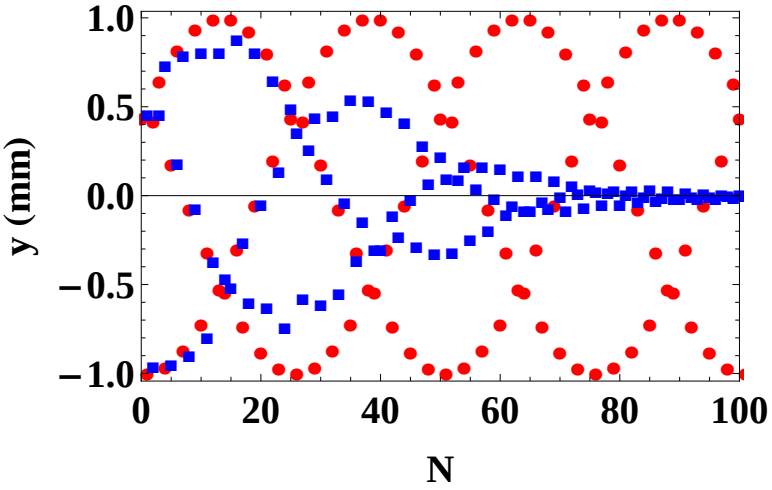

FIGURE 4.29: Plot of transverse position (y) of a bunch vs. turn number (N) for two separate models of the betatron motion. When we first applied the tune calculation methods, we assumed that the turn by turn signal of a bunch is equivalent to the sinusoidal oscillation of a single particle (shown in red). In reality, the BPM will return a signal that is subject to noise, and the measured oscillation of the bunch will appear dampened due to the decoherence of the beam in transverse phase space (blue points). In this instance we have assumed a lattice chromaticity of -10 with a Gaussian momentum spread that has a standard deviation of $\delta = 5\times10^{-4}$, a random error is then added to each position based upon a Gaussian distribution with a standard deviation of 10% of the measured position.

(for a larger momentum distribution, the beam decoheres rapidly and the amplitude of the measured oscillation will tend towards zero after a smaller number of turns).

We now look at the effect of errors in the measurement of the bunch position on the convergence of the tune calculation; we include errors in the turn by turn data, with the error on each measurement described by a Gaussian with a standard deviation given by $\sigma_y = 0$, $0.03y_n$ and $0.06y_n$. For the range of BPM signal lengths and errors considered, we do not see a significant change in the rate of convergence for the DFT and numerical search for DFT maximum tune calculation methods (Figs. 4.31a and 4.31b). However, the NAFF correlator deviates visibly from a convergence rate proportional to $1/N^4$ as the standard deviation of the measurement error is increased (Fig. 4.31c); when 100 turns of data are used, the difference between the tune calculated using NAFF for a signal with and without measurement errors is a couple of orders of magnitude.

Finally we consider the case where there is a decohering bunch and an error on the BPM measurement: the standard deviation of the momentum spread is $\sigma_\delta = 1\times10^{-3}$ and we again apply measurement errors with standard deviations of $\sigma_y = 0$, $0.03y_n$ and $0.06y_n$. The results are shown in Fig. 4.32, where the DFT calculation is again seen to converge at a rate of $1/N$. For the numerical search for the DFT maximum, we see it is the effects of decoherence that dominate in determining the rate of convergence (Fig. 4.32b), whilst for the NAFF correlator (Fig. 4.32c) the effects of both decoherence and measurement errors can be identified. We conclude by presenting the results for convergence of the different tune calculation methods on the same plot, based on a BPM

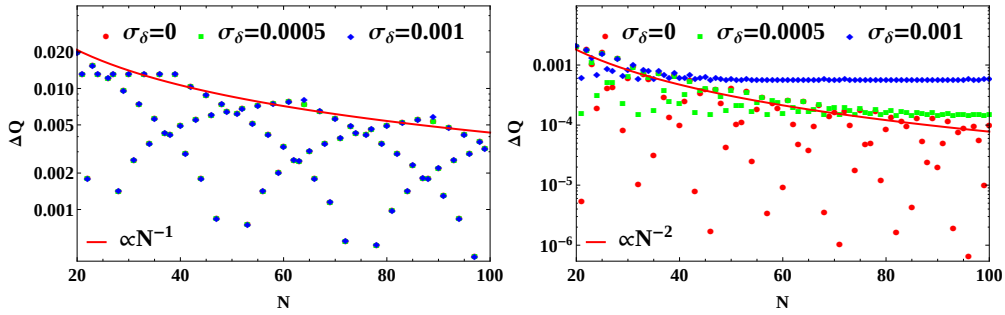

(a) DFT

(b) DFT with numerical search for maximum

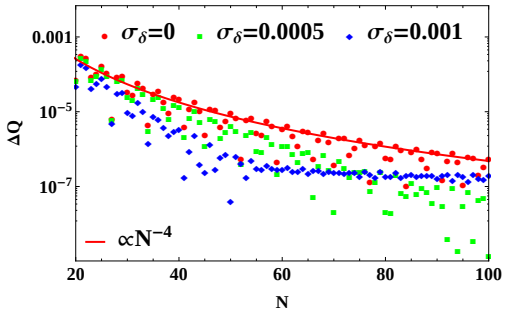

(c) NAFF

FIGURE 4.30: Plots to show the rate of convergence of the three tune calculation methods to the tune when applied to a BPM signal that appears dampened due to decoherence; the simulated BPM signals are based upon momentum distributions with standard deviations of $\sigma_\delta = 0$, 5×10^{-4} and 1×10^{-3}. In the case of the DFT (Fig. (a)), we see that the decoherence of the beam due to a momentum spread does not have a significant impact on the convergence of the calculated tune to the actual tune within the range of N considered (the line to show an error decreasing at a rate of $1/N$ is applicable to all three momentum distributions). When we consider a continuous range of Q, and search for the maximum of the DFT numerically (Fig. (b)) we see that the rate of convergence decreases with increasing momentum spread. In Fig. (c), we see that increasing the momentum spread of the bunch initially improves the rate of convergence of the calculated tune to the actual tune, however in the case of $\sigma_\delta = 1\times10^{-3}$ the convergence appears to stop after some number of turns.

signal with a momentum spread of $\sigma_\delta = 1\times10^{-3}$ and measurement errors of $\sigma_y = 0.06y_n$ (Fig. 4.33): we see that for any number of turns within the range considered, the NAFF correlator calculates the tune to the greatest precision, but that a rate of convergence proportional to $1/N^4$ is unattainable when applying NAFF to EMMA.

When calculating the cell tune, we take in to account that there will be some variation in the phase advance over the individual cells; applying a DFT to the 420 points that are described at the beginning of this section will calculate the betatron tune to an appropriate level of precision. In section 4.5, we calculate the fractional part of the ring tune using approximately 40 turns of data measured at a single BPM; in this instance it is more appropriate to use the NAFF correlator.

The cell tunes calculated by DFT are presented in table 4.2 and in Fig. 4.34 (where we show the previously published NAFF-based results for reference). In Fig. 4.35, we

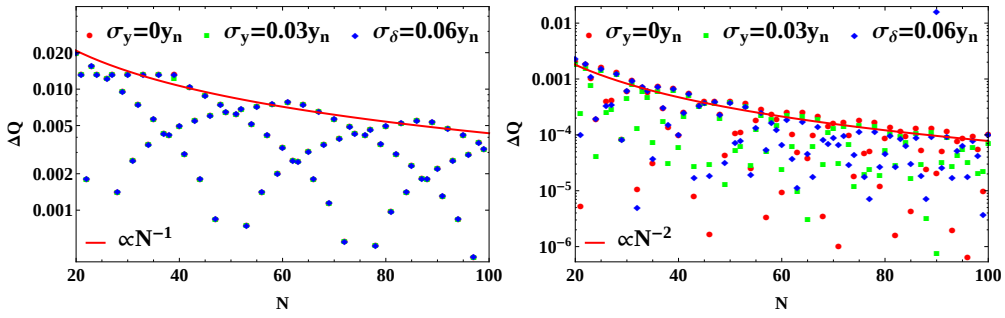

(a) DFT

(b) DFT with numerical search for maximum

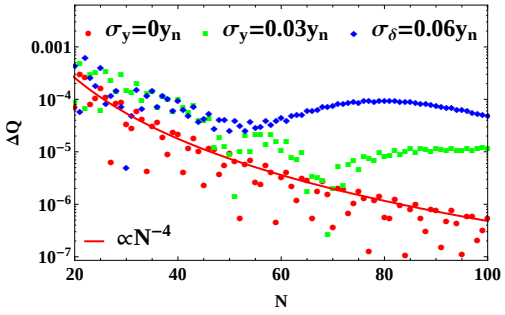

(c) NAFF

FIGURE 4.31: Plots to show the rate of convergence of the three tune calculation methods to the tune when applied to a BPM signal with measurement errors. For the range of turn numbers, N, and errors, σ_y, considered, there is no significant difference in the convergence of the tune calculation for either the DFT (Fig. (a)) or numerical search for DFT maximum (Fig. (b)) methods, which converge at a rate of $1/N$ and $1/N^2$ respectively. The rate of convergence for the NAFF method deviates from being proportional to $1/N^4$ when the standard deviation of the errors is increased (Fig. (c)), however NAFF still offers the best convergence of the tune calculation for the parameters investigated.

TABLE 4.2: Cell tunes were calculated for five equivalent momenta by applying a DFT to 420 data points.

p_{equiv} (MeV/c)	Q_{h}	Q_{v}
12.0	0.251	0.220
14.0	0.210	0.181
16.0	0.181	0.150
18.0	0.160	0.129
20.3	0.148	0.103

show the tune values obtained through the analysis of experimental data with the tune calculated earlier for the field map; the fit to the measured points is discussed in the following section.

Chromaticity

To obtain estimates of the tune and chromaticity for the entire momentum range, a polynomial fit to the measured tune values is carried out, with the appropriate order

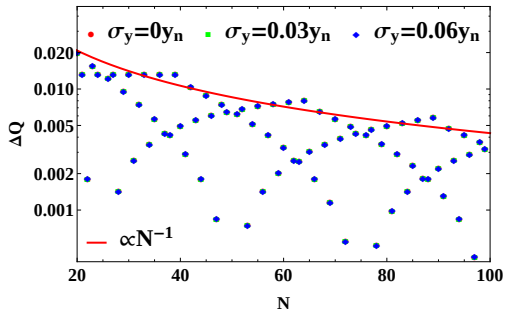

(a) DFT

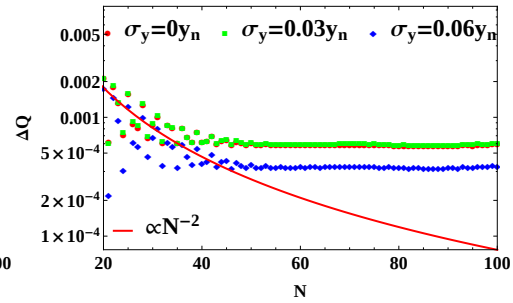

(b) DFT with numerical search for maximum

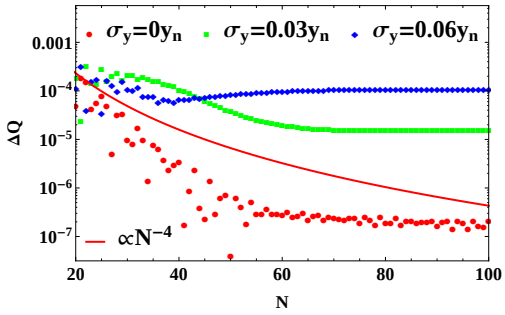

(c) NAFF

FIGURE 4.32: Plots to show the rate of convergence of the three tune calculation methods to the tune when applied to a BPM signal with measurement errors and when the beam is decohering (the bunch has a momentum spread of $\sigma_\delta = 1 \times 10^{-3}$). Again, we see that the convergence of the DFT remains proportional to $1/N$; for the DFT with a numerical search for the maximum we see the effects of the measurement errors, whilst for the NAFF method, we see the effects of both decoherence and measurement errors.

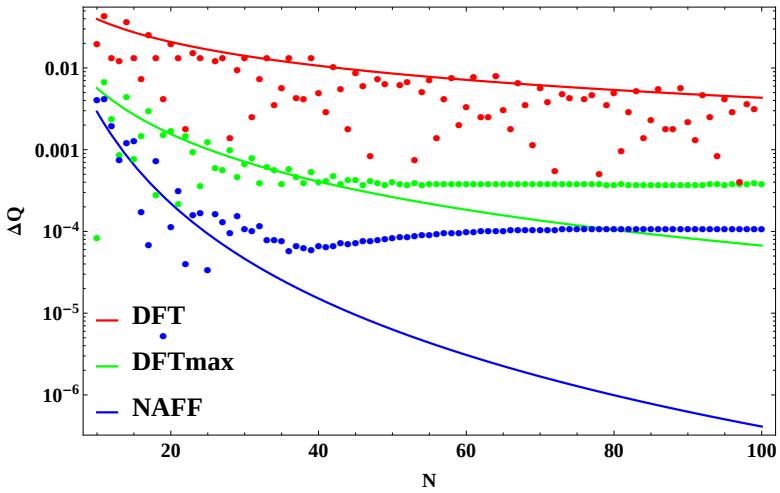

FIGURE 4.33: Plots to show the convergence of the tune calculation methods with increasing length of BPM signal (coloured points) when the signal is based on a momentum spread of $\sigma_\delta = 1 \times 10^{-3}$ and measurement error of $\sigma_y = 0.06y_n$, with lines to show the expected rate of convergence for the methods. For the signal lengths considered, the NAFF correlator gives the best rate of convergence, however, a rate proportional to $1/N^4$ (which was demonstrated in Fig. 4.28) is not attained.

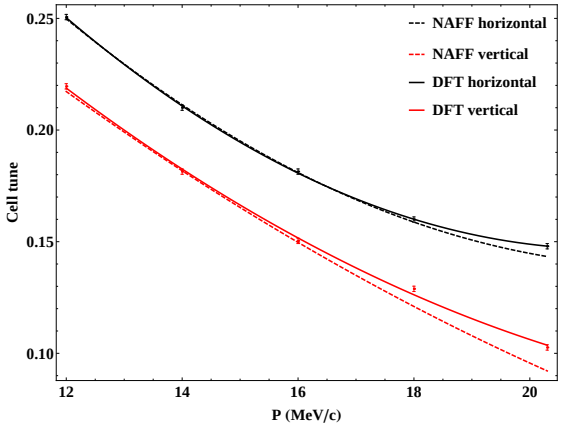

FIGURE 4.34: Measured cell tune vs. equivalent momentum. The tunes calculated through NAFF have previously been published, however it is the values calculated with DFT that are used within this thesis. The DFT was performed on a sample of 420 data points for each momentum.

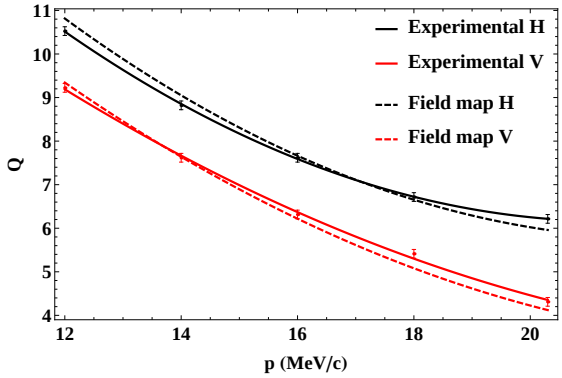

FIGURE 4.35: Plot to show the experimentally measured ring tune (calculated by multiplying the cell tune by 42) alongside the tune calculated by tracking particles of a range of momenta through a field map calculated using nominal EMMA quadrupole currents.

of the polynomial determined by further analysis of the field map data. Initially, a quadratic fit is made to the tunes calculated after tracking particles of 18 different momenta through the field map for nominally powered EMMA quadrupoles (in Fig. 4.35 a quartic fit is shown for this data). Given the model:

$$Q_{y(z)}(p) = a_{y(z)} + b_{y(z)}p + c_{y(z)}p^2, \tag{4.9}$$

then the chromaticity is calculated as

$$\xi_{y(z)}(p) = p\frac{dQ_{y(z)}}{dp} = p(b_{y(z)} + 2c_{y(z)}p). \tag{4.10}$$

We compare the chromaticity given by Eq. 4.10 with the local chromaticity calculated for each of the 18 momenta used when tracking through the field map. The method for calculating the local chromaticity was to track two particles with small fractional

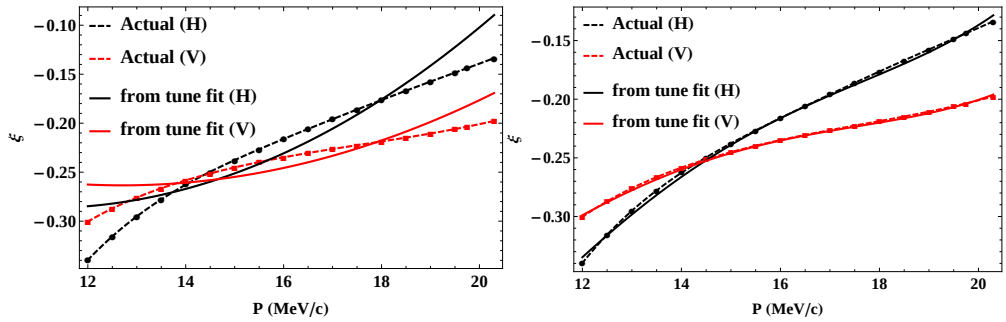

(a) Quadratic fit to the tune data. (b) Quartic fit to the tune data.

FIGURE 4.36: Chromaticity in a field map model of the EMMA lattice. The chromaticity is calculated (points and dashed lines) by tracking two particles with small deviations from the reference momentum and then calculating the difference in tune between the two particles. The solid lines are obtained by first finding a polynomial fit to the tune data, and then finding the chromaticity based on the derivative of the polynomial. It can be seen that a fourth order polynomial is needed in order to accurately describe the chromaticity.

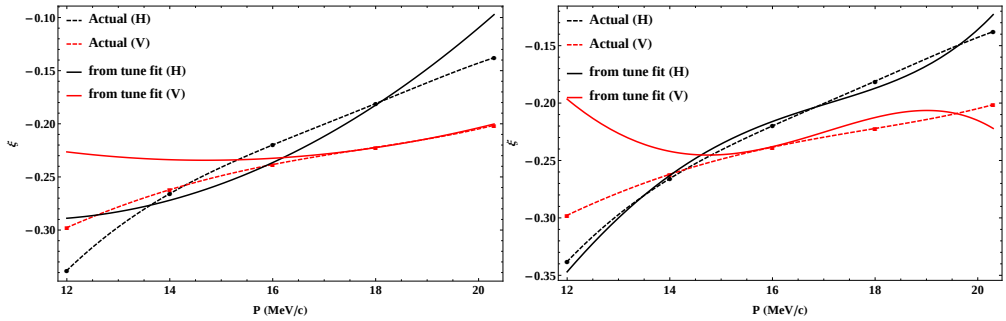

(a) Quadratic (b) Quartic

FIGURE 4.37: Tune measurements have been carried out for five equivalent momenta, meaning that there are insufficient data to determine accurately the coefficients of a quartic fit (demonstrated in the above plots using the equivalent momentum field map data).

momentum offsets of $\delta = \pm 0.001$ around each of the 18 momenta, and then calculate the chromaticity based on the difference in tune between these two particles.

Figure 4.36 shows that a quadratic fit to the tune data does not provide an accurate calculation of the chromaticity and that a quartic fit is more appropriate. However, as the tune has been measured experimentally for just five momenta, it may not be possible to determine the coefficients of the quartic fit with good accuracy (Fig. 4.37). Therefore, a quadratic fit is made to the tune data, and the differences between the actual chromaticity and the chromaticity found through the polynomial fit in Fig. 4.37 are used as an indicator of the systematic error that may be introduced.

The quadratic fit to the measured tune values returns the coefficients (as defined in Eq. 4.9) $a_y = 0.686 \pm 0.012$, $b_y = -0.0505 \pm 0.0015$, $c_y = 0.00118 \pm 0.00005$, $a_z = 0.550 \pm 0.038$, $b_z = -0.0357 \pm 0.0048$ and $c_z = 0.00068 \pm 0.00015$. The measured tune values and polynomial fit are shown in Fig. 4.34.

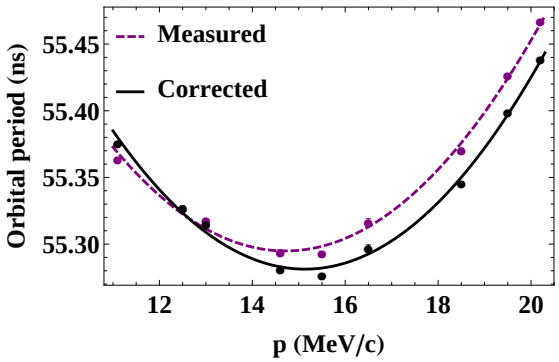

FIGURE 4.38: The orbital period of a particle in EMMA is dependent upon the particle's momentum. This is due to both a change in the single revolution path length and a change in speed. The orbital period dependence on momentum has been measured using equivalent momenta (purple dashed line), which accounts only for the change in path length, a correction has then been applied to account for the change in speed (black line).

4.4.4 Orbital period

The orbital period (or time of flight) has been measured for nine settings of equivalent momentum. The orbital period is dependent upon both the path length of an orbit and the speed of the particle; using equivalent momentum accounts for just the path length dependence (with the true particle momentum being fixed at 12.5 MeV/c), therefore the measurements are corrected for speed during analysis (by scaling each measurement by $\beta(12.5\,\mathrm{MeV/c})/\beta(\mathrm{p_{equiv}})$). The data for the measurements are recorded using an oscilloscope (with a 40 GS/s sampling rate), which has a single EMMA ring BPM and the ALICE rf signal (1.300 GHz) as inputs. The orbital period is found by measuring (for a number of turns) the phase of the rf at which the derivative of the BPM signal crosses zero (i.e. when the voltage at the BPM button is maximum), the phase slip then gives the difference between the rf period and orbital period, allowing for the orbital period to be calculated. A second-order polynomial is fitted to the corrected orbital period data, giving:

$$\mathrm{OP}(p) = (56.67 \pm 0.05) - (0.184 \pm 0.007)p + (0.0061 \pm 0.0002)p^2 \qquad [\mathrm{ns}],$$

where p is the particle momentum (with units of MeV/c). Both the corrected and uncorrected orbital period curves are shown in Fig. 4.38. In Fig. 4.39 we show the corrected experimentally measured orbital period along side the orbital period obtained through field map tracking (again, we track particles of a range of momenta through the field map that represents the nominal quadrupole current settings).

4.4.5 Closed orbit position

In section 4.4.3, we applied DFT and NAFF to a BPM signal in order to calculate the betatron tune. The position of the bunch centroid that is returned by a BPM as a bunch

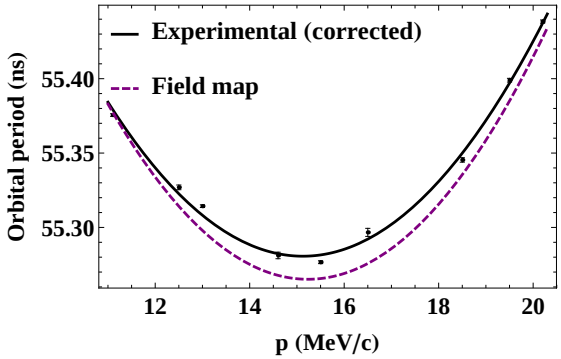

FIGURE 4.39: Plot to show the orbital period measured experimentally (with speed correction) and the orbital period obtained by tracking articles of different momenta through the field map for an EMMA period (with nominally powered quadrupoles).

makes its n^{th} turn in the ring may be broken down into three components:

$$y_n = \langle y_{\text{co}} \rangle + y_{\Delta\text{co}} + y_{\beta,n}, \tag{4.11}$$

where $\langle y_{\text{co}} \rangle$ is the mean closed orbit position for all BPMs, and $y_{\Delta\text{co}}$ and $y_{\beta,n}$ are, respectively, the closed orbit distortion and betatron oscillation at a single BPM. In the case of the tune calculation, we sought to remove the closed orbit terms from Eq. 4.11 so that they do not contribute to the calculated frequency spectrum (e.g. $\langle y_{\text{co}} \rangle$ will give a peak in the frequency spectrum at integer ν): to do this, we focussed on each BPM individually, and subtracted the mean of the BPM signal over all turns for a single BPM from the turn by turn data for the same BPM. We did this for the vertical axis as well as the horizontal axis due to the non-zero closed orbit distortion for the vertical axis. Now we want to find the mean closed orbit position ($\langle \langle y_{\text{co}} \rangle \rangle$); this will be close to zero along the vertical axis for all momenta, therefore we focus on the horizontal axis only. To obtain the mean closed orbit position, we calculate the mean of the closed orbit terms that were found for the individual BPMs.

We find the horizontal momentum-dependent closed orbit with the same BPM data that was used earlier for the tune calculation; this data was collected using the 42 BPMs that are located between the defocusing and focusing quadrupole, which is marked by the black dashed line in Fig. 4.16 (one of the BPMs was not responding, so the calculation is based on data from 41 BPMs). A spread in the closed orbit positions calculated for the BPMs at a given momentum will be due to both statistical errors and closed orbit distortion: the mean closed orbit position across all of the BPMs is presented in Fig. 4.40.

4.5 Further measurements from a decohering signal

A bunch of particles injected into EMMA with some transverse offset from the closed orbit will perform betatron oscillations about the closed orbit. In the absence of acceleration, BPMs will show that the amplitude of oscillation decreases with turn number

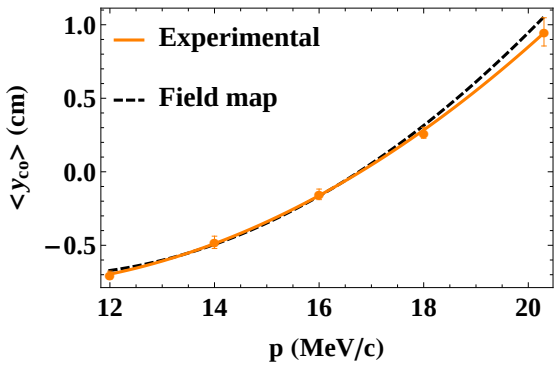

FIGURE 4.40: Plot of closed orbit position vs. momentum at the location of a reference BPM. The orange points show the mean horizontal closed orbit position for 41 of the BPMs located between the defocusing and focusing quadrupoles (with the standard deviation given as the error) vs. momentum; the orange line gives the fit of a second order polynomial to these data points. The dashed black line shows a fit to the closed orbit found after tracking particles of 9 different momenta through the field map for nominally powered EMMA quadrupoles.

(this was introduced in section 4.4.3 when we discussed tune calculation methods). The amplitude of each particle in the bunch has not changed. The particles within the bunch have different betatron frequencies from each other, meaning that their oscillations are no longer in phase after a number of turns. As a result, the centroid of the bunch moves towards the closed orbit. This effect is called decoherence. Decoherence impacts on accelerator performance and has been widely studied [68–72]. Two main causes of decoherence may be identified when looking at the transverse dynamics of a circular accelerator. The first is the momentum spread of the bunch coupled with the chromaticity of the accelerator, whilst the second is the emittance of the bunch coupled with the transverse nonlinearities in the magnetic fields which make up the accelerator.

With no nonlinear magnetic fields (e.g. sextupole) included within EMMA to control chromaticity, we can assume that the rapid decoherence (within tens of turns) of particle bunches observed experimentally is caused by the bunch momentum spread coupled with chromaticity. Decoherence makes the measurement of lattice and bunch values, such as Courant-Snyder parameters and betatron amplitude, difficult. In this section, we will present a method for reconstructing the momentum distribution of particles within a bunch from BPM measurements. Further analysis of the same BPM data allows estimates to be made of the Courant-Snyder parameters and the amplitude of coherent betatron oscillation of the beam. We consider the case where the rf cavities are turned off, meaning that particles ideally make no synchrotron oscillations and the beam energy can be assumed constant. In practice, particles can lose energy (slowly) through beam loading effects; the impacts of these effects is considered in section 4.5.6.

Standard techniques to measure the momentum distribution may be broadly divided into two groups. One group of techniques is largely limited to an estimation of the energy spread. They are based on the assumption that particle distributions within a bunch are Gaussian in the transverse and longitudinal directions. Under these conditions, analytic

solutions for the energy spread are available [68]. The actual implementation can be based on BPM measurements [73, 74] or radiation measurements [75, 76]. In these techniques, the detailed momentum distribution is not obtained.

The other group of techniques measure the momentum distribution. They are based on tomography [77–79]. The measurement setup usually requires an rf cavity, a number of dipole magnets, scintillating screens and cameras. These techniques are not only able to measure the momentum distribution, but they can also reconstruct the longitudinal phase space. This group of techniques are usually destructive - the beam would not continue along the beamline after hitting a screen.

Most commonly, BPMs are used to measure the closed orbit distortion at a given location within a lattice, however they have been applied to measuring other dynamical attributes of an accelerator; for example, given known longitudinal dynamics, BPMs have been used to calculate the chromaticity of a lattice [80]. We present a method for using BPM measurements to reconstruct a momentum distribution, and not just the spread, of a bunch when the lattice chromaticity is known. The method is valid when a bunch decoheres due to the momentum distribution of a bunch. It does not give the longitudinal phase space, but it is not destructive.

4.5.1 Finding the momentum distribution

When a particle bunch is injected into a circular accelerator or storage ring with a large enough transverse offset from the closed orbit (as shown in Fig. 4.41a), then all particles within the bunch may be thought of as being at approximately the same starting phase of betatron oscillation. The individual momenta of particles within the bunch are distributed around a mean value. This, coupled with the chromaticity of the lattice, gives a spectrum of betatron tunes. The range of phase advance per turn for particles results in the spreading of the particles in phase space as particles travel through an increasing number of turns (Fig. 4.41b). Initially, a measurement of the position of the bunch centre of mass (as would be observed with a BPM) remains consistent with the position of a single particle which has been tracked through an equal number of turns, and has starting conditions matching those of the bunch centroid. As the number of turns increases, the spreading of particles in phase space continues, a ring is formed around the closed orbit, and the centre of mass of the bunch tends towards the position of the closed orbit (Figs. 4.41c and 4.41d).

At any point within the lattice the transverse position of the centroid of a mono-energetic bunch is given on the n^{th} turn by:

$$y_n = \sqrt{2\beta_y J_y}\cos(2\pi n Q_y + \phi_0), \tag{4.12}$$

where, for a given transverse axis (horizontal or vertical), β_y and ϕ_0 are the Courant-Snyder beta function and the phase of oscillation on the initial turn ($n = 0$) at a given lattice location. J_y and Q_y are the action and the betatron tune, and are independent

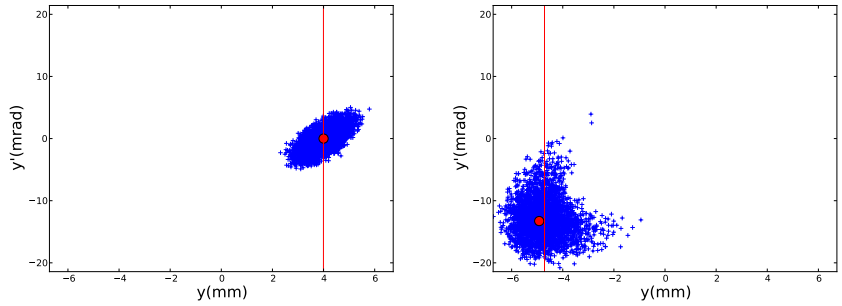

(a) 1st turn

(b) 10th turn

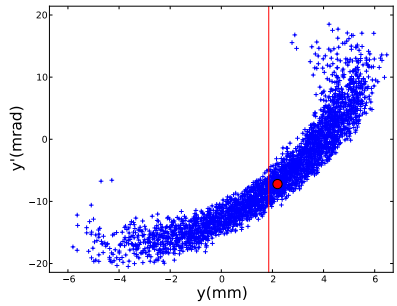

(c) 20th turn

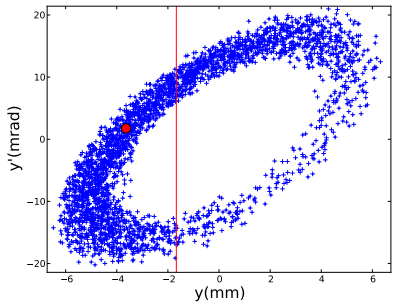

(d) 40th turn

FIGURE 4.41: Simulations showing the decoherence of a particle bunch in horizontal transverse phase space. A total of 4000 particles, with a Gaussian momentum distribution ($\sigma = 50\,\mathrm{KeV/c}$) around the central momentum of 12.1 MeV/c, are tracked through the EMMA lattice using Zgoubi. The bunch centroid has an initial horizontal offset of 4 mm from the closed orbit position. Blue points show the location in phase space of individual particles, the single red point shows the position of an on momentum particle which, at injection, is located at the centre of the bunch, and the red line shows the bunch centre of mass (mean y position) for the specified turn.

of the longitudinal position within the lattice.

To take into account the momentum spread of a bunch, Eq. 4.12 may be modified to Eq. 4.13:

$$y_n \approx \sqrt{2\beta_y J_y} \int_{\delta_{min}}^{\delta_{max}} \cos(2\pi n(Q_y + \xi_y \delta) + \phi_0)\Phi(\delta)d\delta, \tag{4.13}$$

where ξ_y is the linear part of the chromaticity $\left(\frac{\partial Q_y}{\partial \delta}\right)$ and δ is the fractional offset from a reference momentum $\left(\frac{\Delta p}{p_0}\right)$. $\Phi(\delta)$ is the momentum distribution weight function, which gives the relative contribution of different values of δ to the entire bunch. δ_{min} and δ_{max} are the minimum and maximum values of δ which are found within the bunch. The exact shape of the momentum distribution is unknown.

In Eq. 4.13, the closed orbit dependence on momentum and the nonlinear terms of chromaticity are not accounted for. The first of these two factors may be neglected by considering only the vertical z axis (for which dispersion is ideally 0), whilst the nonlinear terms of chromaticity should have little effect provided that both the width of the momentum distribution, $\Phi(\delta)$, and the coefficients of the nonlinear chromaticity

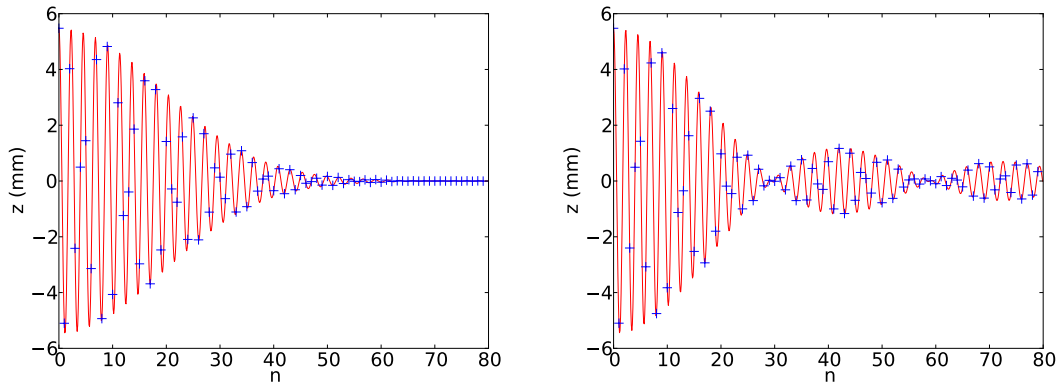

(a) Gaussian momentum distribution

(b) Uniform momentum distribution

FIGURE 4.42: Plots of bunch centroid vs. turn number. The way in which the particle beam decoheres is dependent not only upon the spread of momentum, but also the form of the momentum distribution.

terms are small.

The discrete time Fourier transform of measurements made by BPMs can be used to reveal the frequency spectrum of particles within a bunch. If the source of the decoherence is the momentum spread of the bunch coupled with the chromaticity of the lattice, then an estimate of the momentum distribution can be made. As the form of the momentum distribution is unknown, using both phase space variables (coordinate and momentum) offers the advantage of being able to extract both the even and odd parts of the distribution. It is convenient to use normalised coordinates, f_n, given by:

$$f_n = \hat{z} + i\hat{p}_{z,n}, \tag{4.14}$$

where

$$\hat{z} = \frac{z_n}{\sqrt{\beta_z}},$$
$$\hat{p}_{z,n} = p_{z,n}\sqrt{\beta_z} + \frac{\alpha_z z_n}{\sqrt{\beta_z}}. \tag{4.15}$$

$p_{z,n}$ is the transverse momentum of the bunch centroid measured with respect to the closed orbit transverse momenta along a given axis and α_z is the Courant-Snyder alpha lattice parameter. When taking into account the momentum distribution of the bunch, the progression of the normalised coordinates of the bunch centroid with turn number (n) is given by:

$$f_n = \sqrt{2J_z}e^{-i\phi_0}\int_{\delta_{min}}^{\delta_{max}} e^{-i2\pi n(Q_z+\xi_z\delta)}\Phi(\delta)d\delta. \tag{4.16}$$

Note that, in principle, $\Phi(\delta)$ can be a periodic function of δ, with period $\frac{1}{\xi_z}$. If $\Phi(\delta+\frac{m}{\xi_z}) = \Phi(\delta)$ for integer m, then replacing the limits of the integral in Eq. 4.16 by $\delta_{min} + \frac{m}{\xi_z}$ and $\delta_{max} + \frac{m}{\xi_z}$ leads to the same set of values f_n, for any value of m. This means that mathematically, we cannot determine $\Phi(\delta)$ uniquely for a set of observables f_n. The

best we can do is determine a function $\Phi_p(\delta)$, that has periodicity ξ_z. The important point here is that in the case of purely linear chromaticity, a BPM will observe the same frequency of oscillation for a particle with energy deviation $\delta + m/\xi_z$ (where m is an integer) as for a particle with energy deviation δ. It is therefore necessary to make some assumption about the range of energy deviation for particles in the beam; which is expressed in a change in the limits of the integral in Eq. 4.16 (from $\delta_{\min}$ and $\delta_{\max}$ to $-1/2\xi_z$ and $1/2\xi_z$), and in replacing the function $\Phi(\delta)$ by the periodic function $\Phi_{\mathrm{P}}(\delta)$. With these replacements, Eq. 4.16 becomes:

$$f_n = \sqrt{2J_z}e^{-i\phi_0}\int_{-\frac{1}{2\xi_z}}^{\frac{1}{2\xi_z}} e^{-i2\pi n(Q_z+\xi_z\delta)}\Phi_{\mathrm{P}}(\delta)d\delta. \tag{4.17}$$

In the case that $-\frac{1}{2\xi_z} < \delta_{min}$ and $\delta_{max} < \frac{1}{2\xi_z}$, then $\Phi(\delta)$ and $\Phi_{\mathrm{P}}(\delta)$ will be the same in the range $-\frac{1}{2\xi_z} < \delta < \frac{1}{2\xi_z}$. At this point we observe that Eq. 4.17 is, formally, the inverse Fourier transform of the momentum distribution $\Phi_{\mathrm{P}}(\delta)$. We can consider f_n to be the time domain representation of $\Phi_{\mathrm{P}}(\delta)$. It can be seen in Fig. 4.42 that a Gaussian distribution in $\Phi_{\mathrm{P}}(\delta)$ leads to a Gaussian decoherence signal in the time domain, and that a uniform $\Phi_{\mathrm{P}}(\delta)$ distribution leads to a sinc-like signal.

The chromaticity of EMMA (measured to be ~ -10 [60]) allows for a $\pm 5\%$ momentum spread in the range $\delta = -\frac{1}{2\xi_z}$ to $\frac{1}{2\xi_z}$. Previously, a momentum spread of 100 keV/c (at 15 MeV/c, giving a $\pm 0.3\%$ momentum spread) has been measured in the EMMA injection line [81], meaning that $\Phi_{\mathrm{P}}(\delta)$ should correspond to the momentum distribution of the bunch and should not be affected by the effects that would occur in the case that some particles have $|\delta| > \frac{1}{2|\xi_z|}$. In the case of having an infinite number of samples of the BPM signal, the discrete time Fourier transform (DTFT) can give the momentum distribution, $\Phi_{\mathrm{P}}(\delta)$. For the finite number of BPM data samples obtained experimentally, an estimation of the momentum distribution can be found by:

$$\Phi_{\mathrm{p}}(\delta) \approx e^{i\phi_0}\sum_{n=-N}^{N} f_n e^{i2\pi n(Q_z+\xi_z\delta)}, \tag{4.18}$$

where N is the total number of turns of BPM data to which the calculation is applied, and assuming that the terms for the negative turn numbers may be found by considering the complex conjugate of f_n:

$$f_{-n} = e^{-i2\phi_0} {f_n}^*.$$

Finally, equation 4.18 may be written as

$$\Phi_{\mathrm{p}}(\delta) \approx e^{i\phi_0} f_0 + 2\Re\left(e^{i\phi_0}\sum_{n=1}^{N} f_n e^{i2\pi n(Q_z+\xi_z\delta)}\right). \tag{4.19}$$

4.5.2 Measurement of the initial phase and lattice functions alpha and beta

The expression for reconstructing the momentum distribution given in Eq. 4.19 requires that the initial betatron oscillation phase, ϕ_0, and the Courant-Snyder parameters, α_z and β_z, are known. In this section, we present a method for measuring these parameters where the amplitude of coherent betatron oscillations (measured by the BPMs) is damped because of decoherence.
Firstly, a rotation of $-\Psi_n$ is applied to the normalised phase space coordinates of Eq. 4.15:

$$g_n = (\hat{z}_n + i\hat{p}_{z,n})\, e^{i\Psi_n}. \tag{4.20}$$

If Ψ_n is the betatron phase advance between injection and the n^{th} turn for a mono-energetic bunch ($\Psi_n = 2\pi n Q_z$), then g_n is equal to the normalised phase space coordinates at turn zero (f_0) for all values of n. Accordingly, the argument of g_n will be constant for all n, with value equal to the initial phase of the betatron oscillation, ϕ_0.

In practice, z and $p_{z,n}$ are found through using BPMs and Ψ_n may be calculated after first finding the betatron tune, Q_z, by applying the NAFF correlator to the BPM data. The design values of α_z and β_z at the BPM position are taken as initial estimates. If the experimental values of α_z and β_z do not match the design values, then a plot of the argument of g_n vs. n will show an oscillation around ϕ_0. A fitting procedure, which has the objective of minimising the oscillation of $\arg(g_n)$ vs. n around ϕ_0 and has α_z and β_z as free parameters, can be used to find the experimental values of α_z and β_z.

When the bunch has some momentum distribution, then the transverse coordinate and momentum, z_n and $p_{z,n}$, are described by:

$$z_n = \sqrt{2\beta_z J_z} \int_{-1/2\xi_z}^{1/2\xi_z} \cos(\Psi_n + \psi_n(\delta) + \phi_0)\Phi(\delta)d\delta,$$

$$p_{z,n} = -\sqrt{\frac{2J_z}{\beta_z}} \int_{-1/2\xi_z}^{1/2\xi_z} (\alpha_z \cos(\Psi_n + \psi_n(\delta) + \phi_0) + \sin(\Psi_n + \psi_n(\delta) + \phi_0))\Phi(\delta)d\delta.$$

In this case, Ψ_n is the total phase advance between injection and the n^{th} turn for an on-momentum particle ($\delta = 0$), with on-momentum further defined as being the momentum at which the mean betatron oscillation frequency is found. $\psi_n(\delta)$ is the change in total phase advance due to a particle being off-momentum ($\delta \neq 0$), and is given by $\psi_n(\delta) = 2\pi n \xi_z \delta$.

To show that the argument of g_n in Eq. 4.20 still gives the initial betatron phase when a bunch is not mono-energetic, the ratio of imaginary to real parts of g_n is considered:

$$\begin{aligned}\frac{\Im(g_n)}{\Re(g_n)} &= -\frac{\int_{-1/2\xi_z}^{1/2\xi_z} \sin(\psi_n(\delta)+\phi_0)\Phi(\delta)d\delta}{\int_{-1/2\xi_z}^{1/2\xi_z} \cos(\psi_n(\delta)+\phi_0)\Phi(\delta)d\delta} \\ &= -\frac{\int_{-1/2\xi_z}^{1/2\xi_z} (\sin(\psi_n(\delta))\cos(\phi_0)+\cos(\psi_n(\delta))\sin(\phi_0))\Phi(\delta)d\delta}{\int_{-1/2\xi_z}^{1/2\xi_z} (\cos(\psi_n(\delta))\cos(\phi_0)-\sin(\psi_n(\delta))\sin(\phi_0))\Phi(\delta)d\delta}.\end{aligned} \tag{4.21}$$

It can be seen from equation 4.21 that if:

$$\int_{-1/2\xi_z}^{1/2\xi_z} \sin(\psi_n(\delta))\Phi(\delta)d\delta = 0, \tag{4.22}$$

then:

$$\frac{\Im(g_n)}{\Re(g_n)} = -\tan(\phi_0). \tag{4.23}$$

The conditions for Eq. 4.22 being true are that either the momentum distribution is symmetric around $\delta = 0$ (in which case Eq. 4.23 will be true for all n), or that the momentum distribution is non-symmetrical but the values of $\psi_n(\delta)$ are small and are within the linear part of the sine function in Eq. 4.22 (in which case Eq. 4.23 will be true for some limited range of n, $0 \leq n \leq N$).

4.5.3 Simulation

To demonstrate the ability to reconstruct a momentum distribution, the methods described in sections 4.5.1 and 4.5.2 are applied to data produced through simulation. A total of 4000 particles were tracked through 60 turns of the hard edge representation of the EMMA lattice (as described by table 4.1) by using the Zgoubi tracking code. The transverse dynamical variables of the input particles were distributed randomly on ellipses matched to the lattice at the point of injection and with the action of the individual particles governed by a exponential distribution of width corresponding to an emittance of 0.71 mm mrad [82]. The bunch centroid at injection was located at the closed orbit position along the horizontal transverse axis, and had an offset of 4 mm from the closed orbit position along the vertical transverse axis.

A reference momentum of 12.1 MeV/c was set, and then a Gaussian distribution, with $\sigma_\delta = 0.001$, used to give each particle a random fractional momentum offset, δ. For each revolution of the bunch within the lattice, the mean offset for all particles in position and momentum from the respective transverse closed orbit values was recorded at a location corresponding to the position of one of the BPMs. With the given simulation input parameters, the effects of decoherence are clearly visible in figure 4.43. Applying the methods described in section 4.5.2 for the calculation of the Courant-Snyder parameters to the bunch centroid tracking data, α_z and β_z were found to be -0.79434 and 0.39750 m

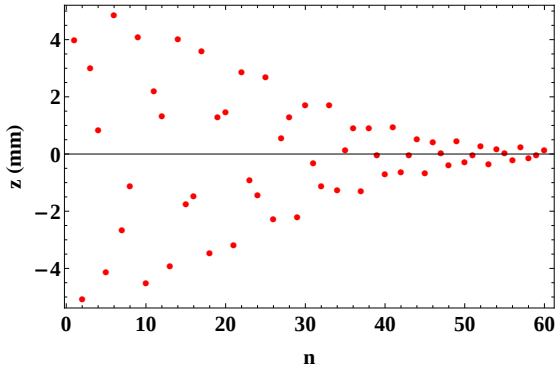

FIGURE 4.43: Vertical offset of bunch centroid from closed orbit position at location of BPM (z) vs. turn number (n). A Gaussian momentum distribution with a standard deviation of 12 keV/c means that the betatron oscillation of the bunch centroid appears significantly damped within 60 turns.

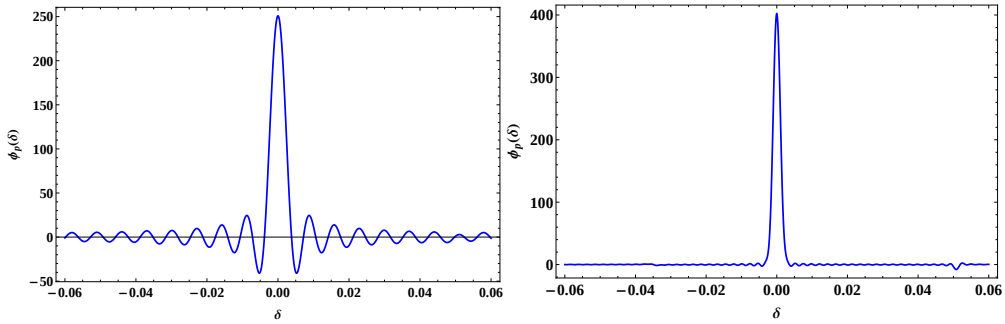

(a) Truncation at 20 turns.

(b) Truncation at 60 turns.

FIGURE 4.44: Plots of the discrete time Fourier transform in cases where the BPM signal is truncated at 20 turns and at 60 turns.

respectively, which is less than 0.01% difference from the values obtained through the tracking of a single particle ($\alpha_z = -0.79437$ and $\beta_z = 0.39749\,\mathrm{m}$).

The accurate reconstruction of the momentum distribution is dependent upon having enough turns to give good resolution in the DTFT integral range $\left(\frac{1}{\xi_z}\right)$, and upon the amplitude of the BPM oscillation signal having approximately converged to zero at the turn number of truncation. In practice, the number of turns for which data are obtained is limited. Therefore, we investigate the effect of truncating the BPM signal at a reduced number of turns on the convergence of the reconstructed momentum distribution.

For truncation at a small number of turns, as shown in Fig. 4.44a, the estimation of the discrete time Fourier transform is distorted by the truncation, and a sinc-like distribution is produced (as discussed earlier in section 4.4.3). As the number of turns increases, it is possible to resolve the features of the momentum distribution (Fig. 4.45), and the influence of truncation becomes less prominent. As it is known that the momentum distribution must have only positive values and that the side lobes of the transform are introduced by the truncation of the BPM signal, then only the positive part of the main lobe of the transformation is taken as being the reconstruction of the momentum distribution.

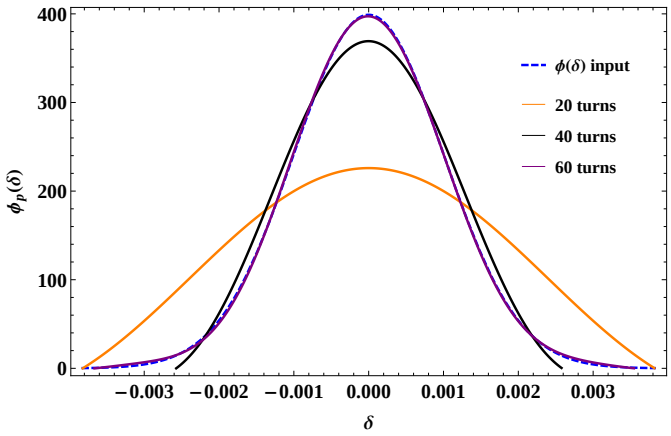

FIGURE 4.45: The reconstructed momentum distribution for truncation at 20, 40 and 60 turns as well as the input form of the momentum distribution, $Phi(\delta)$

.

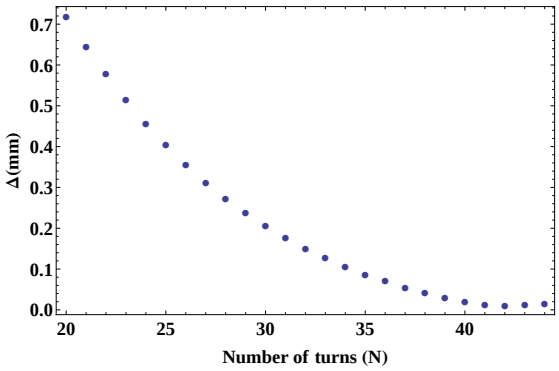

FIGURE 4.46: The mean absolute error (Δ) vs. total number of turns used in reconstructing the momentum distribution.

For each of the truncation settings, the reconstructed momentum distribution was used to reconstruct the simulated 60 turn BPM signal (through fitting the amplitude part of Eq. 4.13).The mean absolute error is then calculated by finding the mean absolute difference between the simulated and reconstructed signal over the first 20 turns, which is taken as a measure of the accuracy of the momentum reconstruction. Figure 4.46 shows that in this case, the mean absolute error falls at a rate proportional to $\frac{1}{N}$.

As a further example, and to show that a non-symmetric momentum distribution may be reconstructed, the simulation was re-run with the particles in the bunch this time having a double peaked Gaussian momentum distribution. The exact form of the momentum probability distribution in this case was:

$$\Phi(\delta) = \frac{0.3}{0.007\sqrt{2\pi}} e^{\frac{-\delta^2}{2(0.0007^2)}} + \frac{0.7}{0.0015\sqrt{2\pi}} e^{\frac{-(\delta-0.001)^2}{2(0.0015^2)}}.$$

For this momentum distribution, the simulated BPM signal converged to approximately zero after around 90 turns. Calculating the Courant-Snyder parameters by using the simulated BPM signal gives $\alpha_z = -0.789$ and $\beta_z = 0.399\,\mathrm{m}$, which deviate by less than 1% from the values obtained through the tracking of a single particle at the mean

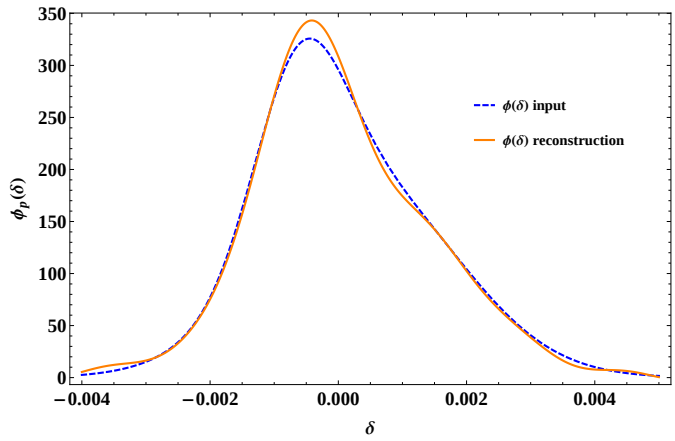

FIGURE 4.47: Reconstruction of a non-symmetrical momentum distribution.

particle momentum ($\alpha_z = -0.794$ and $\beta_z = 0.398\,\mathrm{m}$). In this case, the greater differences between the methods of Courant-Snyder parameter calculation may, in part, be explained by considering the calculation method of the mean betatron tune. For the single peak Gaussian, the mean transverse oscillation frequency coincides with the peak frequency, for the double peak Gaussian, this is no longer the case. This results in a small error being introduced when the peak frequency is used in satisfying Eq. 4.23.

Applying Eq. 4.19 to the full simulated BPM signal produces an accurate reconstruction of the input momentum distribution (Fig. 4.47).

4.5.4 Momentum distribution measurement in EMMA

Within the EMMA lattice there are seventeen pairs of BPMs which are separated only by drift spaces and a vertical corrector magnet. Assuming that the vertical corrector produces only a dipole field, then these pairs of BPMs are appropriate for reconstructing the momentum distribution given that both z_n and z'_n can be obtained by:

$$z'_n \approx \frac{z_n^{(2)} - z_n^{(1)}}{L},$$

where $z_n^{(1)}$ and $z_n^{(2)}$ are the transverse vertical bunch centroid coordinates at BPMs 1 and 2 (respectively) measured with respect to the closed orbit. L is the length of the drift separating BPMs 1 and 2. In practice, the closed orbit is found by averaging the co-ordinate measured at each BPM over many turns. The fractional part of the betatron tune, Q_z, was calculated by applying the NAFF correlator to the signal data from individual BPMs. A value for the chromaticity was obtained using the polynomial defined in section 4.4.3.

Out of the seventeen BPM pairs identified as appropriate, three were selected for reconstructing the momentum distribution. Using more than three of the seventeen pairs at once was not possible due to hardware limitations and the need to use some BPMs for other purposes. The momentum distribution reconstruction method was applied

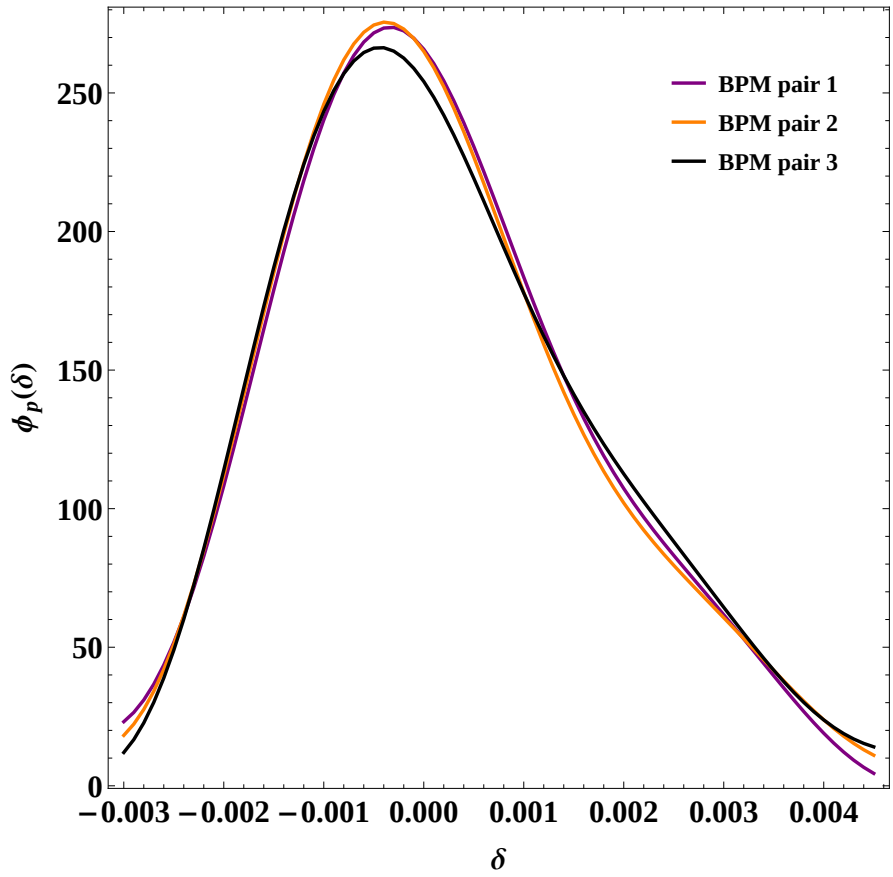

FIGURE 4.48: Reconstruction of the momentum distribution using data from three BPM pairs.

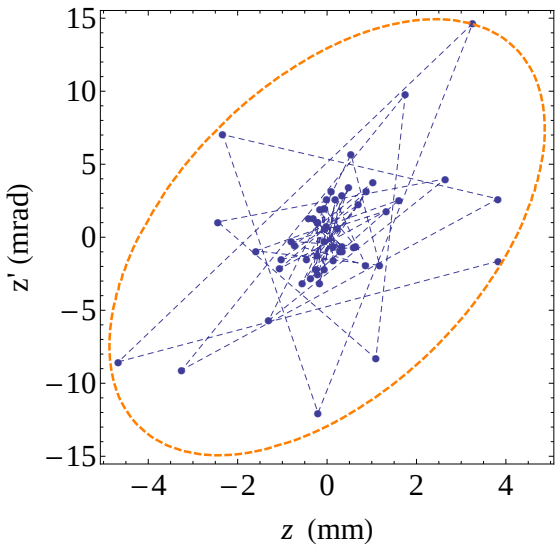

FIGURE 4.49: Blue points represent the measured phase space coordinates, with the dashed blue line included to show the turn by turn progression. The orange dashed line is the phase space ellipse which is given by the fitted action and Courant-Snyder parameters.

TABLE 4.3: Reconstruction of α and β from experimental data.

BPM pair	α	β (m)
1	-0.55 ± 0.04	0.37 ± 0.02
2	-0.70 ± 0.07	0.33 ± 0.02
3	-0.70 ± 0.04	0.37 ± 0.02

independently to the data obtained by each of the three BPM pairs selected, giving three independent momentum distribution measurements per single shot (from injection up until the 60^{th} turn).

For thirty consecutive shots, the Courant-Snyder parameters were calculated at the position of the first BPM in each of the three pairs, giving the results shown in table 4.3. Using the calculated Courant-Snyder parameters, the momentum distribution was then reconstructed. Figure 4.48 shows that there was good agreement of momentum distribution reconstruction between each of the three BPM pairs. Figure 4.49 gives an example of the phase space coordinates calculated for one of the BPMs by using measurements from a BPM pair, as well as showing the phase space ellipse which is drawn using the fitted Courant-Snyder parameters and action. The damping of the measured signal, seen in Fig. 4.49, demonstrates the rapid decoherence of betatron motion of the bunch.

Through use of the same method as for the simulated data, the rate at which the reconstructed momentum distribution converges to the true momentum distribution with increasing turns of BPM data is checked. In the case of experimental data, we see that as the number of BPM measurements used initially increases, the reconstructed momentum appears to converge towards the real momentum distribution (Fig. 4.50). However, as the number of BPM measurements further increases, the mean absolute error also increases, which suggests that the agreement between the real and reconstructed momentum distribution worsens. One possible explanation for this is transient beam loading. During data taking, the revolution frequency of a particle bunch in EMMA was close to a harmonic of the rf cavity resonant frequency. Under such circumstances, the momentum distribution of the particle bunch can change significantly as energy is lost to the rf cavities (this is discussed further in section 4.5.6). Evidence of significant beam loading was observed in the BPM data for the horizontal axis; Taking the mean position of the BPM measurement over a number of turns gives the closed orbit position. When a sliding rectangular window of width 10 turns was used to calculate the change in closed orbit with time, the horizontal closed orbit position was seen to shift towards the centre of the ring with increasing turn number. Although further investigation is required for a proper quantitative understanding, the drift in closed orbit is consistent with the estimated energy loss from beam loading, given the dispersion at the location of the BPM.

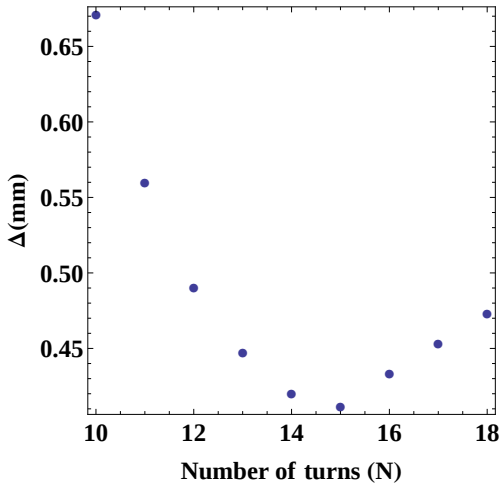

FIGURE 4.50: Convergence of fit betatron oscillation to experimental data.

4.5.5 Synchrotron radiation

So far, it has been assumed that the central momentum of a particle bunch will remain constant throughout the period for which the bunch is tracked by BPMs. In practice, synchrotron radiation and transient beam loading will effect the momentum of the bunch. In this section we consider synchrotron radiation, whilst in following section (section 4.5.6) we investigate the effect of transient beam loading.

The instantaneous synchrotron radiation power for a single particle is given by:

$$P_\gamma = \frac{C_\gamma}{2\pi}\beta_0^2 c\frac{E_0^4}{\rho^2}, \tag{4.24}$$

where, for the particle being accelerated, γ is the Lorentz factor, β_0 and E_0 (MeV) are, respectively, the initial relativistic speed and energy, and ρ is the bending radius [36]. C_γ is a constant, $C_\gamma = 8.846 \times 10^{-14}\,\mathrm{m/MeV^3}$. The energy lost by a particle during one turn in an accelerator is therefore:

$$\Delta E = \oint P_\gamma dt = \oint P_\gamma \frac{ds}{\beta_0 c}. \tag{4.25}$$

If we neglect magnet alignment errors, the bending radius of a particle in EMMA at a distance, s, along its closed orbit is given by:

$$\rho(s) = \frac{p}{q|B_z(s)|}.$$

To achieve alternating gradient focusing, a particle beam travelling through an EMMA cell is first bent with, and then secondly against, the curvature of the ring. Thus, the mean bending radius of a particle is significantly less than the radius of the ring. We find the mean bending radius for 12, 16 and 20.3 MeV/c particles by tracking particles through the field map model, integrating the magnetic fields along the closed orbit and

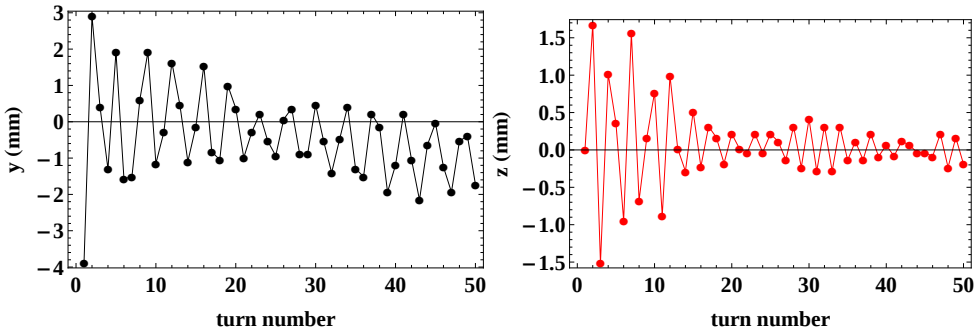

(a) Horizontal (b) Vertical

FIGURE 4.51: Experimental BPM measurements taken for a single shot at an equivalent momentum of 17.6 MeV/c. The horizontal closed orbit (which is approximately the mean BPM position over a number of turns) is observed to shift towards the machine centre with advancing turn number, whilst no shift in closed orbit along the vertical axis is seen. This is behaviour is as expected for a beam losing energy to the cavities through transient beam loading when the revolution period of the bunch is some harmonic of the rf frequency.

then dividing this result by the path length of the closed orbit, s_{co}:

$$\overline{\rho}_p = \frac{p}{q s_{co}} \oint \frac{1}{|B_z(s)|} ds. \tag{4.26}$$

Finally, we combine Eqs. 4.24, 4.25 and 4.26 to estimate the energy lost through synchrotron radiation per particle per turn:

$$\Delta E = \frac{C_\gamma}{2\pi} \beta_0 \frac{E_0^4}{\overline{\rho}_p^2} s_{co}.$$

We find that the energy lost per particle per turn is, respectively, 8.28, 10.4 and 5.98 meV for the 12, 16 and 20.3 MeV/c particles. In the next section, we will see that the changes in particle energy due to synchrotron radiation are negligible compared to the changes due to transient beam loading.

4.5.6 Transient beam loading

Within EMMA there are 19 rf cavities, which are tuned to a resonant frequency of 1.301 GHz. Each time a bunch traverses a cavity it will induce a voltage within the cavity, furthermore, a bunch may encounter the fields induced during previous traversals. Figure 4.51 shows experimentally measured BPM data, which was taken when the beam revolution frequency was close to a harmonic of the rf frequency (at an equivalent momentum of 17.6 MeV/c). Evidence of beam loading seen within the figure includes a shift in the horizontal closed orbit.

A simple modelling technique [83] is applied in order to make an estimate of how beam loading may affect the momentum of particles within a bunch as they circulate within EMMA. We start with a particle bunch, represented by a point charge, being injected into the ring at a time when there is no energy stored in the rf cavities. As the

bunch passes a cavity for the first time, the cavity voltage changes from 0 to V_{b1}, and the induced voltage influences the energy of the particles within the bunch. The extent of this influence can be calculated after considering a simple thought experiment which involves two bunches (of equal charge, q) passing through a cavity at different times:

1. The first bunch passes through the cavity, bringing the cavity voltage from 0 to V_{b1}. The induced voltage, which will be at its most decelerating phase, will act upon the bunch particles as they pass through the cavity. As the amplitude of the cavity voltage changes during the bunches passing, the voltage acting upon the particle is considered to be some fraction, f, of V_{b1}, so that the energy of the first bunch changes by:
$$\Delta E_1 = qfV_{b1} \tag{4.27}$$

2. A second bunch passes through the cavity following a phase advance of π radians in the rf voltage. The phase of the existing cavity voltage (from the first bunch) would now give maximum acceleration for the passing bunch, however, we must also superpose the induced voltage from the second bunch. This time the bunch will see the entire voltage that was induced by the first bunch, and a fraction of the voltage, V_{b2}, induced by itself. The change in energy of the second bunch is therefore:
$$\Delta E_2 = q(fV_{b2} - V_{b1}). \tag{4.28}$$
Given the equal charge of the bunches, the voltage induced by the second bunch will be equal to that induced by the first bunch. As the voltage induced by the first and second bunches add in antiphase, no energy is left stored in the cavity once the second bunch has passed. The change in energy of the second bunch is rewritten:
$$\Delta E_2 = q(fV_{b1} - V_{b1}) = qV_{b1}(f - 1). \tag{4.29}$$
With no energy remaining in the cavity, the conservation of energy means that the energy lost by the first bunch must equal that gained by the second bunch. Equating the energy lost by the first bunch to that gained by the second bunch ($\Delta E_1 = -\Delta E_2$) gives:
$$qfVb_1 = qV_{b1}(1 - f), \tag{4.30}$$
meaning that the fraction of the voltage induced on a single pass of the cavity that acts upon the bunch is $f = 1/2$, and that the energy lost by a bunch to an empty cavity is:
$$\Delta E = \frac{1}{2}qV_b. \tag{4.31}$$

We now introduce two new factors, the shunt impedance, R_s, and the unloaded cavity quality factor, Q_0. The shunt impedance relates the voltage across a cavity to the power

dissipated in the cavity walls (P_c):

$$R_s = \frac{1}{2}\frac{|V_0|^2}{P_c}, \tag{4.32}$$

whilst the cavity quality factor is 2π multiplied by the number of rf cycles it takes to dissipate the energy stored in a cavity (E_c):

$$Q_0 = \frac{\omega_{\mathrm{rf}} E_c}{P_c}. \tag{4.33}$$

Both the shunt impedance and the quality factor have been measured for the EMMA cavities, and are given in table 4.4. The ratio of the shunt impedance to the quality factor, which is given by:

$$\frac{R_s}{Q_0} = \frac{|V_0|^2}{2\omega_{\mathrm{rf}} E_c}, \tag{4.34}$$

can then be used to find the voltage induced across the cavity. Through the conservation of energy, the energy stored within an initially empty cavity must equal the energy lost by a bunch as it passes the cavity. As the charge of the bunch is negative, the induced voltage must be positive, giving:

$$\frac{R_s}{Q_0} = -\frac{V_b^2}{\omega_{\mathrm{rf}} q V_b},$$

and therefore:

$$V_b = -\omega_{\mathrm{rf}} q \frac{R_s}{Q_0}. \tag{4.35}$$

We now want to calculate the effect of the bunch crossing the cavity multiple times. Assuming that the energy lost by a bunch remains stored in the rf cavity for some time, the voltage induced as a bunch makes a first pass of the cavity will oscillate sinusoidally at the rf frequency. When a bunch arrives at the cavity a second time, it will again induce the voltage given by Eq. 4.35, but will also experience the voltage induced during the previous traversal. The voltage seen by a particle passing through a cavity is dependent upon the phase of an electromagnetic field when the particle arrives at the cavity (it is assumed that the change in phase of the rf as a particle crosses the cavity is negligible). The energy lost by the particle to the cavity immediately effects the electric field within the cavity, but not the magnetic field. For this reason it is convenient to write the cavity voltage as a phasor:

$$V(t) = V_0 e^{i\omega_{\mathrm{rf}} t}, \tag{4.36}$$

where the real and imaginary components of Eq. 4.36 at a given time are representative of the electric and magnetic fields, respectively. Each time the bunch encounters the cavity, the induced voltage (Eq. 4.35) is added to the real part of Eq. 4.36. Furthermore, the Ohmic losses of energy from the cavity with time are accounted for. The decay of

stored energy after a time, t, is given by:

$$E_{\mathrm{c}}(t) = E_{\mathrm{c},0} e^{\frac{-\omega_{\mathrm{rf}} t}{Q_0}}, \tag{4.37}$$

and as $V_0 \propto \sqrt{E_{\mathrm{c}}}$, then:

$$V(t) = V_0 e^{\frac{-\omega_{\mathrm{rf}} t}{2Q_0}}. \tag{4.38}$$

Therefore, the voltage experienced by a particle on the n^{th} encounter of a particular cavity at a time T is:

$$V_n(T) = V_b \left(\frac{1}{2} + e^{i\omega_{\mathrm{rf}} \Delta t_{n-1}} e^{\frac{-\omega_{\mathrm{rf}} \Delta t_{n-1}}{2Q}} + ... + e^{i\omega_{\mathrm{rf}} \Delta t_1} e^{\frac{-\omega_{\mathrm{rf}} \Delta t_1}{2Q}} \right), \tag{4.39}$$

where $\Delta t_n = T - t_n$. The change in energy of a bunch is:

$$\Delta E_b = q\Re(V_n(T)) \tag{4.40}$$

The methods for the modelling technique can now be summarised as follows:

1. EMMA has 19 rf cavities, which are located in every second cell, with the exception of the injection and extraction cells (the 2^{nd} and 26^{th} cells respectively). We use the experimentally measured orbital period vs. momentum curve that has not been corrected for speed; this is as we want to estimate the time taken for an infinitesimally short bunch to travel between cavities when using equivalent momenta.

2. The bunch begins at the first cavity within the lattice. Every time a bunch passes a cavity, it induces a voltage in the cavity which is given by Eq. 4.35. The voltage experienced by the bunch at the first cavity is given by Eq. 4.39, and the bunch energy changes according to Eq. 4.40. The updated particle momentum determines the travel time of the bunch to the second cavity.

3. The voltage across each cavity is independent of the voltage across the other cavities. The bunch induces and sees the same voltages across the second cavity as was the case for the first cavity (i.e. there is no superposition of the voltages induced in different cavities). Again, the time for the bunch to reach the following cavity is based upon the new particle momentum. The process is continued for the remaining cavities in the ring, with the energy lost by the bunch being the same at all cavities during the first revolution in the ring.

4. The bunch returns to the first cavity after some time that is dependent on the varying bunch momentum during the first revolution. The voltage induced during the second pass of the cavity adds to the decaying voltage induced during the first pass. The net voltage once the bunch has left the cavity may be greater or less than the voltage across the cavity before the bunch entered depending upon the

TABLE 4.4: EMMA rf cavity.

EMMA cavity parameters	
Number of cavities	19
Length	110 mm
Resonant frequency	1.301 GHz
Voltage	120 − 180 kV
Quality factor (Q)	23000
Shunt impedance (R_s)	3.4 MΩ

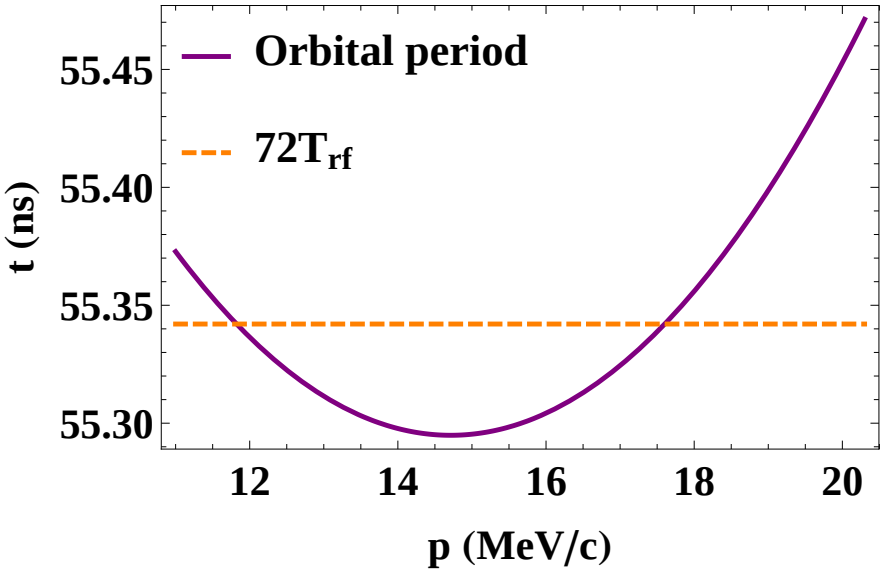

FIGURE 4.52: Plot to show the measured orbital period curve (without speed correction) and the time period for the 72$^{\text{nd}}$ harmonic of the rf frequency ($72T_{\text{rf}}$). The two lines intersect at 11.82 MeV/c and 17.60 MeV/c, and the effects of beam loading can be expected to be greatest at these momenta when using the EMMA equivalent momenta settings.

phase advance of the initial voltage. The process is repeated for the remaining cavities and then for the desired number of turns.

Two specific cases are considered, firstly when the initial particle momentum is 17.35 MeV/c and secondly for an initial momentum of 18.35 MeV/c. When the initial particle momentum is set to 15.85 MeV/c, the resonant frequency of the cavity can be described by $\omega_{rf} \approx 71.95\omega_b$, where ω_b is the revolution frequency of the bunch. In this instance, the phase of the induced rf voltages at which the particle bunch arrives at the cavities changes quickly, and as a result synchrotron oscillations are performed through a small range of δ (Fig. 4.53a).

For an initial momentum of 17.85 MeV/c, the particle revolution frequency is close to a harmonic of the rf resonant frequency, $\omega_{rf} \approx 72.01\omega_b$. This time the bunch remains within the decelerating phase of the induced rf voltages for approximately 50 of the 60 turns of tracking, and there is a more significant decrease in δ (Fig. 4.53b).

Finally, we investigate the effect of transient beam loading on the propagation of the transverse phase space variables. The cavity model within Zgoubi does not have

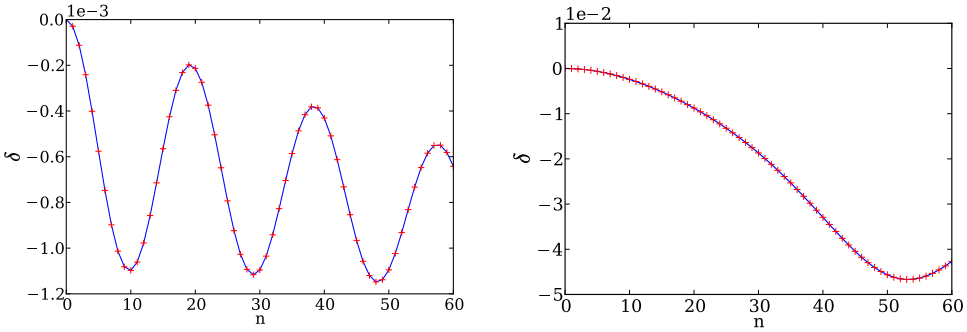

(a) $p_{\text{equiv}} = 15.85\,\text{MeV/c}$ (b) $p_{\text{equiv}} = 17.85\,\text{MeV/c}$

FIGURE 4.53: Simulated effects of beam loading upon particle momentum with the EMMA cavity parameters given in table 4.4 and injection at equivalent momenta of 15.85 MeV/c (a) and 17.85 MeV/c (b). When injecting at 17.85 MeV/c, the bunch orbital frequency is given by $\omega = 71.95\omega_{\text{rf}}$ and δ after 60 turns has a magnitude no greater than 10^{-3}. If the injection momentum is 17.85 MeV/c, then $\omega = 72.01\omega_{\text{rf}}$ and we see δ grow to the order of 10^{-2}.

the features necessary for this study, so we instead use a python script that updates the momentum of particles within the bunch at each cavity according to the beam loading model developed, and then uses Zgoubi to track the transverse dynamical variables between cavities. A bunch of 4000 particles is tracked through 80 turns of the hard edge model: at the start of tracking, the mean bunch momentum is 19.02 MeV/c (with the spread of particles around the mean described by a Gaussian with a standard deviation of $\delta = 0.001$) and the bunch centroid is offset from the closed orbit by 1 mm, we consider an rf frequency of 1.299 GHz, which gives $\omega_{\text{b}} = 72.02\omega_{\text{rf}}$ for this model. The approximation of representing the bunch by a point charge in the beam loading model means that although the mean momentum of the bunch will change, the momentum distribution around the mean is constant. A plot of the bunch centroid vs. turn number for the first 50 turns is given in Fig. 4.54, where we see that the closed orbit along the horizontal axis shifts by approximately −0.5 mm. Next we apply a sliding rectangular window that has a width of 16 turns to the simulated BPM data for the vertical axis; as we move the window along the BPM data, we calculate the tune for each sample (i.e. for turns 1 to 16, 2 to 17, etc.) and the calculated tune is used as an estimate for the tune at the midpoint of the sample (meaning that we start with a tune calculation for the 8^{th} turn). In Fig. 4.55, we plot the tunes calculated through NAFF alongside the tune calculated using the known values for the mean momentum of the bunch and chromaticity of the lattice; we see that the NAFF method for tune calculation offers a reasonable estimate for the tune shift of the bunch. On the basis of this result, we assume that the NAFF correlator could be applied to experimental data to give an estimate of the change in the mean bunch momentum for cases where the chromaticity is known.

If we use the BPM data to reconstruct the momentum distribution, and then fit the remaining parameters of the model for the betatron oscillation to the BPM data, we see that the mean absolute error (Fig. 4.56) behaves in the same fashion as for

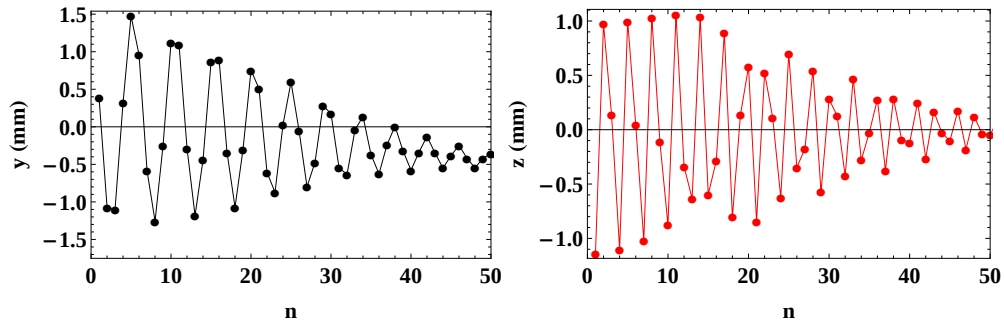

(a) Horizontal (b) Vertical

FIGURE 4.54: Simulated bunch centroid position vs. turn number when transient beam loading is taken into account. We track 4000 particles through the design hard edge representation of EMMA, the bunch momentum initially has a Gaussian distribution with standard deviation of $\delta = 0.001$ around the mean momentum of 19.02 MeV/c. An rf frequency that gives $\omega_b = 72.02\omega_{rf}$ is selected. We see that the closed orbit shifts by approximately -0.5 mm during tracking.

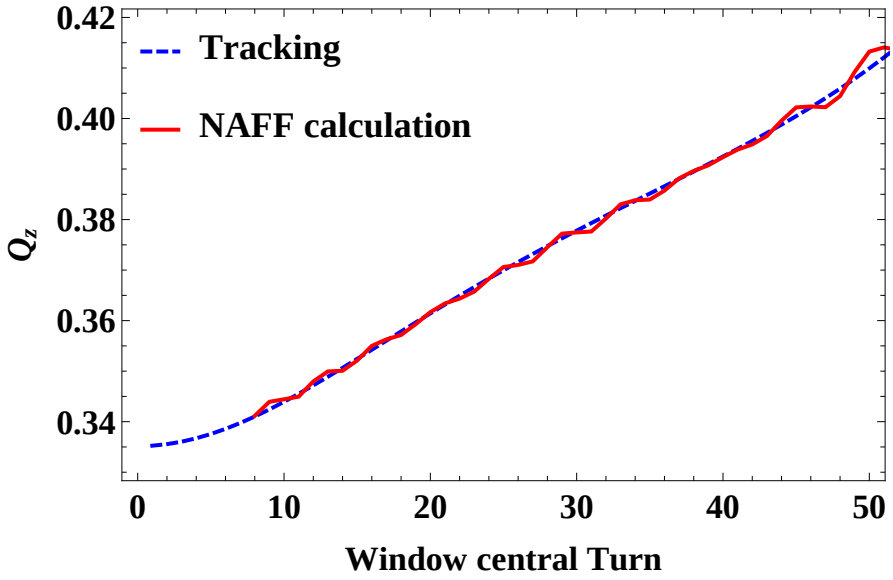

FIGURE 4.55: Plot to show the tune calculated when the NAFF correlator is applied to a subset (16 turns long) of the simulated BPM data (Fig. 4.54b) vs. the turn number at the centre of the subset (red line). We also show the change in tune calculated using the change in bunch momentum during tracking and the lattice chromaticity (blue dashed line); it is noted that the line based on the NAFF correlator is in good agreement with the line based on known values for the beam momentum and lattice chromaticity. The NAFF correlator has been applied to experimental data in order to estimate the impact of beam loading on the bunch momentum.

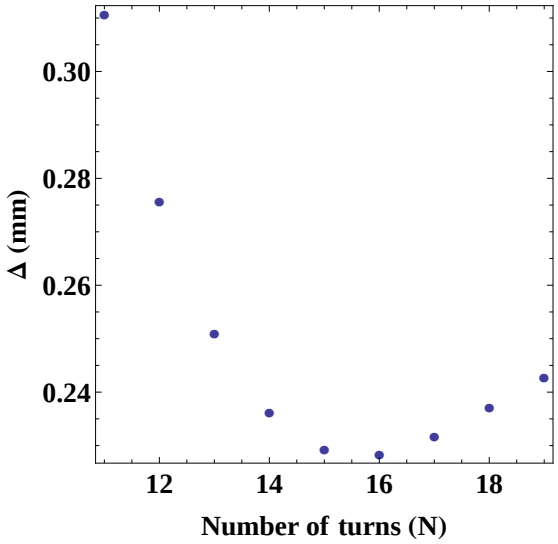

FIGURE 4.56: Convergence of fit betatron oscillation to data from a simulation that includes the model for beam loading. The behaviour of the convergence when applying the reconstruction methods to this simulation data is similar to that found for experimental data (Fig. 4.50).

the experimental data (Fig. 4.50). In order to reconstruct the momentum distribution accurately, it is important to minimise effects such as beam loading. Ensuring that the particle revolution frequency is not a harmonic of the rf resonant frequency, through either careful selection of the injection momentum or by detuning the rf cavities, offers a partial solution to this problem.

4.5.7 Conclusions from lattice parameter reconstruction

A method for calculating the Courant-Snyder parameters and reconstructing the bunch momentum distribution in non-zero chromaticity machines has been presented.
When applied to data produced through simulation, the methods were shown to accurately calculate the Courant-Snyder parameters, reconstruct the input momentum distribution and then reconstruct the BPM signal of the bunch.
When applied to BPM data collected with EMMA an estimate of the momentum distribution can be found. However, the reconstructed momentum distribution does not converge to a true momentum distribution when increasing turns of BPM data are used (as was demonstrated with simulated data). Transient beam loading and the bunch revolution frequency being a harmonic of the rf cavity resonant frequency offer a reason as to why the reconstructed momentum distribution may not converge as expected.
In future, the concept could be tested more rigorously, on EMMA or another non-zero chromaticity machine, by taking more data after first ensuring that the rf cavities have been detuned.

4.6 Amplitude dependent orbital period and tune

Non-scaling FFAGs are being considered for accelerating particle (e.g. muon) beams with a large emittance, however, early tracking studies demonstrated that a large transverse emittance leads to a blow up of the longitudinal emittance in machines such as EMMA [84]. The cause of this blow up is the increase in path length that is experienced by a particle of greater transverse amplitude and the subsequent increase in orbital period per turn. The amplitude dependence of orbital period means that the particles within a bunch experience different rates of phase slippage with the accelerating rf, which goes unrectified due to the lack of longitudinal focusing in the serpentine channel. For a proton therapy machine, the transverse emittance of a bunch during acceleration should be small enough for the amplitude dependence of orbital period not to impact upon the performance of the machine significantly, however, experimental results on this subject are presented as a further test of computer modelling. The orbital period dependence on amplitude has been found to be linked to the chromaticity of the lattice, and to first order is given by:

$$\Delta t = \frac{2\pi}{c}(\xi_y J_y + \xi_z J_z). \tag{4.41}$$

4.6.1 Experimental method and results

The relationship between orbital period and transverse amplitude was investigated experimentally whilst EMMA was set up for an equivalent momentum of 17.6 MeV/c. The transverse emittance of the injected bunch is expected to be of a size similar to that previously measured (0.71 mm mrad), and so smaller than the EMMA design acceptance of 3000 mm mrad, allowing for the bunch centroid to be injected at some offset from the closed orbit without a significant loss of particles. The transverse oscillation amplitudes for each axis were varied independently, with the horizontal amplitude controlled by adjusting the injection septum strength and the vertical amplitude by using a corrector in the EMMA injection line. Prior to collecting data for the horizontal axis the vertical amplitude was minimized, then, during data collection, no further attempt made to affect the vertical amplitude (with the converse being true when taking data for the vertical axis). For each axis a number of strengths for the perturbing magnet were defined, and the data then collected in a random order (i.e. not going low to high in magnet strength) so as to reduce the probability of a orbital period shift due to a gradual slip in the injection momentum being attributed to the change in transverse amplitude. During the initial analysis, it was found that accurately calculating the horizontal amplitude is difficult due beam loading effects, and so emphasis is placed on the results for the vertical axis.

Using the methods of section 4.5, the amplitude of the coherent oscillation (Fig. 4.57a) and beta function were calculated at the location of a reference BPM (Fig. 4.57b) for up to 50 injection cycles per vertical corrector strength (due to the BPMs not triggering reliably for large amplitude oscillations, or beyond a certain number of turns). The

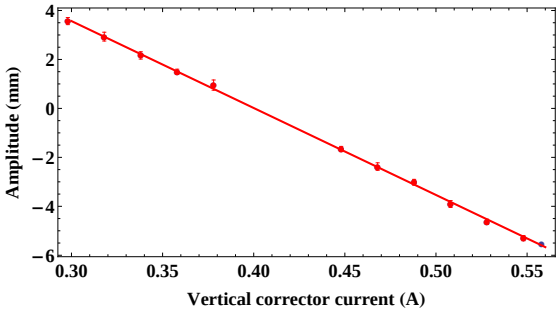

(a) Vertical amplitude vs. vertical corrector current.

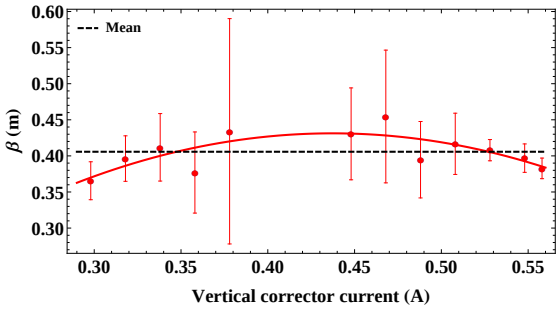

(b) Calculated vertical β function vs. vertical corrector current.

FIGURE 4.57: Plots of vertical betatron oscillation amplitude (a) and calculated vertical β function (b) vs. the current through a vertical corrector magnet that is located in the EMMA injection line. Given the small emittance of the particle bunch provided by ALICE, the vertical corrector can be used to effectively vary the injection amplitude of the coherent betatron oscillation in EMMA. The oscillation amplitude and β function were calculated using the methods described in section 4.5. From these measurements, we take the mean β function across all corrector currents (shown by the black dashed line in (b)) and the oscillation amplitude to calculate the vertical action for each corrector current (Fig. 4.58).

mean oscillation amplitude for each corrector current and the mean β function across all corrector currents ($\bar{\beta}$) were then used to calculate the vertical action for each setting (Fig. 4.58). The orbital period for each vertical corrector setting was calculated using the method described in section 4.4.4, which gives the results shown in Fig. 4.59. A plot of the calculated orbital period vs. vertical action (Fig. 4.60) for three BPM pairs suggests a linear relationship that is in agreement with theory. We see some differences between the relationship of orbital period vs. vertical action that is obtained for the three different BPM pairs. This can possibly be explained by an assumption made during reconstruction of the betatron motion, which is that decoherence begins the first time a particle bunch passes a BPM pair. In reality, decoherence will begin from the point of injection. Further investigation would be required to verify this idea, however the closer agreement between the relationship obtained using BPM pairs 1 and 2 (which are located in consecutive cells just beyond the injection septum) than between pairs 1 and 3 or 2 and 3 (where pair 3 is located $\sim$ 20 cells later than pair 2) offers some supporting evidence. We also note that the shape of the parabola shown in the

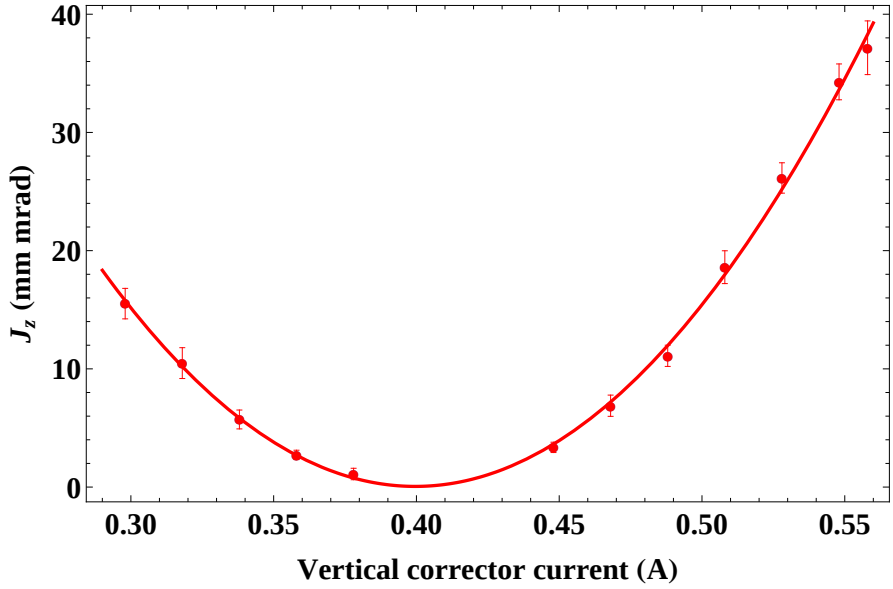

FIGURE 4.58: Vertical action vs. vertical corrector current. The vertical action is calculated as $J_z = A^2/(2\bar{\beta})$, where A and $\bar{\beta}$ are the amplitudes and mean vertical β function shown in Fig. 4.57.

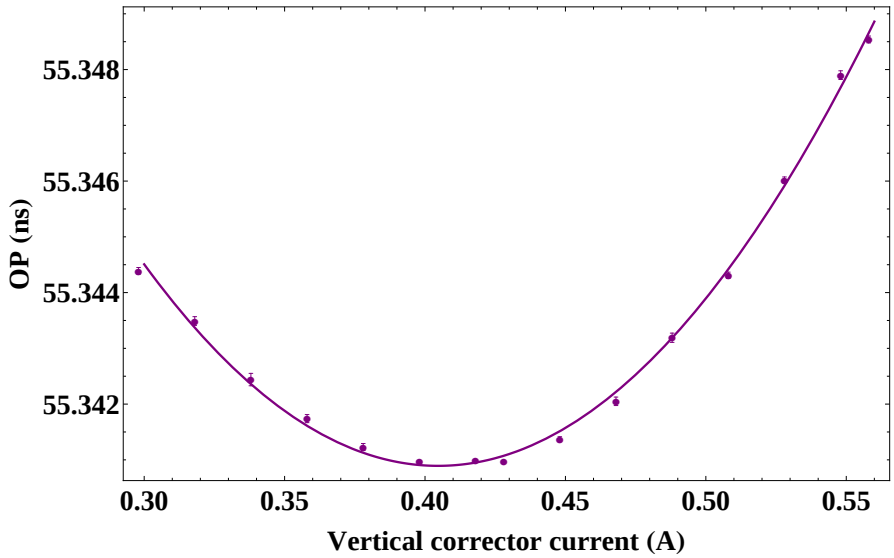

FIGURE 4.59: Orbital period vs. vertical corrector current at an equivalent momentum of 17.6 MeV/c.

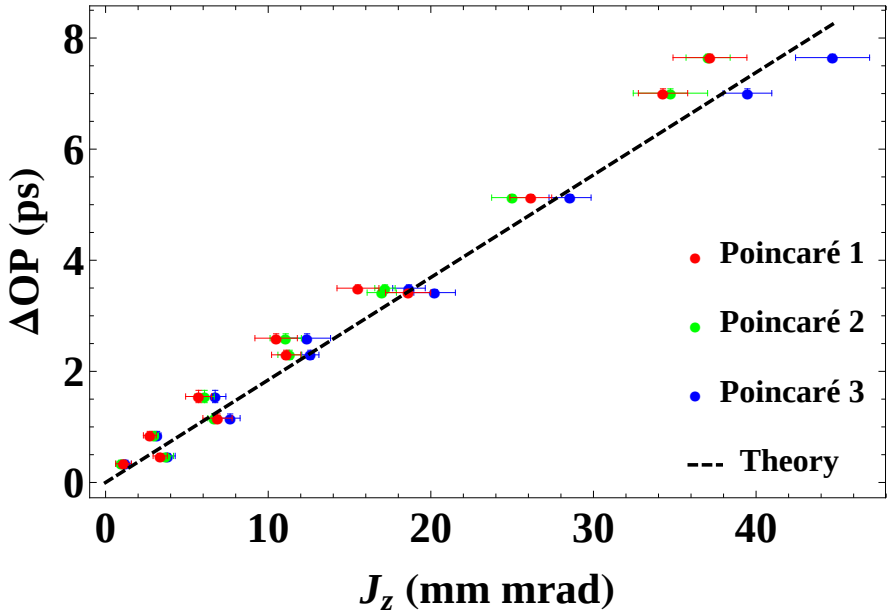

FIGURE 4.60: Change in orbital period vs. vertical action. Here, the chromaticity for the theoretical relationship is taken from the polynomial fit of the tune measured using equivalent momenta.

plot of vertical action vs. corrector current (Fig. 4.58) can still be seen in the plots of orbital period vs. vertical action for individual BPM pairs (Fig. 4.60). This suggests that the vertical corrector setting also has some influence over the horizontal oscillation amplitude of the beam injected into EMMA (Fig. 4.61). Finally, whilst following the methods for calculating the orbital period vs. transverse action using experimental data, the betatron tune shift with amplitude was also obtained. Although this should be a small effect in a machine which is based on linear multipole fields, the use of the NAFF correlator allowed for the tune to be calculated to sufficient precision with the limited number of turns of data available Fig. 4.62. Experimental measurement of the tune shift with amplitude may also be possible due to the contribution of sextupole fields within the fringe field regions (tracking studies found that the rate of amplitude detuning was significantly higher for the field map model than for the hard edge model).

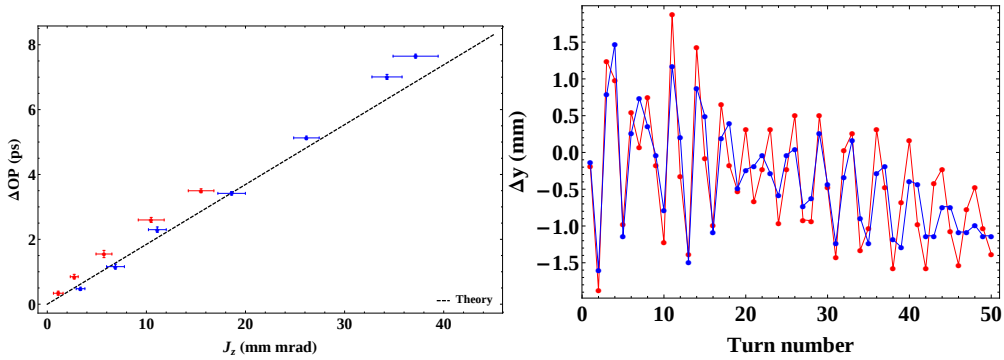

(a) Orbital period dependence on vertical action.

(b) Horizontal offset from closed orbit vs. turn number.

FIGURE 4.61: In Fig. (a) it can be seen that the data points left of the parabola minima of Fig. 4.58 (above red) would have a different line of best fit to those from the right (above blue). The data in the above example is taken from Poincaré 1 (Fig. 4.60), however the same behaviour is seen for all three Poincaré pairs. In Fig. (b), BPM data for the horizontal centroid position is compared for two vertical corrector settings that give a similar vertical action, but at different sides of the parabola minima, these are 0.318 A (red) and 0.488 A (blue). It appears that the horizontal amplitude for the 0.318 A setting is larger than for 0.488 A, suggesting that the vertical injection amplitude affects the horizontal amplitude, and offering a possible explanation for the parabola being visible in Fig. (a).

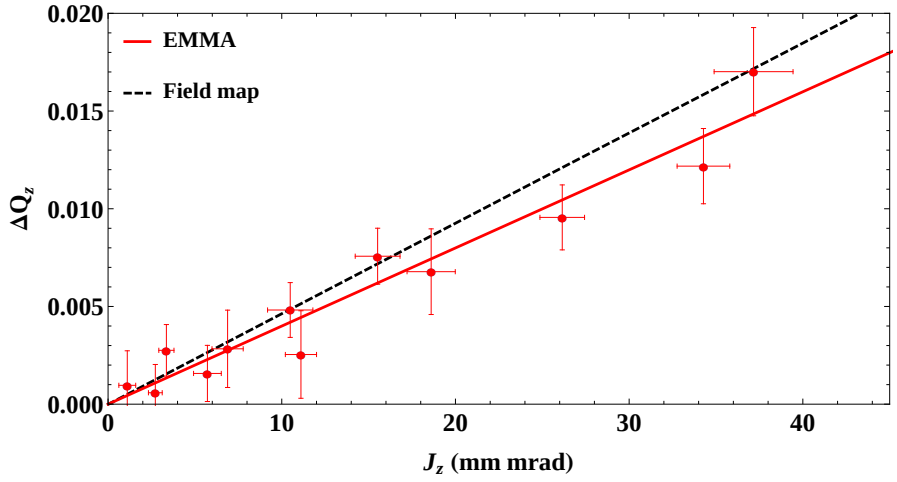

FIGURE 4.62: Change in vertical betatron tune vs. vertical action calculated experimentally and using a field map. The rate of amplitude detuning should be small for the nsFFAG EMMA, however it proved possible to measure this effect by applying the NAFF correlator to the small number of turns of data. We find reasonable agreement between the experimental and field map based calculations.

Chapter 5

Synchrotron design study

In this chapter the performance of a synchrotron for proton therapy will be investigated. The key performance metrics are the energy range and control of energy variation, intensity and emittance of the extracted beam, extraction efficiency and the speed with which the extracted beam can be turned on and off. The extracted beam energy is determined by the magnet strengths, and is not a major technical issue; apart from issues affecting the energy spread of the extracted beam, we do not consider energy issues in depth. Other properties of the extracted beam are closely associated with the technique used to achieve extraction. Possible techniques include kicker-based extraction methods, and methods based on the use of resonances. The latter are of particular interest in proton therapy because they provide a better capability of conforming the radiation dose to the tumour volume.

One example of resonant extraction is the use of sextupoles to limit the region of stable particle motion in phase space. The size and shape of the boundary of the stable region (the separatrix) is determined by a variety of parameters, including the tunes (set by the quadrupole strengths) and the sextupole configuration. Extraction may be achieved by adjusting the quadrupoles to bring the tune closer to a third-integer resonance, thereby shrinking the separatrix and causing particles to cross into the unstable region of phase space. However, in practice this method offers poor control over the extraction rate. As an alternative, the magnet strengths can be fixed, and particles can be pushed across the separatrix using a transverse rf deflecting cavity. This is the technique used for example at the Heidelberg Ion-Beam Therapy Center [26]. Experience shows that this method provides very good intensity control, and the ability to switch the extracted beam off very quickly, which is useful for respiratory gating. This is the technique on which we focus in this chapter, because of the advantages that it provides.

Our goal is to describe the technique of resonant extraction using a transverse rf cavity, and evaluate the capability it provides for control of the extracted beam. This will provide a basis for comparison between synchrotrons and FFAGs for proton therapy. A simple lattice (proposed by Fukumoto in 1989 [85]) is used as the starting point for the investigation. After briefly describing the lattice and presenting the main parameters,

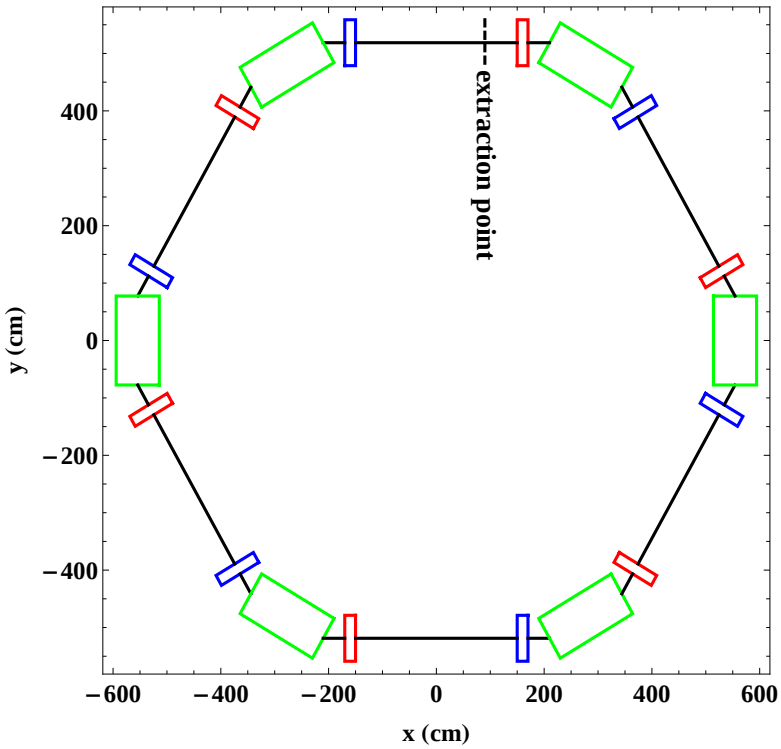

FIGURE 5.1: Layout of the synchrotron lattice. Solid blue and red rectangles mark the defocusing and focusing quadrupoles, respectively, and green rectangles mark the dipole magnets. The dashed black line shows the plane at which extraction will be considered.

we proceed to describe the extraction process in more detail. Simulation results are presented as the theory is developed.

5.1 The synchrotron lattice

The lattice is a six-cell ring of circumference 35 m, with each cell having a DOFB (defocusing–drift–focusing–bend) structure. The footprint with the locations of the dipole and quadrupole magnets is shown in Fig. 5.1. The main parameters of the lattice are given in table 5.1. The β functions and dispersion for one periodic cell are shown in Figs. 5.2 and 5.3 respectively.

During acceleration, the dipole fields ramp linearly at a rate of 2.6 T/s (with the dipole having a peak field of 1.5 T for 230 MeV protons), keeping the equilibrium orbit fixed with changing magnetic rigidity. The change in orbital period of a reference particle during acceleration therefore only depends on particle speed (Fig. 5.4), and, given a harmonic number of 1, the rf period is increased to match the orbital period. The synchronous phase varies from 20° to 30° and the accelerating voltage from 450 V to 300 V within one acceleration cycle in order to keep the height of the rf bucket (when measured in momentum) approximately constant (Fig. 5.5). Important components in the lattice for the extraction process are the sextupoles, the transverse rf deflecting cavity, and components (typically electrostatic and magnetic septa) for capturing particles

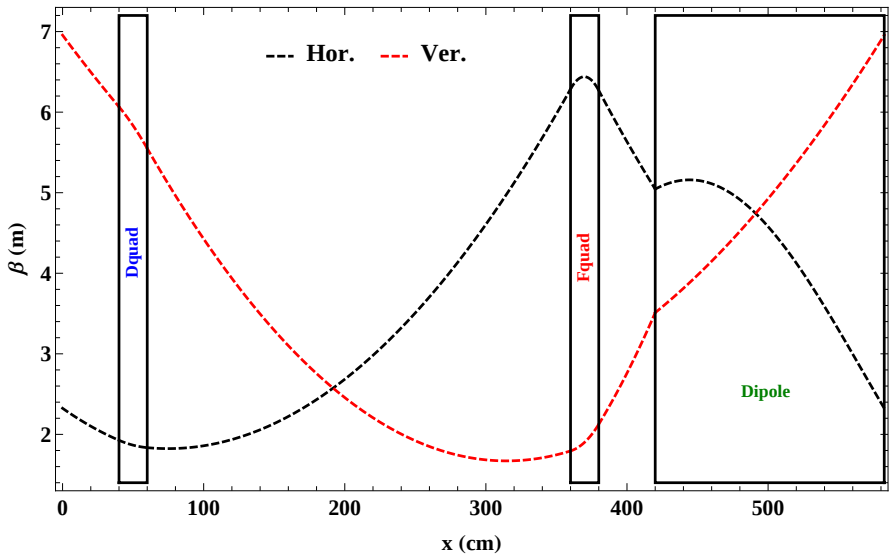

FIGURE 5.2: β functions from the start of the defocusing quadrupole to the end of the focusing quadrupole. Extraction is initially considered at a point in the long drift (60 cm before the focusing quadrupole).

that reach a certain betatron amplitude and directing them into a transport line. The position in the lattice of the first extraction septum is linked with the design of the rest of the system; in this study we consider the septum only as far as its location is concerned. The perturbation sextupoles are located at a point in the long drifts where the horizontal β function is large; for simplicity, the extraction point is initially located at the same relative position in a cell as the perturbation sextupoles. It is assumed that a particle is extracted once its horizontal position exceeds a certain reference value at the chosen location of the septum. Our main focus will be on the effects of the sextupoles and the transverse deflecting cavity that are used to drive particles to large horizontal amplitudes.

Following the acceleration of protons to the required treatment energy, the extraction process can be broken down into three steps:

1. Quadrupoles are used to tune the lattice close to a third-integer horizontal betatron tune.

2. Sextupoles are energised in order to perturb the beam to form a separatrix (of specified size, shape and orientation) in phase space.

3. Particles are pushed across the separatrix using a transverse deflecting cavity.

5.2 Tuning the lattice close to resonance

The size of an area enclosed by a separatrix depends upon both the proximity of the tune working point to the resonant condition and the strength of the sextupole fields. Once the particles have been accelerated to the treatment energy, the quadrupole strengths

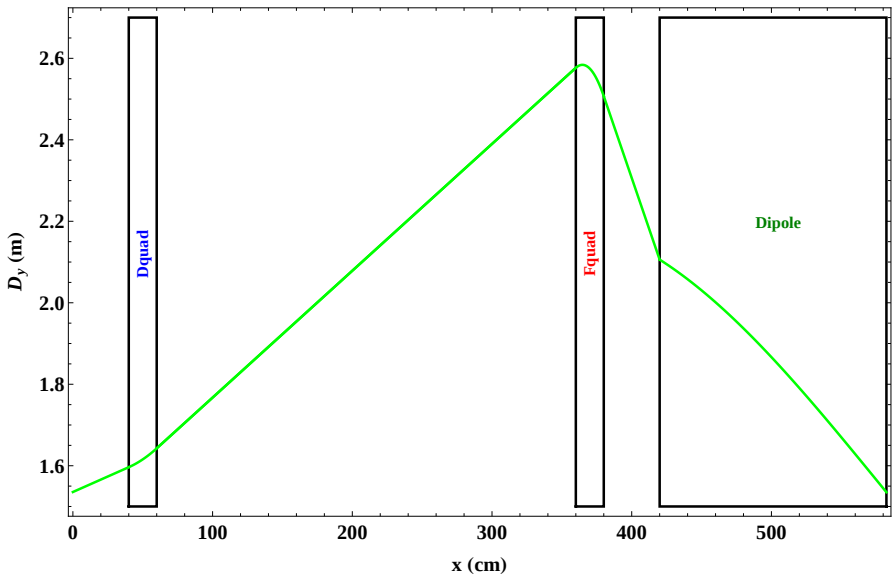

(a) Dispersion.

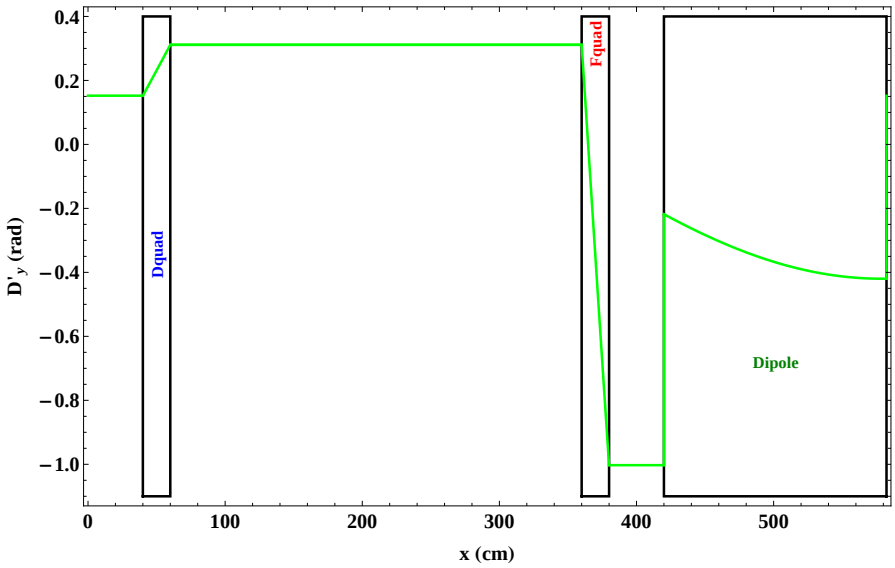

(b) Angular dispersion.

FIGURE 5.3: Plots to show the horizontal dispersion and angular dispersion vs. longitudinal position. For a perfectly aligned lattice, the vertical dispersion is zero.

are altered so that the horizontal tune is close to a third-integer. This was achieved in the tracking study by finding a response matrix, m, which satisfies:

$$\begin{pmatrix} \Delta Q_y \\ \Delta Q_z \end{pmatrix} = \begin{pmatrix} \frac{\partial Q_y}{\partial K_{1,f}} & \frac{\partial Q_y}{\partial K_{1,d}} \\ \frac{\partial Q_z}{\partial K_{1,f}} & \frac{\partial Q_z}{\partial K_{1,d}} \end{pmatrix} \begin{pmatrix} \Delta K_{1,f} \\ \Delta K_{1,d} \end{pmatrix} = m \begin{pmatrix} \Delta K_{1,f} \\ \Delta K_{1,d} \end{pmatrix}.$$

Inverting m and then calculating the product of the inverted matrix with the desired tune shift provides an estimate of the required change in quadrupole strengths,

$$\begin{pmatrix} \Delta K_{1,f} \\ \Delta K_{1,d} \end{pmatrix} = m^{-1} \begin{pmatrix} \Delta Q_y \\ \Delta Q_z \end{pmatrix},$$

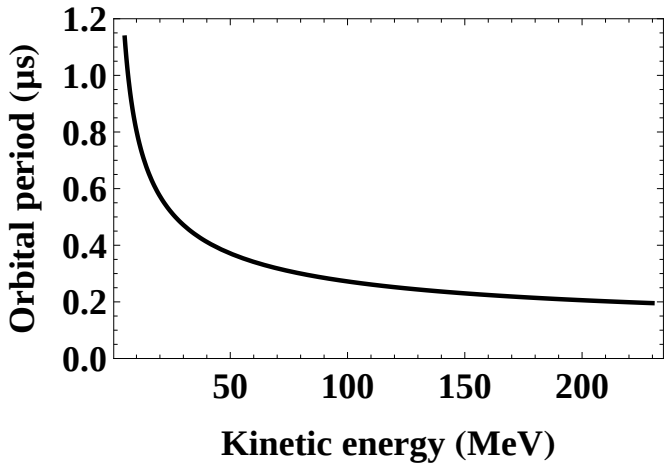

FIGURE 5.4: Orbital period vs. kinetic energy for the proton synchrotron. The path length is fixed with respect to reference energy, and the orbital period of a reference particle depends only on particle speed.

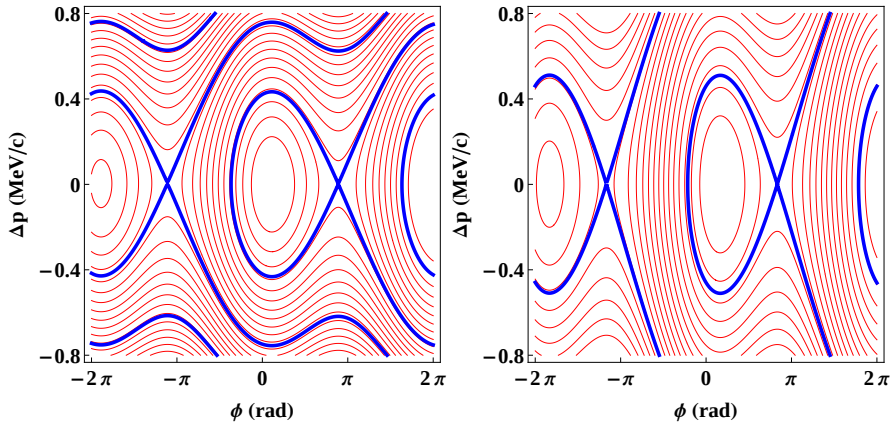

(a) 5 MeV. (b) 230 MeV.

FIGURE 5.5: Longitudinal phase space plots for the proton synchrotron. The shifts in synchronous phase, rf voltage and rf frequency keep the height of the rf bucket approximately constant (at around 1 MeV/c) during acceleration. We note also that the lattice remains below transition for the full momentum range.

where $K_{1,f}$ and $K_{1,d}$ are the strengths of the focusing and defocusing quadrupoles, respectively, and Q_y and Q_z are the horizontal and vertical tunes. The estimate is then improved by carrying out several iterations of the procedure; with the desired tunes in this case being achieved to 3 decimal places after 3 iterations. The horizontal betatron tune is moved to 1.672, whilst the vertical tune is kept at 1.850. The final quadrupole strengths are given in table 5.2.

5.3 Introduction of sextupole perturbations

Sextupole magnets are routinely used to correct the chromaticity of a lattice. In this study a method for perturbing the beam with sextupoles (and forming a separatrix)

TABLE 5.1: Basic parameters of the proton synchrotron model.

Lattice	
Circumference	34.939 m
Periods	6
Lattice type	DOFB
Bend radius	1.55 m
Bending magnet length (measured along reference path)	1.623 m
Dipole field strength	0.2087 m^{-1}
Quadrupole lengths	20 cm
Quadrupole strengths (H/V)	2.562/-0.492 m^{-2}
Dynamical properties	
Normalised rms emittance	1.5π mm mrad
Injection energy	5 MeV
Injection momentum spread (δ)	$\pm$ 0.001
Extraction energy	60-230 MeV
Tune (Q_y/Q_z)	1.800/1.850

whilst not affecting the chromaticity of the lattice is followed; the same method is applied at the Heavy Ion Medical Accelerator in Chiba (HIMAC) [86]. An additional advantage of the method used is that it allows control over the orientation of the separatrix in phase space. In this section, we present simulation results to show the effect of sextupoles on the beam motion, and develop the theory to explain the observed behaviour. The theory also provides a basis for optimisation of the extraction efficiency.

5.3.1 Chromatic and geometric effects of a sextupole pair

Initially, a single sextupole is included within the lattice. As the magnitude of the perturbation to a beam is dependent upon the β function at the location of the sextupole, the sextupole is placed at a point where the horizontal and vertical β functions are large and small, respectively, (sx1 in Fig. 5.6). Figure 5.7a shows how the horizontal chromaticity of the lattice depends on the sextupole strength. Now a second sextupole (sx2 in Fig. 5.6), this time of arbitrary fixed strength, is included in the lattice at a position separated from sx1 by a horizontal phase advance of πQ_y. The periodicity of the lattice ensures that the β functions at the locations of sx1 and sx2 are equal. It is now seen that when the strength of sx1 is equal and opposite to that of sx2, the chromaticity of the lattice returns to that obtained when no sextupole fields are present (Fig. 5.7b). Despite the chromatic effect of the sextupoles cancelling, the geometric perturbation remains, and a separatrix is formed in phase space (Fig. 5.8).

5.3.2 Control of a separatrix through multiple sextupole pairs

Now a second sextupole pair is included. The sextupoles in the original pair (sx1 and sx2) are fixed to being equal and opposite in strength (and are relabelled as dsx1 and fsx1 in Fig. 5.12 respectively), the sextupoles in the second pair (dsx2 and fsx2) are also

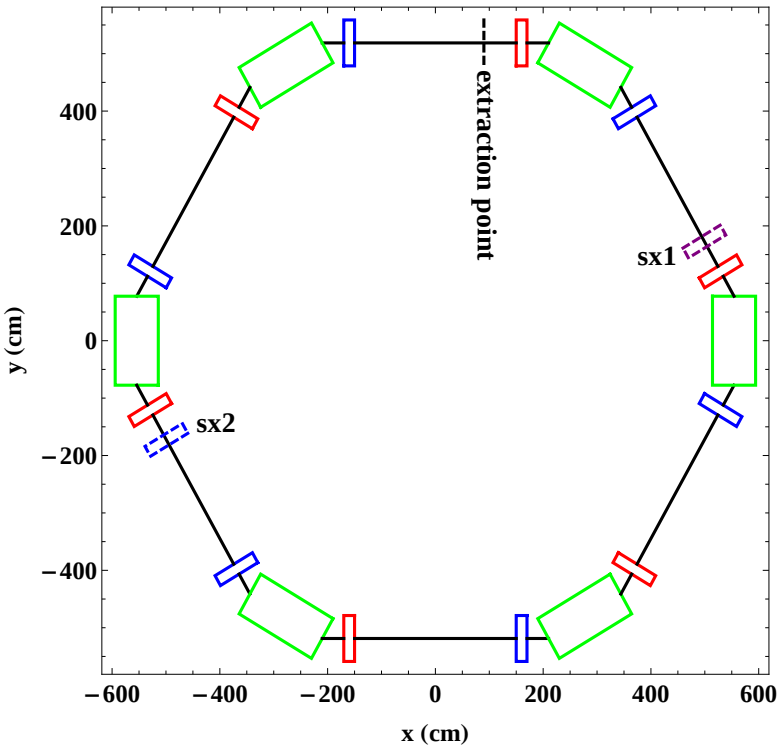

FIGURE 5.6: Initially a single sextupole, sx1, is included. Later a second sextupole, sx2, is inserted at a horizontal phase advance of πQ_y from sx1.

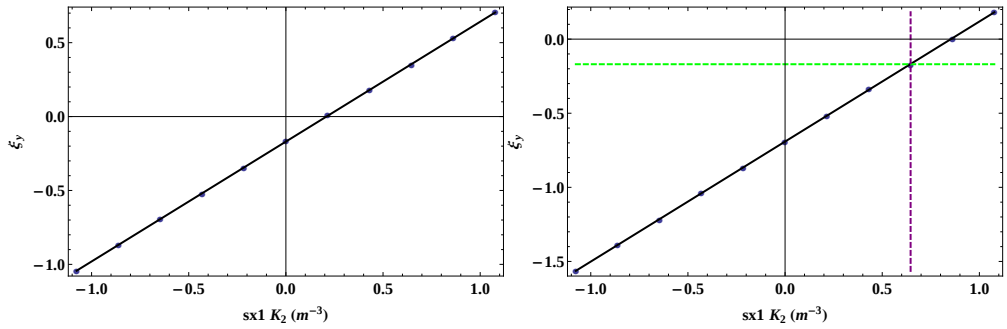

(a) Single sextupole.

(b) Sextupole pair.

FIGURE 5.7: A single sextupole (sx1) is inserted into the lattice; the chromaticity of the lattice is dependent upon the strength of the sextupole (a). A second sextupole (sx2) is inserted at at a horizontal phase advance of πQ_y from sx1 ; the strength of sx2 is fixed arbitrarily at $K_2 = -0.65\,\mathrm{m}^{-3}$) (with $-K_2$ for sx2 shown in (b) by the purple dashed line), it can be seen that when the strength of sx1 is equal and opposite to that of sx2, then the chromaticity of the lattice is the same as when no sextupole perturbation is included (green dashed line).

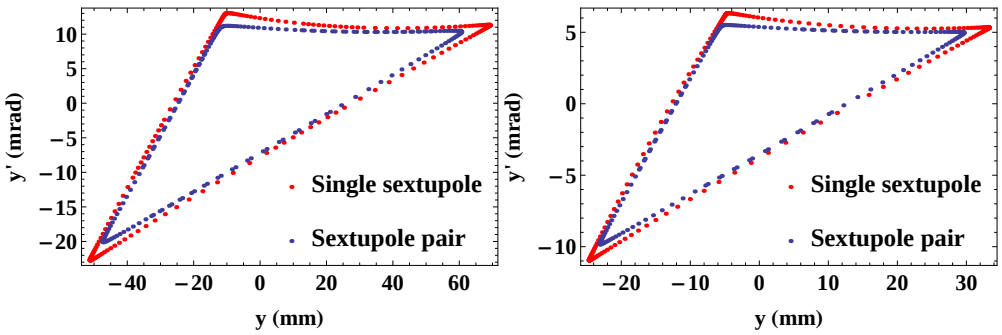

(a) $K_2 = 2.15\,\mathrm{m}^{-3}$ (b) $K_2 = 4.3\,\mathrm{m}^{-3}$

FIGURE 5.8: The separatrix produced by a pair of sextupoles of strengths $+K_2$ and $-K_2$ and separated by a horizontal phase advance of πQ_y is approximately equivalent to that produced by a single sextupole of strength $2K_2$ (demonstrated in Fig. (a) and (b)). The area enclosed by the separatrix is dependent upon the sextupole strength.

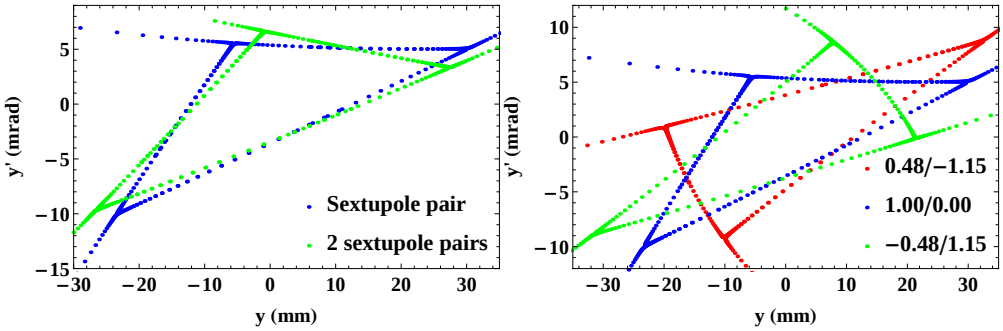

(a) (b)

FIGURE 5.9: Figure (a) compares the separatrix formed in Fig. 5.8b with a single sextupole pair to that formed when a second sextupole pair (as shown in Fig. 5.12) is included. The sextupole strengths for the single pair are $\pm K_2$, whilst the strengths of the two pairs are $\pm K_2/\sqrt{3}$, leading to separatrices that contain approximately the same areas, but are rotated with respect to one another. When two sextupole pairs are included, the rotation can be controlled by varying the ratio of the sextupole strengths in the first pair to those of the second pair. This is demonstrated in Fig. (b), where the separatrix orientation is shown for three sextupole strength settings. Multiplying the values in the legend by K_2 gives the sextupole strength for dsx1/dsx2 in each case.

separated by a horizontal phase advance of πQ_y from one another and are equal and opposite in strength. If the two families of sextupoles have the same strengths ($\pm K_2/\sqrt{3}$), then the area contained by the separatrix will be approximately the same as that obtained with a single sextupole pair with strengths $\pm K_2$, however the separatrix will be rotated (Fig. 5.9a). Furthermore, the orientation of the separatrix can be controlled by varying the ratio of the strengths of the sextupoles in the first pair to those in the second pair (Fig. 5.9b).

The orientation of the separatrix is important in maximising the number of particles that are successfully extracted from the ring (extraction efficiency). Relative sextupole strengths of $0.665/-1.15$ for dsx1/dsx2 provide a separatrix that meets the requirements (which are discussed further in section 5.3.3) for good extraction efficiency. In

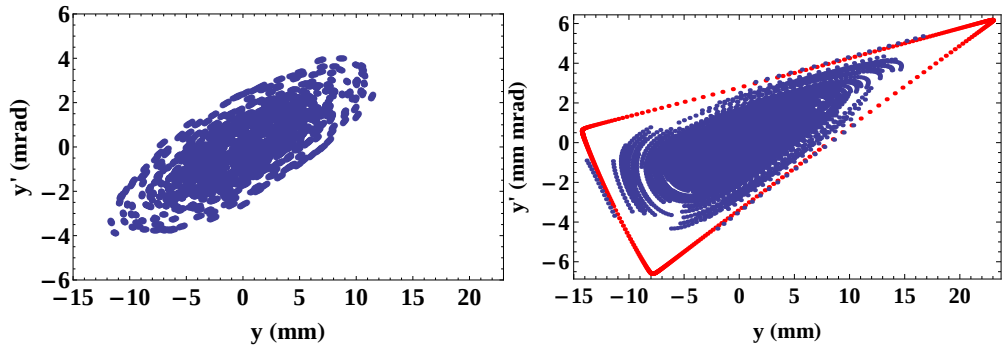

(a) Untuned (b) Extraction

FIGURE 5.10: Figure (a) shows the phase space portrait at the extraction point of a 60 MeV particle bunch with 1.5π mm mrad normalised emittance, untuned refers to the lattice being in the nominal acceleration state (i.e. quadrupole strengths as given under the injection heading of table 5.2). Fig. (b) shows the same bunch once the lattice has been shifted towards a third-integer resonance, and the perturbing sextupoles have been activated.

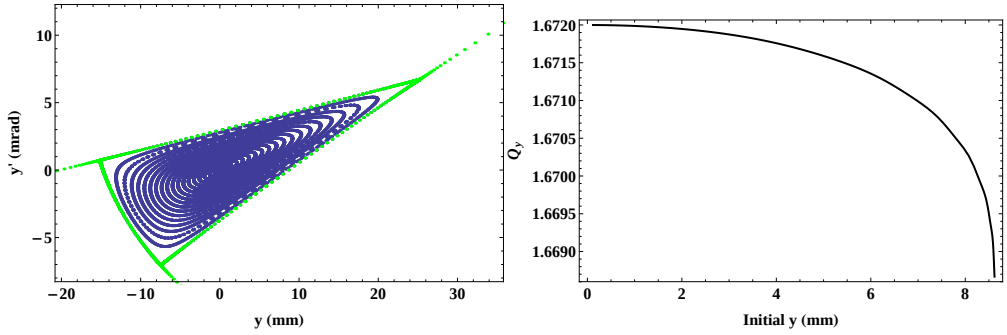

(a) (b)

FIGURE 5.11: The initial tracking coordinates of a particle are varied from (y, y') = (0 mm, 0 mrad) to a point close to the separatrix (8.07 mm, 0 mrad) (Fig. (a)). The rate of change of tune shift with amplitude increases rapidly as the particle start condition approaches the separatrix (Fig. (b)).

the following section, a theory for the effects of sextupole perturbations whilst at a third-integer tune is described, and further optimisations introduced. The betatron tune and sextupole perturbation strength will then be modified to give a separatrix that allows particles within a 60 MeV $\pm 1\%$ proton bunch (of normalised emittance of $1.5\,\pi$ mm mrad) to follow stable trajectories, but with the largest action particles close to the separatrix (Fig. 5.10).

5.3.3 Theory of sextupole perturbations

In the previous section the application of sextupole perturbations at a third-integer tune for slow extraction is demonstrated. In this section the process is explored analytically. This provides some insight into the effects of the sextupoles, and leads to methods that allow further optimisation of the slow extraction process. The description given here largely follows that in [87].

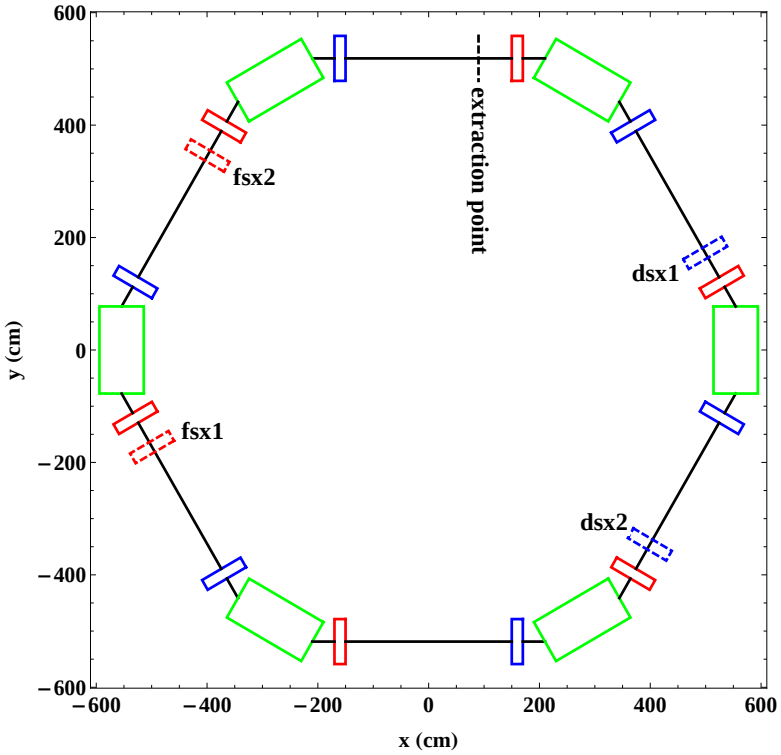

FIGURE 5.12: Dashed lines show the features related to extraction; including the black dashed line which shows the Poincaré plane, and the dashed blue and red rectangles that represent the perturbing sextupoles. Sextupoles dsx1 and fsx1 are of equal and opposite strength (as are dsx2 and fsx2), and the ratio of the strength of dsx1 to dsx2 determines the orientation of the separatrix in phase space.

The field in a sextupole is given by (see table 2.1):

$$B_y = -\frac{\partial^2 B_y}{\partial z^2} yz,$$
$$B_z = \frac{1}{2}\frac{\partial^2 B_z}{\partial y^2}(y^2 - z^2),$$

Assuming that the length of a sextupole is short compared to the betatron wavelength, the effect of a sextupole on the beam dynamics may be approximated as a single kick in the transverse momentum. In this case, the transverse dynamical variables are updated at the longitudinal midpoint of the sextupole according to:

$$\Delta y = \Delta z = 0,$$
$$\Delta y' = \frac{B_z L}{B\rho} = \frac{1}{2}\frac{L}{B\rho}\frac{\partial^2 B_z}{\partial y^2}(y^2 - z^2) = \frac{1}{2}LK_2(y^2 - z^2),$$
$$\Delta z' = -LK_2 yz,$$

where L is the length of the sextupole, and $B\rho$ is the beam rigidity. Given Eq. 4.15, which relates the variables y, y', z and z' to the normalised variables $\hat{y}$, $\hat{y}'$, $\hat{z}$ and $\hat{z}'$, the

effect of a sextupole kick in normalised coordinates can be written as:

$$\Delta\hat{y}' = \frac{1}{2}\beta_y^{3/2} L K_2 (\hat{y}^2 - \frac{\beta_z}{\beta_y}\hat{z}^2) = S\left(\hat{y}^2 - \frac{\beta_z}{\beta_y}\hat{z}^2\right), \tag{5.1}$$

$$\Delta\hat{z}' = -2\frac{1}{2}\beta_y^{3/2} L K_2 \frac{\beta_z}{\beta_y}\hat{y}\hat{z} = -2S\frac{\beta_z}{\beta_y}\hat{y}\hat{z},$$

with S representing the normalised sextupole strength,

$$S = \frac{1}{2}\beta_y^{3/2} L K_2. \tag{5.2}$$

It can be seen from Eq. 5.1 that if a sextupole is placed at a point where β_y is significantly larger than β_z, then the effect of the sextupole may be approximated by analysis of the horizontal motion only (provided that the vertical tune does not lead to further resonances).

After n turns in an accelerator, the normalised horizontal coordinates of an unperturbed particle (without the sextupole) are transformed as:

$$\begin{pmatrix} \hat{y}_n \\ \hat{y}'_n \end{pmatrix} = \begin{pmatrix} \cos(2\pi n Q_y) & \sin(2\pi n Q_y) \\ -\sin(2\pi n Q_y) & \cos(2\pi n Q_y) \end{pmatrix} \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} = M_n \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix}. \tag{5.3}$$

If the fractional part of the tune is offset from one third or two thirds of an integer by a small difference ΔQ_y, then, from Eq. 5.3, the propagation of a particle through three turns is given by:

$$\begin{pmatrix} \hat{y}_3 \\ \hat{y}'_3 \end{pmatrix} = \begin{pmatrix} \cos\psi & \sin\psi \\ -\sin\psi & \cos\psi \end{pmatrix} \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} \approx \begin{pmatrix} 1 & \psi \\ -\psi & 1 \end{pmatrix} \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} = \Psi \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix}, \tag{5.4}$$

where $\psi = 6\pi\Delta Q_y$.

At this point the effect of the sextupole perturbation on the transverse oscillation amplitude may be calculated as follows. A scenario where a sextupole kick is encountered every third turn is used. The transfer matrix for three different cases is calculated: where the sextupole kick occurs after the first, second and third turns (cases A, B and C, respectively). This method is valid provided that the perturbations are small, and ultimately has the advantage of producing simple expressions for the change in $\hat{y}$ and $\hat{y}'$ over three turns. The final phase space variables in A, B and C are:

$$\begin{aligned}
\begin{pmatrix} \hat{y}_3 \\ \hat{y}'_3 \end{pmatrix}_A &= M_2 \left(M_1 \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} + \begin{pmatrix} 0 \\ S\hat{y}_1^2 \end{pmatrix} \right), \\
\begin{pmatrix} \hat{y}_3 \\ \hat{y}'_3 \end{pmatrix}_B &= M_1 \left(M_2 \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} + \begin{pmatrix} 0 \\ S\hat{y}_2^2 \end{pmatrix} \right), \\
\begin{pmatrix} \hat{y}_3 \\ \hat{y}'_3 \end{pmatrix}_C &= \Psi \begin{pmatrix} \hat{y}_0 \\ \hat{y}'_0 \end{pmatrix} + \begin{pmatrix} 0 \\ S\hat{y}_3^2 \end{pmatrix},
\end{aligned}$$

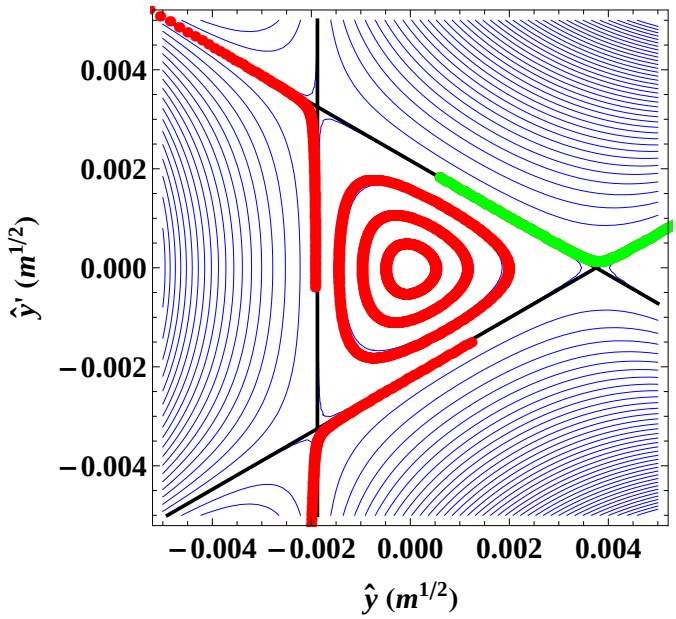

FIGURE 5.13: Phase space portrait in normalised coordinates at the location of the perturbing sextupole. The horizontal tune is 0.668 and the perturbing sextupole is of strength $K_2 = 8.62\,\mathrm{m}^{-3}$. Blue lines show contours of constant Hamiltonian, whilst red points are results obtained through numerical tracking. The progression of a particle that is just outside of the separatrix is marked in green for every third turn. This (green) path can calculated approximately using Eqs. 5.5.

where M_1, M_2 and Ψ are matrices that transport particles through one, two and three turns respectively. Summing the changes in $\hat{y}$ and $\hat{y}'$ in the three cases A, B and C, and then taking the first order terms in S and ψ gives formulae for calculating the effect of a sextupole on the phase space variables of a particle every three turns when the betatron tune is close to a third-integer. The formulae (which are commonly referred to in the literature as the spiral step and spiral kick equations) are:

$$\Delta\hat{y}_3 = \psi\hat{y}'_0 + \frac{3}{2}S\hat{y}_0\hat{y}'_0, \tag{5.5}$$

$$\Delta\hat{y}'_3 = -\psi\hat{y}_0 + \frac{3}{4}S(\hat{y}_0^2 - \hat{y}_0'^2).$$

The motion of a system that obeys the spiral step and kick equations may be approximated by Hamilton's equations with an appropriate Hamiltonian. We write:

$$\Delta\hat{y}_3 \simeq \frac{d\hat{y}}{dt}\Delta t, \quad \Delta\hat{y}'_3 \simeq \frac{d\hat{y}'}{dt}\Delta t. \tag{5.6}$$

Taking $\Delta t = 1$ for three turns, and, given Hamilton's equations (Eq. 3.1), we find:

$$\frac{d\hat{y}}{dt} = \frac{\partial H}{\partial \hat{y}'} = \psi\hat{y}' + \frac{3}{2}S\hat{y}\hat{y}',$$
$$\frac{d\hat{y}'}{dt} = -\frac{\partial H}{\partial \hat{y}} = -\psi\hat{y} + \frac{3}{4}S(\hat{y}^2 - \hat{y}'^2).$$

Integration of these equations with respect to $\hat{y}$ and $\hat{y}'$ then leads to the Hamiltonian:

$$H = \frac{\psi}{2}(\hat{y}^2 + \hat{y}'^2) + \frac{S}{4}(3\hat{y}\hat{y}'^2 - \hat{y}^3). \tag{5.7}$$

The fixed points are determined by $\frac{\partial H}{\partial J} = \frac{\partial H}{\partial \phi} = 0$, where J and ϕ are the action angle variables. The nature of the fixed points (stable or unstable) are determined by whether $\frac{\partial^2 H}{\partial J^2}$ and $\frac{\partial^2 H}{\partial \phi^2}$ have the same or opposite signs. The unstable fixed points in this case are found to be at $H = (2\psi)^3/(27S^2)$. The value of the Hamiltonian is constant on a separatrix. Therefore Eq. 5.7 can be used to obtain the equation for the separatrix:

$$H - \frac{(2\psi)^3}{27S^2} = \left(\frac{S}{4}\hat{y} + \frac{\psi}{6}\right)\left(\sqrt{3}\hat{y}' + \hat{y} - \frac{4\psi}{3S}\right)\left(\sqrt{3}\hat{y}' - \hat{y} + \frac{4\psi}{3S}\right) = 0. \tag{5.8}$$

Eq. 5.8 describes an equilateral triangle in normalised phase space. This is the shape of the separatrix in the approximation that we ignore higher-order terms in S and ψ when obtaining the spiral step and kick equations. The minimum distance from one side of the separatrix to the stable fixed point at the centre of the separatrix is:

$$h = -\frac{2\psi}{3S}. \tag{5.9}$$

The equation for the separatrix at the location of a sextupole can then be written as

$$\hat{y}\cos(\alpha) + \hat{y}'\sin(\alpha) = h,$$

where α is the angle measured anti-clockwise from the $\hat{y}$ axis to a line perpendicular to the separatrix. From Eq. 5.8, the solutions for α are $\pi/3$, π and $5\pi/3$, as can be seen in Fig. 5.13. As the observation point is moved around the ring, the separatrices rotate in phase space according to the phase advance between the sextupole and observation point ($\Delta\mu$). The equation for a separatrix at an arbitrary observation point is therefore:

$$\hat{y}\cos(\alpha - \Delta\mu) + \hat{y}'\sin(\alpha - \Delta\mu) = h. \tag{5.10}$$

The advantage of slow extraction using third-integer resonance is that a particle will remain on a stable orbit until it is pushed across the separatrix. Once in the unstable region of phase space, the particle amplitude grows predictably according to the spiral step and kick equations. As the amplitude of oscillation in normalised phase space is given by $A = \sqrt{\hat{y}^2 + \hat{y}'^2}$, for a particle on resonance ($\psi = 0$), adding the spiral step and

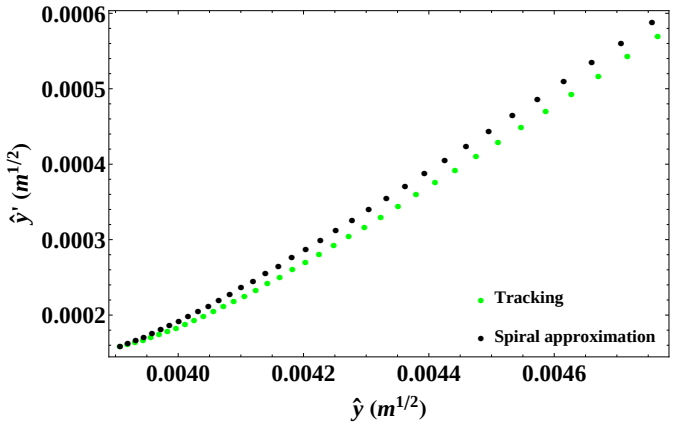

FIGURE 5.14: Comparison of spiral arm obtained through tracking (as in Fig. 5.13) with that calculated by the spiral approximation of Eq. 5.5.

kick equations in quadrature gives the change in amplitude over three turns,

$$\begin{aligned}\Delta A &= \sqrt{\Delta\hat{y}^2 + \Delta\hat{y}'^2}, \\ &= \sqrt{\left(\frac{3}{2}S\hat{y}_0\hat{y}'_0\right)^2 + \left(\frac{3}{4}S(\hat{y}_0^2 - \hat{y}_0'^2)\right)^2}, \\ &= \sqrt{\frac{9}{16}S^2(\hat{y}_0^4 + 2\hat{y}_0^2\hat{y}_0'^2 + \hat{y}_0'^4)}.\end{aligned}$$

Hence:

$$\Delta A = \frac{3}{4}SA^2. \tag{5.11}$$

As the change in amplitude increases with amplitude, eventually the change over three turns will be large enough to allow a beam to cross the thin wire of an electrostatic septum. The purpose of the electrostatic septum is to provide a transverse momentum kick to the beam, which then translates to a change in the horizontal position after some phase advance of the beam. The horizontal displacement permits the beam to enter the field region of a magnetic septum without hitting the (relatively thick) blade of the magnetic septum. Within this thesis, the dynamics of the extraction process from the electrostatic septum onwards are not considered in detail, and particles are considered to have been extracted once they have reached a threshold horizontal position at the location of the electrostatic septum. The location of the electrostatic septum and the orientation of the separatrix at this point are key factors in maximising the fraction of beam that is successfully extracted from a machine. The septum should be positioned where the horizontal β function is large, and the separatrix should be oriented in phase space so that one arm of the separatrix makes a 45° angle with the normalised horizontal ($\hat{y}$) axis. The 45° angle is selected to meet a number of criteria. In particular, optimisation of this angle ensures that particles approaching the septum blade are more likely, on a later turn, to cross the blade without hitting it, while travelling at the desired angle to be cleanly extracted. The orientation of the separatrix in normalised phase space

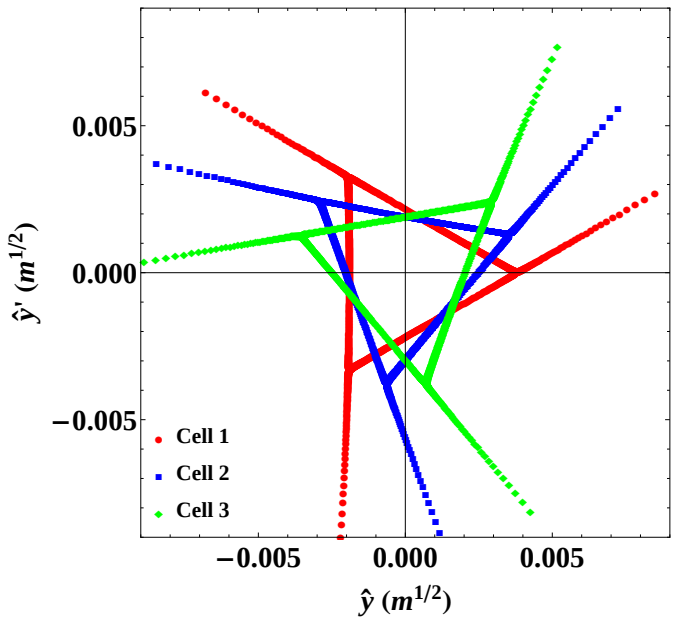

FIGURE 5.15: The separatrix rotates clockwise in phase space with increasing phase advance. Given a horizontal tune of $\approx 5/3$, then the separatrix in normalised phase space viewed at the same relative position within successive cells rotates clockwise by $\frac{5\pi}{9}$ per cell.

at the location of the electrostatic septum is determined by the phase advance between the perturbing sextupole and the septum. Figure 5.13 shows a phase space portrait (in normalised coordinates) for the lattice at the longitudinal midpoint of the sextupole. Numerical tracking results obtained with Zgoubi are marked in red, apart from one arm of the separatrix, which is relevant for Fig. 5.14, that is marked in green.

The system for slow extraction described in section 5.3.2 used a number of sextupoles in order to perturb the beam. The combination of these sextupoles determined both the orientation of, and the area enclosed by, the separatrix, whilst not affecting the chromaticity of the lattice. So far, the theory developed refers to a perturbation by a single sextupole of strength S. In our model, the combined effect of a number of sextupoles with different strengths and at different locations can be represented by a single "virtual" sextupole. The strength of the virtual sextupole, S_v, is given by adding the strengths of the real sextupoles whilst taking into account the phase advance between them:

$$S_v e^{3i\mu_{y,v}} = \sum_n S_n e^{3i\mu_{y,n}}, \tag{5.12}$$

where S_v is a real number. Hence:

$$S_v^2 = \left(\sum_n S_n \cos(3\mu_{y,n})\right)^2 + \left(\sum_n S_n \sin(3\mu_{y,n})\right)^2. \tag{5.13}$$

If two sextupoles of strengths S_1 and S_2, are separated by a horizontal phase advance of $\Delta\mu = \mu_2 - \mu_1$, then the strength of the virtual sextupole is:

$$\begin{aligned} S_v^2 &= (S_1\cos(3\mu_1) + S_2\cos(3\mu_2))^2 + (S_1\sin(3\mu_1) + S_2\sin(3\mu_2))^2 \\ &= S_1^2\cos^2(3\mu_1) + 2S_1S_2\cos(3\mu_1)\cos(3\mu_2) + S_2^2\cos^2(3\mu_2) \\ &\quad + S_1^2\sin^2(3\mu_1) + 2S_1S_2\sin(3\mu_1)\sin(3\mu_2) + S_2^2\sin^2(3\mu_2). \end{aligned}$$

Given that $\cos^2\theta + \sin^2\theta = 1$:

$$S_v^2 = S_1^2 + 2S_1S_2\cos(3\mu_1)\cos(3\mu_2) + 2S_1S_2\sin(3\mu_1)\sin(3\mu_2) + S_2^2, \tag{5.14}$$

and as

$$\begin{aligned} \sin\theta_a\sin\theta_b &= \frac{1}{2}(\cos(\theta_a - \theta_b) - \cos(\theta_a + \theta_b)), \\ \cos\theta_a\cos\theta_b &= \frac{1}{2}(\cos(\theta_a - \theta_b) + \cos(\theta_a + \theta_b)), \end{aligned}$$

then,

$$\begin{aligned} S_v^2 &= S_1^2 + S_1S_2(\cos(3\mu_1 - 3\mu_2) + \cos(3\mu_1 + 3\mu_2)) \\ &\quad + S_1S_2(\cos(3\mu_1 - 3\mu_2) - \cos(3\mu_1 + 3\mu_2)) + S_2^2 \\ &= S_1^2 + 2S_1S_2\cos(3\Delta\mu) + S_2^2. \end{aligned} \tag{5.15}$$

If the sextupoles are separated by an even multiple of $\pi/3$, then the geometric perturbations will add. In the case where sextupoles are separated by πQ_y, where $Q_y = n \pm \frac{1}{3}$, the geometric perturbations will add in phase when n is an odd integer and in anti-phase when n is an even integer.

Now we consider the contribution of a pair of sextupoles to the chromaticity of the lattice. Starting with Eq. 2.6, which gives the chromaticity of a lattice, the effect of sextupole fields on the horizontal and vertical chromaticities is:

$$\begin{aligned} \Delta\xi_y &= \frac{1}{4\pi}\oint \beta_y K_2 D_y \mathrm{d}s, \\ \Delta\xi_z &= -\frac{1}{4\pi}\oint \beta_z K_2 D_y \mathrm{d}s. \end{aligned}$$

After making the thin lens approximation, the above equations are written as:

$$\begin{aligned} \Delta\xi_y &= \frac{1}{4\pi}\sum_n \beta_{y,n} L K_{2,n} D_{y,n}, \\ \Delta\xi_z &= -\frac{1}{4\pi}\sum_n \beta_{z,n} L K_{2,n} D_{y,n}, \end{aligned}$$

where $LK_{2,n}$ is the integrated strength of the n^{th} sextupole. For two sextupoles (at locations of equal dispersion and β function), the change in horizontal chromaticity in

terms of normalised sextupole strength is:

$$\Delta\xi_y = \frac{1}{4\pi}\frac{2D_y(S_1+S_2)}{\sqrt{\beta_y}},$$

meaning that the contributions of individual sextupoles add to the chromaticity regardless of whether n is even or odd, and that the introduction of chromatic effects and geometric perturbations can be separated (as was demonstrated in section 5.3.2). Specifically, when n is even, two sextupoles that are separated by πQ_y will contribute to the chromaticity of the lattice and not introduce geometric perturbations if they have the same strength. However, if they have opposite strengths, then there will be no net chromatic effects but the pair will introduce a geometric perturbation. The phase of the virtual sextupole can then be calculated as

$$\tan(3\mu_{y,v}) = \frac{\sum_n S_n \sin(3\mu_{y,n})}{\sum_n S_n \cos(3\mu_{y,n})}. \tag{5.16}$$

In the case of the sextupole pair with equal or opposite strengths, the virtual sextupole is located at the position of one of the real sextupoles. Figure 5.12 depicts the case where two pairs of sextupoles are used to produce a geometric perturbation. Each pair consists of sextupoles with opposing strength, that are separated by a phase advance of πQ_y, thus ensuring that there is no net change in chromaticity. The aim of using two sextupole pairs is to provide control over the orientation of the separatrices in phase space whilst not affecting the overall strength of the geometric perturbation. Each pair is now replaced by a virtual sextupole, and the rotation of the separatrix (relative to the separatrix produced by a single sextupole at the position of S_1) is given by

$$\tan(3\mu_{y,v}) = \frac{S_2 \sin(3\mu_2)}{S_1 + S_2\cos(3\mu_2)}.$$

For the six cell lattice, separating two virtual sextupoles by a single cell, $\Delta\mu = \pi Q_y/3$, gives a geometric perturbation

$$S_v^2 = S_1^2 + S_1S_2 + S_2^2,$$

and a separatrix oriented according to

$$\tan(3\mu_{y,v}) = -\frac{\sqrt{3}S_2}{2S_1+S_2},$$

thus allowing the separatrix to be rotated by $\pm 30°$. Together with the rotation through $\pm 180°$ that can be achieved by reversing the polarity of the virtual sextupole, this is sufficient to achieve any required orientation of the separatrix, as shown in Fig. 5.16.

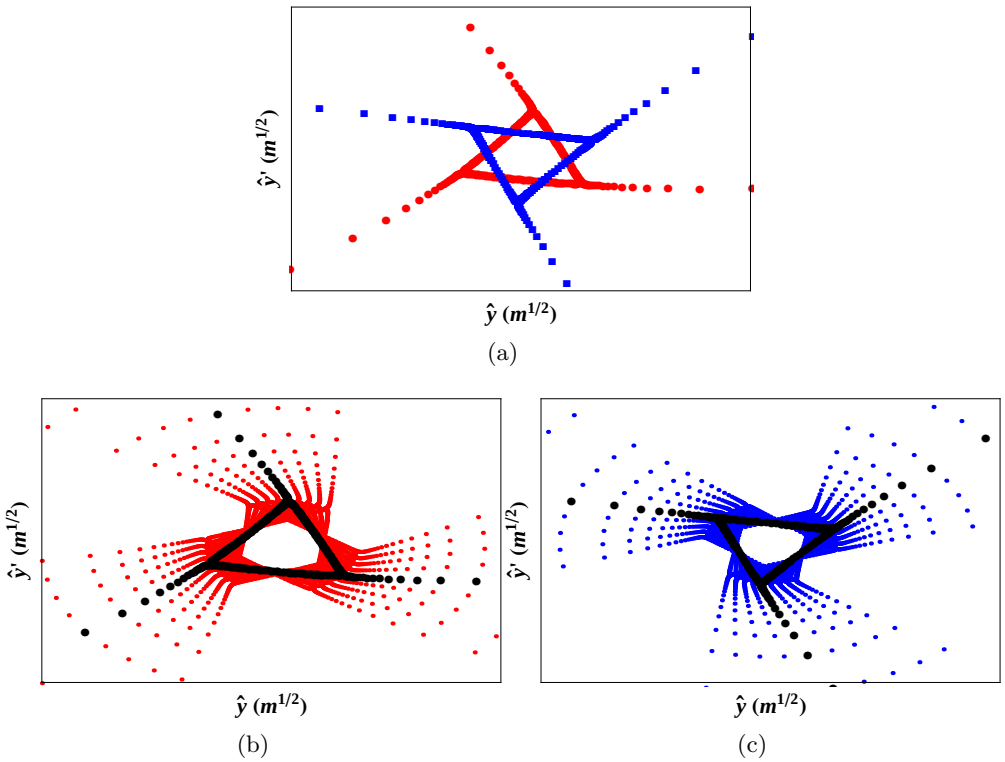

FIGURE 5.16: Using two pairs of sextupoles to perturb the beam allows the separatrices at the extraction plane to be tuned to any desired orientation in normalised phase space. The sextupoles within a pair have opposite strengths and a phase advance of πQ_y between them, ensuring that they make no contribution to the chromaticity of the lattice. For a single sextupole pair, there are two possible orientations for the separatrix at the extraction plane (differing by 180°). These orientations are shown in (a), and are determined by the polarity of the sextupole that is located in the cell before the extraction plane. Including a second pair allows for the separatrices of each orientation in (a) to be rotated by ±30° (b) and (c) whilst maintaining a constant perturbation strength. In (b) and (c) the separatrices of (a) are shown in black, the coloured points then indicate the range of orientations possible for the given example with two pairs of sextupoles.

5.3.4 Off-momentum particles

So far, little attention has been paid to the dynamics of particles away from the design momentum, other than ensuring that the lattice chromaticity is not affected by the perturbation sextupoles. Figure 5.17a shows the separatrices for particles with $\delta = \pm 0.001$ as well as $\delta = 0$; dispersion at the extraction point means that the separatrices are offset from one another, whilst the area enclosed by the separatrix for higher momentum is smaller than that for low momentum due to the negative chromaticity of the lattice (meaning a higher momentum particle travelling along its equilibrium path has a tune closer to the resonant condition than a particle at lower momentum).

The extraction efficiency depends on the effective thickness of the electrostatic septum wire, which is determined by the alignment and length of the electrostatic septum as well as the physical thickness of the wire. Initially considering only the on-momentum particles, if the septum is rotated to match the beam divergence at the extraction point,

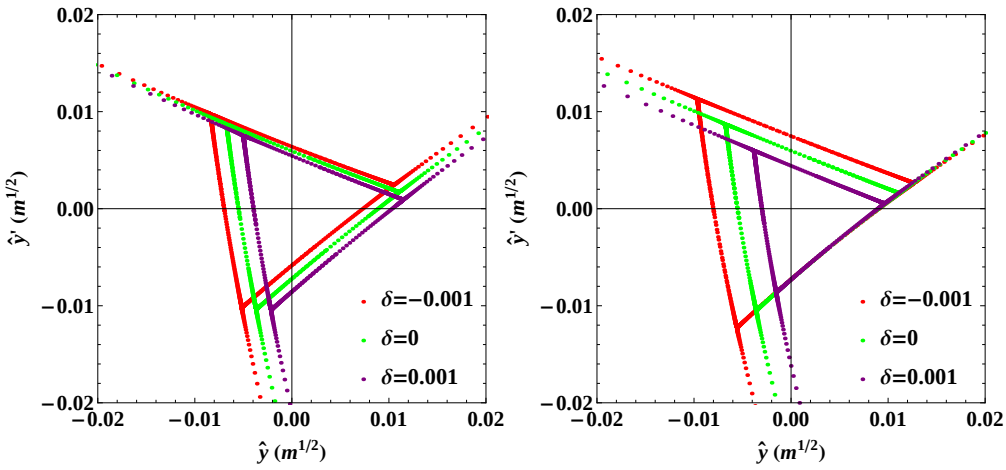

FIGURE 5.17: The separatrices for off-momentum particles at the extraction point before (a) and after (b) setting the chromaticity so that the Hardt condition is true. When the Hardt condition is satisfied, the number of particles that are lost by hitting an electrostatic septum may be reduced.

then the gap that needs to be jumped over three turns (Eq. 5.5) for a particle to be successfully extracted is equal to the physical thickness of the septum wire. Figure 5.17a shows that for the current lattice configuration, the separatrices for off-momentum particles are not aligned with that of an on-momentum particle, meaning that the divergence of the beam at the location of the septum depends on momentum. This will reduce the extraction efficiency as the effective thickness of the septum wires is increased for both high and low momentum particles. Consider a low momentum particle that does not quite enter the septum on a given turn. If the divergence of the particle is greater than the angle made between the septum and reference axis, then the particle may be lost at some point along the length of the septum. For a high momentum particle that just enters the septum, and has a divergence that is less than the angle made between the septum and the reference axis, if the electric field does not change the divergence of the particle rapidly enough, then this particle will also be lost at some point along the length of the septum. For off-momentum particles, the transverse oscillation of a particle can be written so as to include the first order dispersion:

$$\hat{y}_\beta = \hat{y} - \hat{D}\delta,$$
$$\hat{y}'_\beta = \hat{y}' - \hat{D}'\delta,$$

where $\hat{D}\delta$ and $\hat{D}'\delta$ are respectively the offsets in closed orbit location and divergence due to dispersion. The effect of a sextupole on the beam dynamics is now:

$$\Delta\hat{y}'_\beta = S\hat{y}^2 = S(\hat{y}_\beta + \hat{D}\delta)^2, \tag{5.17}$$

and the Hamiltonian including the dispersion term is written as:

$$H = \frac{1}{2}\left(\psi - 3S\hat{D}\delta\right)(\hat{y}_\beta^2 + \hat{y}_\beta'^2) + \frac{S}{4}(3\hat{y}_\beta\hat{y}_\beta'^2 - \hat{y}_\beta^3). \tag{5.18}$$

When including a momentum offset, the equation of a line that represents one side of the separatrix is:

$$(\hat{y} - \hat{D}\delta)\cos(\alpha - \Delta\mu) + (\hat{y}' - \hat{D}'\delta)\sin(\alpha - \Delta\mu) = h. \tag{5.19}$$

The value h, which is defined in Eq. 5.9, is proportional to the tune distance from a third-integer resonance ($\Delta Q_y = \psi/6\pi$). The equation for a separatrix for an off-momentum particle is then found by including the chromaticity of the lattice in Eq. 5.19, by setting:

$$h = \frac{-4\pi}{S}(\Delta Q_y - \xi_y\delta). \tag{5.20}$$

Equation 5.19, for the off-momentum separatrix, is then:

$$(\hat{y} - \hat{D}\delta)\cos(\alpha - \Delta\mu) + (\hat{y}' - \hat{D}'\delta)\sin(\alpha - \Delta\mu) = -\frac{4\pi}{S}(\Delta Q_y - \xi_y\delta). \tag{5.21}$$

The equation for one arm of the separatrix (i.e. one value of α) can be made independent of momentum at the extraction point of the lattice ($\Delta\mu$). This is true when:

$$\hat{D}\cos(\alpha - \Delta\mu) + \hat{D}'\sin(\alpha - \Delta\mu) = -\frac{4\pi}{S}\xi. \tag{5.22}$$

This equation defines the Hardt condition. In the case of the example synchrotron, the Hardt condition is imposed at the extraction point by suitable adjustment of the chromaticity of the lattice. This is achieved by introducing two extra sextupoles that are equal in strength and separated by a phase advance of πQ_y; the location of these sextupoles can be seen in Fig. 5.19, where there is a 10 cm drift space separating the chromatic sextupoles and the neighbouring perturbation sextupoles. As the area enclosed by the separatrix is dependent on the tune distance from resonance, ΔQ, a non-zero chromaticity can mean that the area of the stable region enclosed by a separatrix changes rapidly with momentum (Fig. 5.17b). It is intended particles with momentum deviation within the range $\pm 0.1\%$ can be extracted from the lattice; bringing the chromaticity close to 0 limits the change in the stable area for this range of momenta. Considering a separatrix arm that is oriented at an angle of $\pi/4$ radians to horizontal axis ($\alpha = 5\pi/3$ and $\Delta\mu = -\pi/12$), Eq. 5.22 becomes:

$$\frac{1}{\sqrt{2}}(\hat{D} - \hat{D}') = -\frac{4\pi}{S}\xi. \tag{5.23}$$

If the normalised dispersion and angular dispersion are equal, then the chromaticity that satisfies the Hardt condition is zero. Figure 5.18 shows the normalised dispersion and angular dispersion in the long drift of the proton synchrotron. It can be seen that the

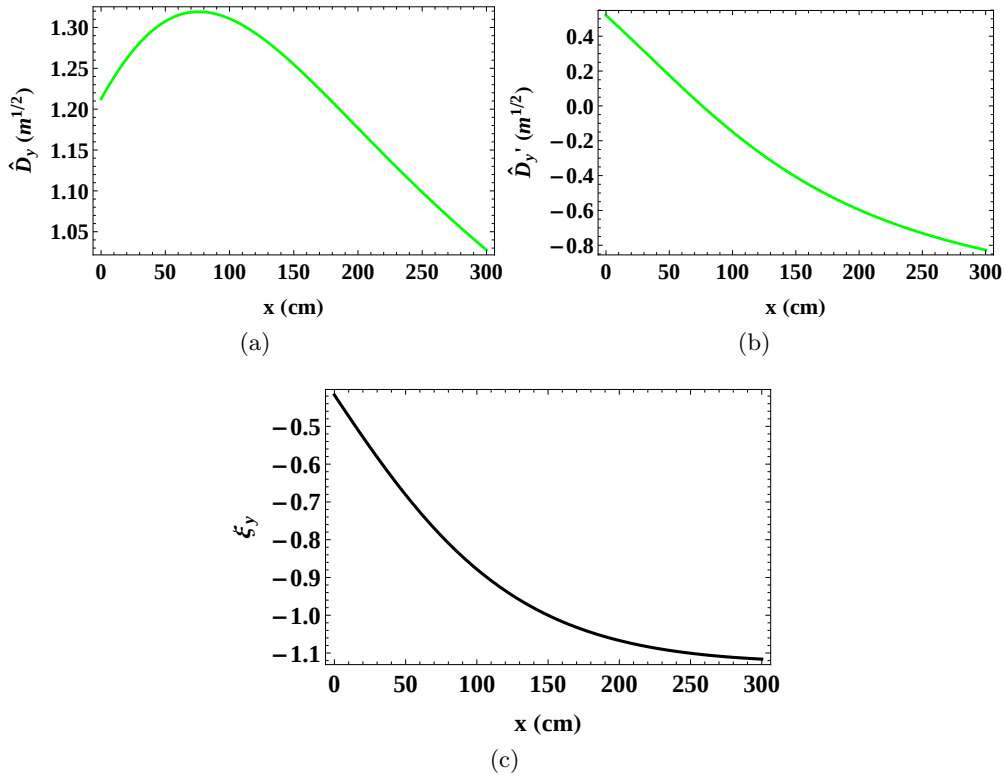

FIGURE 5.18: Plots to show the normalised dispersions in the long drift of the initial proton synchrotron.At a position where a separatrix arm is oriented at $\pi/4$ from the horizontal axis in normalised phase space, the chromaticity that meets the Hardt condition will be zero if the normalised dispersion and angular dispersion are equal. Figure 5.18c shows the chromaticity required to meet the Hardt condition versus the location along the long drift of the extraction point. For the proton synchrotron, the chromaticity required to meet the Hardt condition is closest to zero at the start of the long drift.

chromaticity required to meet the Hardt condition will be closer to zero if the extraction point is towards the beginning of the long drift. A new extraction position located 40 cm into the long drift of the first cell is selected (Fig. 5.19). To maintain a rotation of approximately $\Delta\mu = \pi/4$ for the separatrix, the relative sextupole strengths are now set to dsx1/dsx2=0.69/0.457. The off-momentum separatrices for this new lattice configuration are shown in Fig. 5.17b.

5.4 Final lattice parameters

The parameters of the lattice that is used to demonstrate extraction will now be presented. It is not intended that these parameters will give the optimum extraction for this machine. Instead the techniques and theories developed in the previous sections have been considered and applied to produce a lattice where the beam dynamics give favourable extraction efficiency; this is to say that further optimisation based upon the theory presented may improve to some extent the extraction efficiency of this machine.

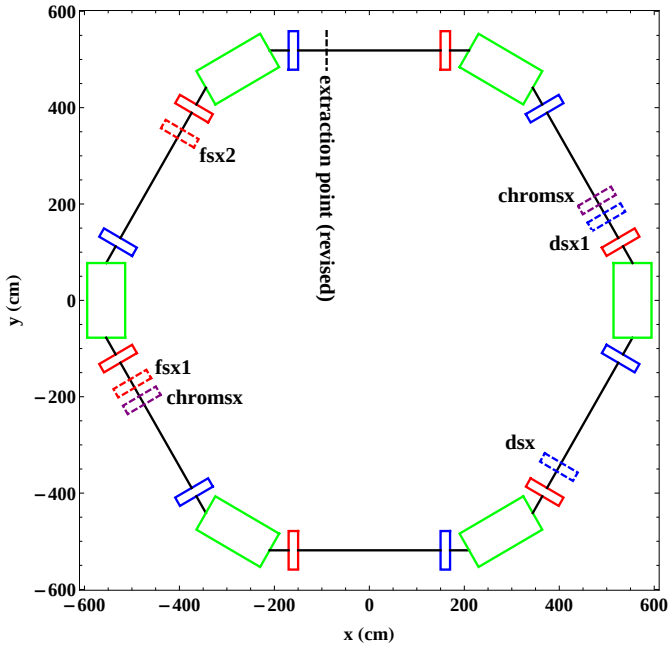

FIGURE 5.19: Plot of the proton synchrotron including the chromatic sextupoles (labelled 'chromsx') and the revised extraction position. The chromatic sextupoles are used to control the chromaticity so that the Hardt condition is satisfied, the position of one side of the separatrix is then constant for some range of momentum offset. The revised extraction point is selected to that the chromaticity required to meet the Hardt condition is closer to zero.

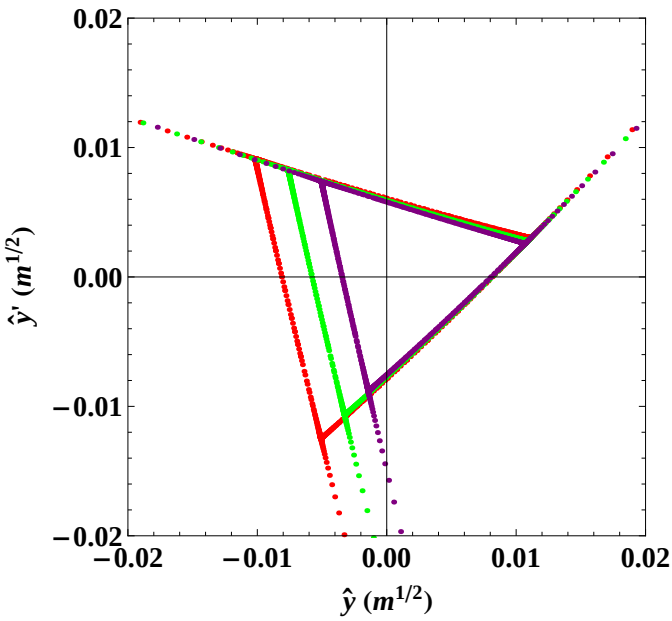

FIGURE 5.20: Plot to show the off-momentum separatrices for the revised extraction point. The normalised dispersion and angular dispersion are closer to being equal at the new extraction point than at the original extraction point; given Eq. 5.23, the chromaticity required to satisfy the Hardt condition is now closer to zero than at the original extraction point. The result of this is that the off-momentum and on-momentum separatrices are more similar in terms of the area they contain than was the case in Fig. 5.17b.

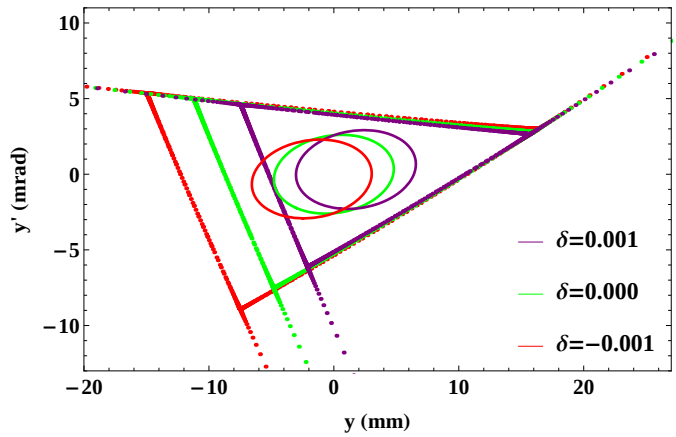

FIGURE 5.21: Plot to show the alignment of the separatrix arm at the extraction point once the lattice tunes have been moved close to resonance. The solid ellipses mark the rms emittance for a 60 MeV bunch ($\delta = \pm 0.001$ and 0) before the lattice is set for extraction.

The steps taken in modifying the lattice from the basic design (presented in Fig. 5.1) can be summarised as:

1. Move the horizontal tune close to one third of an integer. The selected tune is above third-integer, which means that when the beam is perturbed by a sextupole field, the separatrix formed will rotate clockwise in phase space with increasing phase advance (the separatrix rotates anti-clockwise if the tune is beneath third-integer).

2. Select an extraction point that minimises $\hat{D} - \hat{D}'$.

3. Introduce perturbation sextupoles. These sextupoles are arranged so that they do not affect the chromaticity of the lattice, but allow for the orientation of the separatrix to be set as required at the extraction point. The sextupoles are tuned so that only one arm of the separatrix intercepts an electrostatic septum, with this arm at an angle of $\approx \pi/4$ radians to the horizontal axis.

4. Introduce chromatic sextupoles. These sextupoles are arranged so that they affect the chromaticity of the lattice, but do not introduce geometric perturbations. The chromaticity is set so as to satisfy the Hardt condition at the extraction point. This ensures that the orientation of the relevant arm of the separatrix does not change with particle momentum (Fig. 5.21). A further objective is to keep the chromaticity value small, which helps maintain a constant area enclosed by the separatrix with respect to changes in momentum.

5. Optimisation. The area enclosed by the separatrix is dependent upon the tune difference from resonance, ΔQ, and the sextupole perturbation strength. The rate of amplitude growth and the chromaticity required to satisfy the Hardt condition are both dependent upon the sextupole perturbation strength. A large sextupole

TABLE 5.2: Parameters of the synchrotron model. The extraction quadrupole strengths move the lattice close to a horizontal third-integer tune. The sextupoles energised during extraction control the size and rotation of a separatrix.

Lattice	
Circumference	34.939 m
Periods	6
Lattice type	DOFB
Bend radius	1.55 m
Bending magnet length	1.623 m
Dipole field strength	0.2087 m^{-1}
Quadrupole lengths	20 cm
Sextupole lengths	20 cm
Injection	
Injection energy	5 MeV
Injection momentum spread (δ)	$\pm$ 0.1%
Normalised rms emittance	1.5 π mm mrad
Quadrupole strengths (H/V)	2.56/-0.492 m^{-2}
Tune (Q_h/Q_v)	1.800/1.850
Chromaticity (ξ_h/ξ_v)	-0.277/-2.59
Extraction	
Extraction energy	60-230 MeV
Quadrupole strengths (H/V)	2.319/-0.407 m^{-2}
Tune (Q_h/Q_v)	1.672/1.850
Chromaticity (ξ_h/ξ_v)	-0.772/-2.189
Sextupole strengths (dsx1/fsx1)	$-3.864/3.864\,\mathrm{m}^{-3}$
Sextupole strengths (dsx2/fsx2)	$-2.558/2.558\,\mathrm{m}^{-3}$
Sextupole strengths (chromsx)	$-1.775\,\mathrm{m}^{-3}$

perturbation strength helps to achieve a large rate of amplitude growth outside the separatrix. However, a large tune separation from resonance is then needed to ensure the area within the separatrix is large enough to encompass the beam. If the tune spread is large, then it is difficult to achieve efficient extraction using the transverse rf deflecting cavity. The tune spread and sextupole perturbation strengths are optimised taking into account these conflicting requirements. The chromaticity to satisfy the Hardt condition is then achieved by adjusting the chromatic sextupoles.

The lattice configuration gives a virtual sextupole of strength, $S_v = 11.7\,\mathrm{m}^{-1/2}$ and phase advance of $\Delta\mu = -0.21$ rad. The final locations of the sextupoles are shown in Fig. 5.22. The theoretical separatrix for the final lattice (with a horizontal tune of 1.672) at the extraction point is shown in Fig. 5.22. The normalised dispersion and angular dispersion at the extraction point are found through tracking to be $\hat{D} = 1.202\,\mathrm{m}^{1/2}$ and $\hat{D}' = -0.309\,\mathrm{m}^{1/2}$. After applying the theory for the off-momentum separatrices, we find that the horizontal chromaticity required to meet the Hardt condition for the given lattice configuration is $\xi_y = -0.54$. However, given the small difference between theory and tracking (Fig. 5.22), the chromaticity found through tracking that

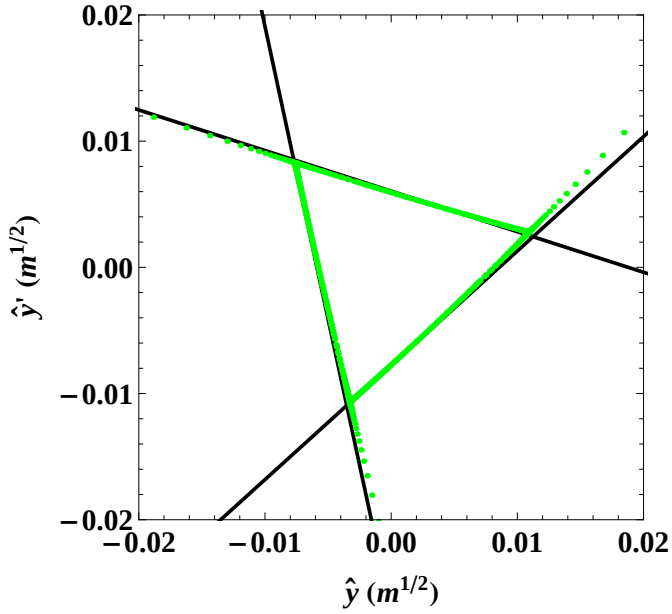

FIGURE 5.22: The black lines show the theoretical on-momentum separatrix obtained for the extraction point of the final lattice configuration using Eq. 5.10. The green markers show the separatrix obtained through tracking. The small differences are accounted for by the approximations made when developing the theory.

minimises the separation in phase space of the relevant off-momentum separatrices is $\xi_y = -0.77$. The chromatic sextupole strengths required for a chromaticity of $\xi_y = -0.77$ are $K_2 = -1.775\,\mathrm{m}^{-3}$

5.5 Transverse deflecting cavity

The slow extraction process that is introduced in section 2.8.1 and discussed in more detail in this chapter requires an rf cavity that provides a momentum kick transverse to the longitudinal axis of the cavity (Fig. 5.23). As no such element was available within the Zgoubi code, a new element, based upon the existing model of an rf cavity for longitudinal acceleration, was created. The existing element updates the kinetic energy of a particle according to Eq. 2.47: this approach works well given that the longitudinal momentum should remain greater than zero throughout tracking. As the transverse momentum may change direction under the influence of the rf kicker, the new rf cavity element updates the momentum of the particle directly rather than the kinetic energy. The change in momentum (Δp_y) of a charged particle as it travels through an electric field that is directed along the transverse horizontal axis is:

$$\Delta p_y = qE_y\Delta t,$$

where q is the charge of the particle, E_y is the electric field along the transverse horizontal axis and Δt is the time taken for a particle to travel through the cavity. If we take a rectangular cavity of length l as an example and the particle has a velocity $\mathbf{v}$, then the

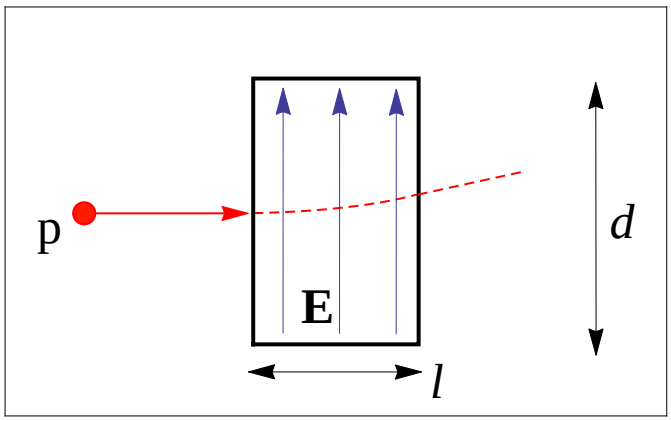

FIGURE 5.23: A transverse deflecting cavity of horizontal width d and longitudinal length l is modelled as a momentum kick (Eq. 5.24) at the location of the cavity.

above equation may be written as:

$$\Delta p_y = qE_y \frac{l}{v_x},$$

where v_x is the velocity of the particle along the longitudinal axis of the cavity. As the electric field is related to the voltage V_y by:

$$E_y = \frac{V_y}{d},$$

where d is the transverse width of the cavity, then the change in momentum in terms of the cavity voltage is:

$$\Delta p_y = q\left(\frac{V_y l}{d}\right)\frac{1}{v_x}.$$

As with the longitudinal cavity, the field varies sinusoidally, and so the change in momentum due to the transverse cavity is:

$$\Delta p_y = \frac{1}{v_x} q V_0 \sin(\phi), \tag{5.24}$$

where ϕ is the phase of the rf voltage at which the particle arrives, and V_0 is based on the peak voltage across the cavity gap:

$$V_0 = V_{y,\max}\frac{l}{d}.$$

In the following simulations, values of V_0 of up to 2 kV are applied, which corresponds to a maximum angular kick of 2 μrad (for a 60 MeV particle). This is in keeping with literature on the HIMAC synchrotron.

5.6 Simulation of beam extraction

5.6.1 Beam interrupt time

The ability to pause beam extraction is important when delivering particle therapy; for example, when treating moving tumours, respiratory gating sees the beam switched off at certain times during the respiration cycle, and prevents unwanted dose being given to healthy tissue. Being able to pause extraction means that dose delivery can be stopped and started without the need for an acceleration cycle in-between. In order to limit the dose to healthy tissue, the time taken to reduce the output current from the nominal treatment value to zero should be short. Here we assume that the time taken to pause extraction is given by the time taken for particles in the unstable region of phase space to be intercepted by the electrostatic septum.

The time taken for a particle to move from the separatrix to cross the blade of an extraction septum in the case that the area enclosed by the separatrix is zero (i.e. tune is exactly on resonance) can be estimated using Eq. 5.11. To obtain an estimate of the time taken to pause the beam in practice, simulations were carried out. A total of 5000 particles were located randomly, given a uniform distribution, around the theoretical separatrix (Fig. 5.24). The particles were then tracked for up 600 turns without rf, with particles moving beyond a transverse horizontal position of 5 cm at the extraction point considered as having been extracted. Of the 5000 particles, 2581 remained within the stable region of phase space, whilst the rest were extracted. The time at which each particle is extracted is recorded, and the results presented as a histogram (Fig. 5.25). It can be seen that all of the extracted particles have left the machine before 70 μs, which corresponds to approximately 280 turns in the machine. This time is comparable to that found experimentally for the HIMAC synchrotron (50 μs) [88].

5.6.2 Extraction with rf perturbation

For particle extraction through third-integer resonance with rf knockout, the particles within a bunch initially follow stable trajectories: extraction only begins once the transverse deflection cavity is activated. Each time a particle passes the cavity, it will receive a small kick in momentum that is dependent upon both the phase and voltage of the rf (Eq. 5.24). If the kicks add coherently over a number of turns then there will be a significant change in betatron action, which can ultimately lead to the particle crossing the separatrix into the unstable region of phase space.

The effects of the perturbation from the transverse deflecting cavity on the trajectories of particles following paths within the separatrix have been studied in simulation. Some initial results are presented in Fig. 5.26. In this first simulation, two particles are tracked: one having initial coordinates $y = 4\,\text{mm}$, $y' = 0\,\text{mrad}$ and the second with $y = 7.5\,\text{mm}$, $y' = 0\,\text{mrad}$. The particles are initially tracked through 4000 turns with the transverse rf voltage set to $V_0 = 0$, and the results in action-angle variables shown in Fig. 5.26a. The action is seen to vary with phase angle (and by a greater amount

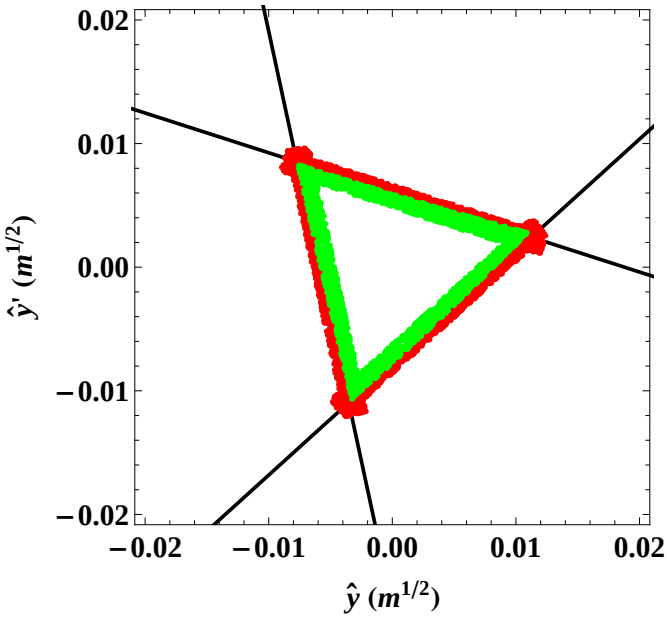

FIGURE 5.24: A total of 5000 particles were tracked for up to 600 turns in order to find an estimate of the maximum time taken for extraction to be paused. The solid black lines represent the separatrix obtained through theory by using Eq. 5.10, green points represent the initial phase space coordinates of the particles that remained on a stable orbit and red points are the start coordinates of the particles that were extracted after some number of turns.

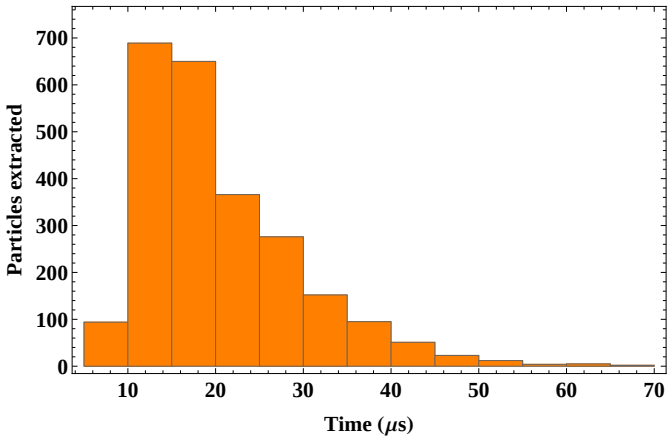

FIGURE 5.25: Histogram showing the numbers of particles extracted in successive 5 μs intervals from the start of tracking. The starting phase space coordinates of the particles that are included in the histogram are given by the red points in Fig. 5.24. It can be seen that extraction of the sample particles has stopped by 70 μs.

in the case of the larger amplitude particle), this is due to the triangular distortion of trajectories in phase space resulting from the influence of the geometric sextupoles. The rf voltage is then set to $V_0 = 2\,\mathrm{kV}$, and tracking carried out under two conditions of the deflecting cavity frequency: first when the rf frequency is set to $f_{\mathrm{rf}} = 1.6718 f_{\mathrm{rev}}$ and then when $f_{\mathrm{rf}} = 1.6707 f_{\mathrm{rev}}$. In both cases the initial phase of the rf is $\phi = \pi/4$. These two frequencies are based on the betatron tune shift with amplitude (Fig. 5.11b), with the first frequency chosen to be at a harmonic of the betatron oscillation frequency of the low amplitude particle and the second at a harmonic of the high amplitude particle. Results for the simulations are shown in Figs. 5.26b and 5.26c respectively, where the areas enclosed by particles in phase space are seen to oscillate for both rf frequencies, with the greatest oscillation amplitude in each case observed for the particle that has initial betatron frequency that is a harmonic of the rf frequency. For the initial particle start coordinates and rf frequencies applied, neither of the particles crosses the separatrix during tracking. Further, it can be seen that an rf frequency that resonates with a particle that is close to the separatrix will have only a small effect on the amplitude of a particle that is close to the closed orbit. In practice, the rf perturbation is split into two components; the first is a field with a frequency that varies in time, the purpose of which is to diffuse particles from the centre of the bunch outwards towards the separatrix, the second is a field of fixed frequency that pushes particles over the separatrix.

The slow extraction process takes place over a number of seconds, which corresponds to tens of millions of turns in the accelerator. Efforts to find the time structure of the extracted beam, using the tracking methods already presented, have not been made due to the CPU time involved. The emittance and extraction efficiency are calculated by tracking particles that are close to the separatrix at the start of the simulation. The effects of a number of variables on the extraction of particles will be investigated through simulation, these are:

1. The phase of transverse rf at the start of the simulation.
2. The septum position.
3. The amplitude of the perturbing voltage.
4. The particle momentum.

In all of these cases, tracking is carried out for 600 turns. The initial coordinates for the particles are the same as those for the surviving particles from the study on the beam interrupt time. In order to estimate the extraction efficiency, an electrostatic septum wire of width 0.1 mm is considered; for a given extraction amplitude, y_{ext}, particles that have a horizontal position of $y_{\mathrm{ext}} \leq y \leq y_{\mathrm{ext}} + 0.1\,\mathrm{mm}$ are said to have been lost at the septum, whilst particles with $y > y_{\mathrm{ext}} + 0.1\,\mathrm{mm}$ are assumed to be successfully extracted.

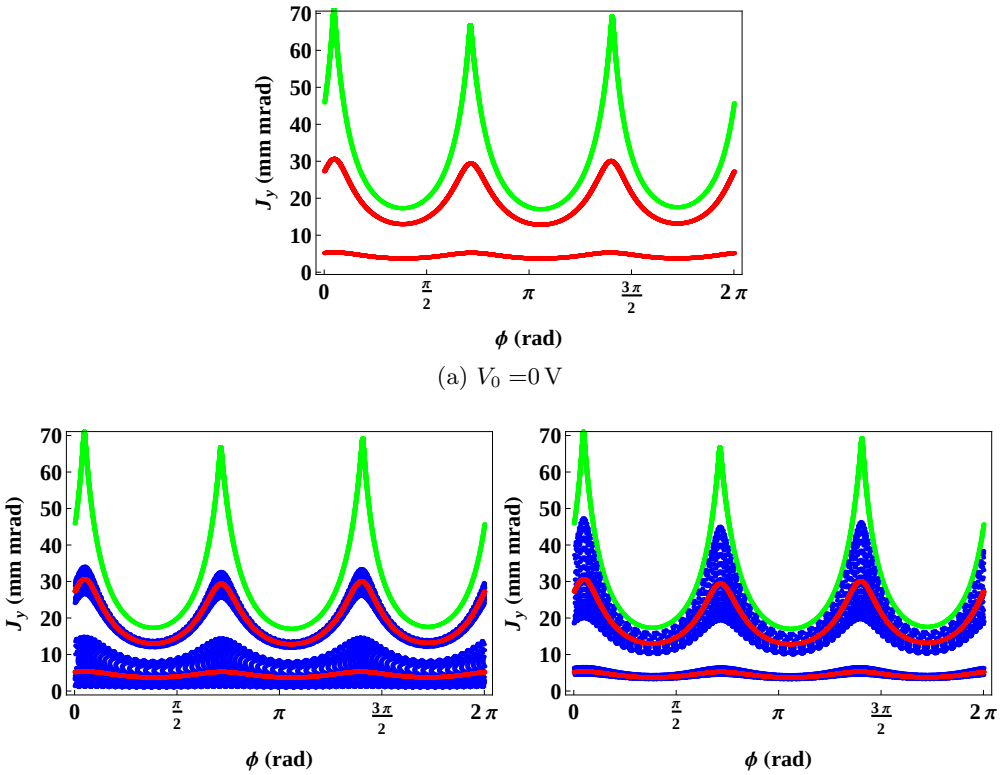

(a) V_0 =0 V

(b) V_0 =2 kV, Q_h = 1.6718

(c) V_0 =2 kV, Q_h = 1.6707

FIGURE 5.26: Plots to show the influence of the transverse deflection cavity on the amplitude of oscillation. Two particles, with starting coordinates of y = 4 mm, y' = 0 mrad and y = 7.5 mm, y' = 0 mrad, are tracked for 4000 turns. The trajectory of the particles through phase space in action angle variables at the extraction point is shown in red in (a) for the case when V_0 of the deflection cavity is set to 0 V, the separatrix is marked green. In this first plot, the variation of the action with angle is seen due to the geometric perturbation by the sextupole field (with the variation greater at larger amplitudes where the phase space trajectory tends towards being triangular). For (b) and (c), the rf voltage, V_0, has been set to 2 kV, and the phase of the rf at the start of the simulation is $\pi/4$, the rf frequencies are set according to $f_{\rm rf} = Q_y f_0$, where Q_y is the betatron tune and f_0 is the revolution frequency of a particle. The change in revolution frequency with amplitude is small compared with the tune shift with amplitude, which changes rapidly as a particle approaches the separatrix (Fig. 5.11b). In (b), $f_{\rm rf}$ has been set according to Q_y = 1.6718, which corresponds to the tune of the particle that starts at y =4 mm; the red plot shows the particle trajectories when V_0 = 0 V, whilst the blue regions indicate the variation in the action for each phase space angle for the specified rf frequency and voltage. In (c), $f_{\rm rf}$ has been set according to Q_y = 1.6707, which corresponds to the tune of the particle that starts at y =7 mm. As can be expected, the change in action over the 4000 turns is greatest for whichever particle is closest to the resonant condition; however, in both figures there is insufficient amplitude growth to cause the particle to be extracted.

Initial rf phase and septum position

The stable particles from the beam interrupt time investigation (2581 in total) are used as starting positions for a tracking simulation including kicks from a transverse deflecting rf cavity. In all cases, the perturbation voltage is set to 2 kV and the frequency to $1.67 f_{\mathrm{rev}}$, where f_{rev} is the revolution frequency of a particle with zero action in transverse phase space. The initial phase of the rf is varied from 0 radians to 1.75π radians in intervals of $0.25\,\pi$ radians. Figure 5.27 shows the starting positions of particles that are extracted during 600 turns of tracking for each of the different starting phases. It can be seen that the initial phase of the rf determines the position in phase space of the particles that are extracted first. This dependence on initial rf phase is taken as an indication of the long term behaviour of particles during the extraction process. Emittance and extraction efficiency are key metrics for the quality of extracted beam. The position of the electrostatic septum is important for both of these measures; placing the septum at a greater horizontal position increases the extraction efficiency, but also increases the emittance. The emittance and extraction efficiency are calculated for different septum positions, for each of the initial transverse rf phases. In the case of each setting, tracking of individual particles is carried out until the particles have a transverse horizontal position greater than that of the septum wire edge at the extraction location. If the particle arrives at the wire within 0.1 mm of the septum wire edge, then the particle is considered to be lost, if the particle arrives beyond 0.1 mm of the septum wire edge, then the particle is considered to be part of the extracted beam. The emittance of the extracted beam is then calculated from:

$$\epsilon_y = \sqrt{\langle y^2 \rangle \langle y'^2 \rangle - \langle y y' \rangle^2}, \tag{5.25}$$

where $\langle . \rangle$ indicates the mean over all particles. The position of the electrostatic septum is varied from 20 mm to 30 mm at intervals of 1 mm; the results for the emittance and extraction efficiency averaged over all initial phases of the rf transverse deflecting cavity are shown in Fig. 5.28. It is seen that the horizontal emittance of the extracted beam increases with transverse septum position, but remains small (of the order of 0.01 mm mrad) for all septum positions investigated, and that the extraction efficiency increases from 89% to 98% with increasing septum position. However, it is noted that whilst the emittance of the extracted beam remains small, the transverse width of the beam becomes large. Estimates of the Courant-Snyder parameters for the extracted beam are given by:

$$\alpha_y = -\frac{\langle y y' \rangle}{\epsilon_y}, \tag{5.26}$$

$$\beta_y = \frac{\langle y^2 \rangle}{\epsilon_y}, \tag{5.27}$$

$$\gamma_y = \frac{\langle y'^2 \rangle}{\epsilon_y}. \tag{5.28}$$

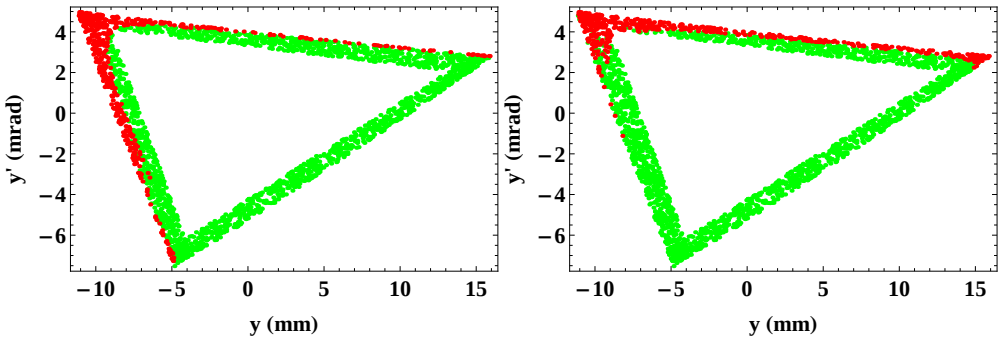

(a) $\phi_0 = 0\pi$ radians

(b) $\phi_0 = \frac{1}{4}\pi$ radians

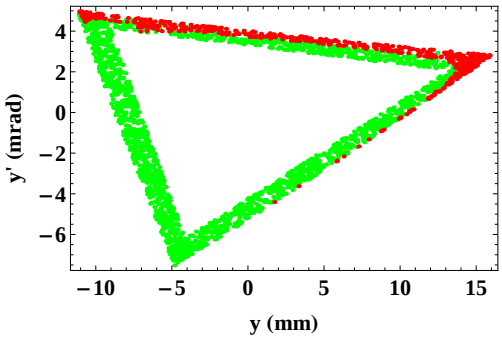

(c) $\phi_0 = \frac{1}{2}\pi$ radians

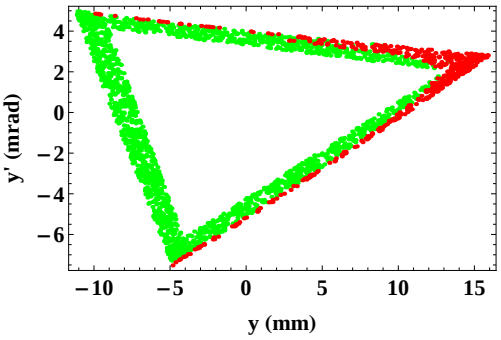

(d) $\phi_0 = \frac{3}{4}\pi$ radians

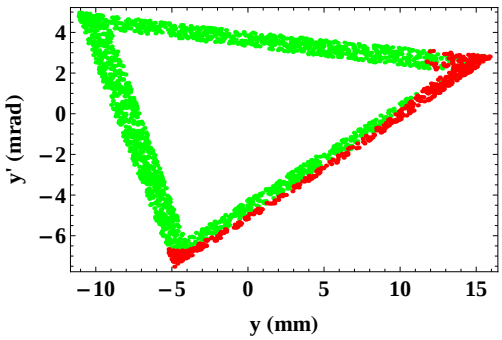

(e) $\phi_0 = \pi$ radians

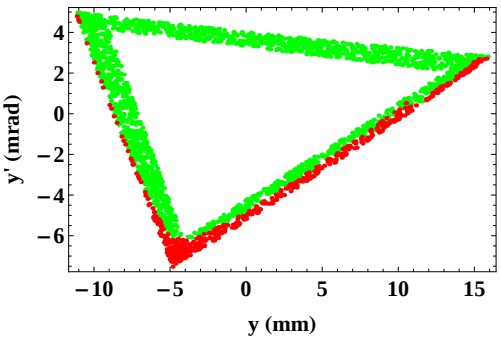

(f) $\phi_0 = \frac{5}{4}\pi$ radians

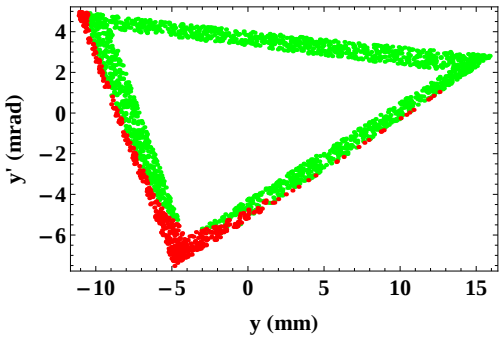

(g) $\phi_0 = \frac{3}{2}\pi$ radians

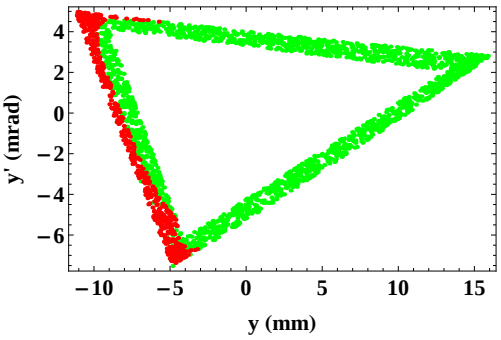

(h) $\phi_0 = \frac{7}{4}\pi$ radians

FIGURE 5.27: The initial particle coordinates are the same in each of the above plots, and are taken from the surviving particles in the beam interrupt time study. In each case, the transverse rf is set to 2 kV and a frequency of $1.67 f_{\text{rev}}$. The initial phase of the transverse rf is varied between 0 and $1.75\,\pi$ radians in intervals of $0.25\,\pi$ radians. The above plots each give the initial start coordinates of the particles in phase space, with the particles that survive 600 turns of tracking shown in green, and those that are extracted shown in red. It can be seen that the order in which particles are removed from the separatrix is dependent upon the initial phase of the rf.

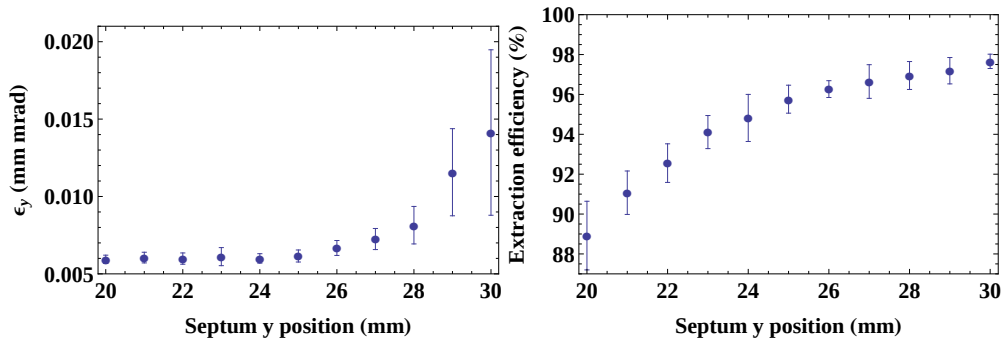

(a) (b)

FIGURE 5.28: Plots showing horizontal emittance and extraction efficiency vs. transverse septum position. The horizontal emittance and extraction efficiency presented and the associated errors are the mean and standard deviation of the values found over the range of initial rf phases simulated. Figure 5.28a shows that despite there being over a two fold increase in the horizontal emittance for the septum positions considered, the emittance remains small in comparison to the requirements of extracted beam. Figure 5.28b shows that the extraction efficiency increases from 89% to 97% (approximately) over the range of septum positions.

The combination of large transverse horizontal amplitudes and small emittances implies large horizontal β functions; however, it is expected that along the vertical axis, the emittance and Courant-Snyder parameters should be of a magnitude comparable to those found in the ring. It is anticipated that the combination of these factors could lead to challenges in the matching of a transport line to the extracted beam. One approach is to transport the extracted beam along a line that is mismatched horizontally; the rotation of the extracted beam in horizontal phase space can then be exploited to deliver a beam that is of a small size to the patient. Figure 5.29 shows the transverse horizontal phase space for three different septum positions.

A reference septum position of 25 mm is selected for further investigation, as beyond this point there is a marginal improvement in extraction efficiency for comparatively large increases in extracted beam emittance and width. The errors in Fig. 5.28 are based upon the standard deviation of the extraction emittances and efficiencies calculated for the different initial rf phases. For each initial rf phase, approximately 20% of the input particles are extracted within 600 turns, which corresponds to ≈ 500 particles on which to base an estimate of the emittance and the extraction efficiency. The errors shown in Fig. 5.28 are therefore statistical in nature rather that originating from underlying physics. With the numbers of particles used in tracking, there appears to be little dependence of extraction efficiency and emittance on initial rf phase. Therefore is is assumed that a single initial rf phase may be used to gain insight into the effects of the remaining variables to be investigated.

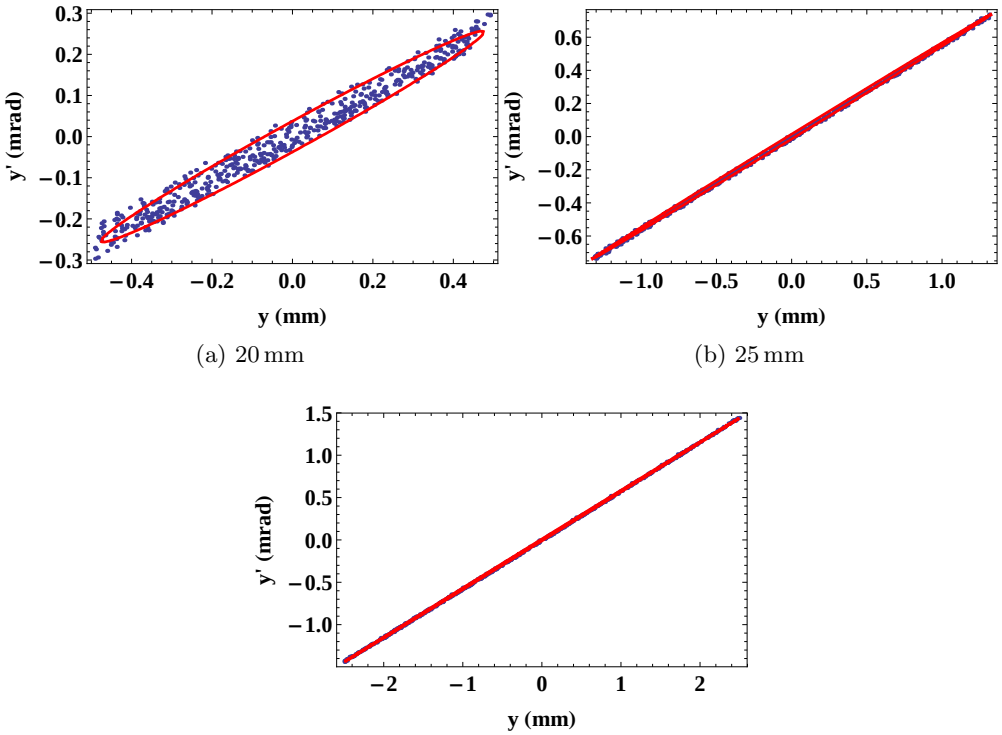

(a) 20 mm (b) 25 mm

(c) 30 mm

FIGURE 5.29: Horizontal phase space plots for the extracted beam given septum positions of 20 mm, 25 mm and 30 mm. Blue points show individual phase space coordinates of extracted particles, whilst the red ellipse has an area of $2\pi\epsilon_y$, and is based on Eqs. 5.25 and 5.26. The horizontal emittance increases as the septum position is moved away from the centre of the ring, but remains small in comparison to the extraction requirements. A small emittance combined with increasing amplitude of the extracted beam implies large β functions in horizontal phase space, which could complicate the matching of a transport line to the extraction point.

Amplitude of the rf transverse deflecting cavity voltage

Control of the intensity of the extracted beam is an important feature of the slow extraction process. By adjusting the amplitude of the electric field in the transverse deflecting cavity, control of the extracted beam intensity over very short time scales can be achieved. At the Heidelberg Ion Therapy centre a feedback loop between a beam current monitor located at the treatment nozzle and the amplitude control for the perturbing rf field allows for a beam of almost constant intensity to be delivered to a patient. Here we investigate the impact of the amplitude of the rf cavity voltage on the dynamics of the extracted particles. The investigation begins again with the particles that were found have stable trajectories in simulations of the beam interrupt time. Tracking is carried out over 600 turns in the ring for four settings of the transverse deflecting cavity voltage, 1 kV, 2 kV, 3 kV and 4 kV; the initial phase and frequency of the rf is maintained at $0.75\,\pi$ radians and $1.67\,f_{\mathrm{rev}}$ respectively. The initial phase space coordinates of the particles that are extracted within 600 turns are shown in Fig. 5.30, where it can be seen that the larger perturbation voltage causes particles that are initially further from the

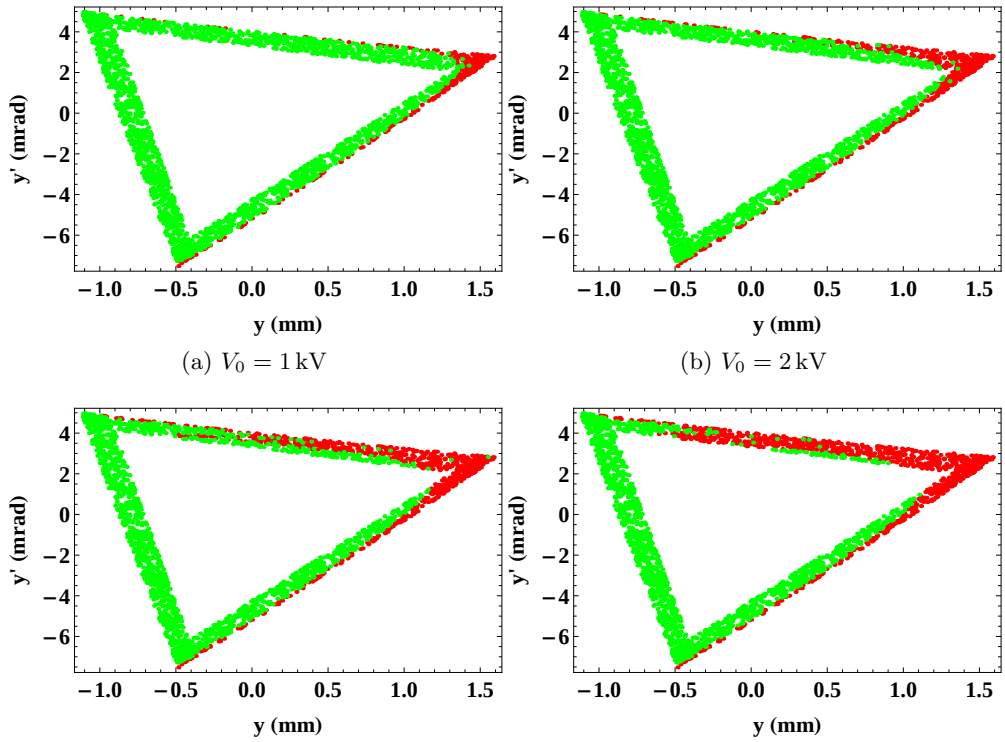

(a) $V_0 = 1\,\mathrm{kV}$ (b) $V_0 = 2\,\mathrm{kV}$

(c) $V_0 = 3\,\mathrm{kV}$ (d) $V_0 = 4\,\mathrm{kV}$

FIGURE 5.30: The amplitude of the perturbation voltage is used to control the intensity of the extracted output. In the above example, the transverse rf voltage is varied from 1 kV to 4 kV, whilst the rf frequency is fixed at $1.67\,f_{\mathrm{rev}}$ and the initial phase is $0.75\,\pi$ radians. The fraction of particles extracted during 600 turns (shown in red) increases with the amplitude of the perturbation voltage, varying between 12 % (for $V_0 = 1\,\mathrm{kV}$) and 38 % (for $V_0 = 4\,\mathrm{kV}$) of the sample particles.

separatrix to be extracted during the simulation. Within 600 turns, between 12 % (for 1 kV) and 38 % (for 4 kV) of the initial particles were extracted. It is seen in Fig. 5.31 that the magnitude of the perturbation voltage also impacts upon the emittance of the beam. Estimates of the emittance, calculated using Eq. 5.25, are in the order of 10 times larger for the beam extracted using a 4 kV perturbation than for the beam extracted using a 1 kV perturbation.

Particle momentum

Finally we consider the momentum spread of the bunch, and its impact on extraction. We use values for the fractional offset in momentum of $\delta = \pm 0.01$ and 0, and particle tracking is carried out for each setting of δ separately. To determine the starting location of particles in phase space, we use the equation for the off-momentum separatrix (Eq. 5.21) in order to locate 5000 particles around the theoretical separatrix for each δ; as with the beam interrupt time investigation (section 5.6.1), we then track these particles for 600 turns and take the particles that have not reached the extraction amplitude as being on stable orbits (in each case there are approximately 2500 stable particles). We then use the initial starting coordinates of the stable particles in a tracking study

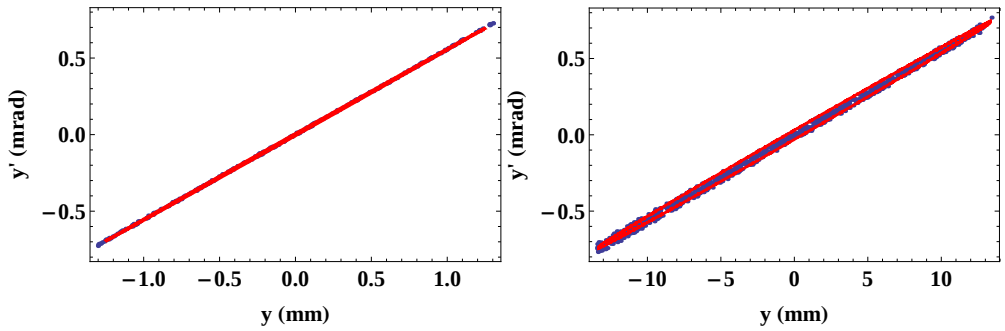

(a) $V_0 = 1\,\mathrm{kV}$ (b) $V_0 = 4\,\mathrm{kV}$

FIGURE 5.31: Increasing the magnitude of the rf perturbation leads to an increase in the width of the extracted beam. The beam extracted using a 4 kV perturbation (b) has an emittance that is to the order of 10 times greater than the beam extracted using 1 kV (a).

that includes rf perturbation: the rf voltage is set to 2 kV and the initial rf phase to $0.75\,\pi$ radians. The results of tracking show that for $\delta = \pm 0.01$ and 0 the percentage of particles extracted during 600 turns are respectively 17%, 21% and 25%. The preference for the extraction of high momentum particles shown by tracking can be explained by the negative horizontal chromaticity of the lattice; a high momentum particle that is following the closed orbit will have a tune closer to the resonant condition than a low momentum particle, this gives a smaller distance in phase space between the closed orbit and the separatrix for a high momentum particle than for a low momentum particle.

5.7 Summary and conclusions for the synchrotron design study

There are currently 56 hadron therapy facilities operating worldwide; 23 of these use synchrotrons in order to create particle beams of appropriate treatment energy, with the remaining facilities using cyclotrons [89]. A key advantage of the synchrotron is that it allows for variable beam extraction energy, which can lead to improved conformity of radiation dose to a tumour volume when compared to a cyclotron based design. A drawback of a simple synchrotron is that it has a pulsed output, which can lead to longer treatment times than for a cyclotron based system when using spot scanning to deliver dose; movement of the patient means that longer treatment times may be associated with poorer dose distribution within the tumour. Resonance based methods are used in synchrotrons to extend the extraction of beam from occurring over a single turn to occurring over many millions of turns (corresponding to a few seconds), which helps to overcome the problem of the pulsed output from a synchrotron.

In section 5.1 we introduced a simple synchrotron lattice as the basis of a study of the third-integer resonance extraction technique. In section 5.2, we manipulate the strength of the quadrupoles within the lattice in order to move the horizontal tune close to a third-integer resonance whilst keeping the vertical tune away from resonances. This

resonance is then driven by the sextupoles that are included in the lattice in section 5.3. During section 5.3, we see that families of sextupoles can be positioned and powered in order to control the lattice chromaticity and geometric perturbations independently; we offer a theory for the control of the separatrix formed in phase space during third-integer resonant extraction, which is verified through tracking studies in Zgoubi. The theory is then used to optimise the lattice for slow extraction, including controlling the chromaticity so that the Hardt condition is satisfied (which reduces particle loss during the extraction process). The final magnet configuration for the test synchrotron lattice is given in section 5.4. We then introduce the transverse deflecting cavity that is used to push particles out of the stable region of horizontal phase space (section 5.5); once outside of the stable region, particles travel along one of the arms of the separatrix that is formed when the tune is close to third-integer resonance and the motion of the particle is perturbed by a sextupole field.

Finally, in section 5.6, we simulate the extraction of the treatment beam from our optimised synchrotron. Initially, we look at the beam interrupt time (section 5.6.1), which tells us how much time passes between the deflecting cavity being powered down and no more beam being extracted from the machine; through tracking, we found an interrupt time comparable to that of an existing proton facility. In section 5.6.2 we look at the impact of the transverse position of an electrostatic septum on the extracted beam emittance, as well as the dependence of beam extraction on the initial phase and amplitude of the transverse rf perturbation; we find that the initial rf phase controls the order in which particles are extracted given different initial locations in phase space, and that the rf amplitude controls the intensity of beam extraction. The position of the electrostatic septum can be chosen in order to optimise a balance of extraction efficiency and the horizontal emittance of the extracted beam (in particular the beam width). In all cases studied, the horizontal emittance of the extracted beam is significantly smaller than is required for proton therapy (even though it increases with increasing rf amplitude).

The findings of this study into the design and performance of a proton synchrotron are used as a basis for comparison for the FFAG designs in chapter 6.

Chapter 6

FFAG Design Study

In this chapter the performance of a FFAG accelerator that was designed for the application of proton therapy will be investigated through simulation. We will look at the same performance metrics as in the case of the proton synchrotron in chapter 5. These include the energy range and control of energy variation, intensity and emittance of the extracted beam, extraction efficiency and the speed with which the beam can be turned on and off. As the magnetic fields within an FFAG are fixed during acceleration, particles across the entire momentum range of the accelerator may be following stable orbits at any given time. However, the longitudinal dynamics within proton FFAG designs are usually subject to an rf frequency that changes with time so that the orbital frequency of a design particle remains synchronised with some harmonic of the rf frequency. Without the requirement of ramping magnetic fields, the pulsed output of the FFAG can be of a much higher frequency than that of the synchrotron, which is potentially beneficial when treating patients. In this chapter, we start our investigation by looking more closely at the non-scaling FFAG PAMELA that was first introduced in section 2.9.2; in particular we will consider extraction methods that have already been proposed for this accelerator, which include kicker-based extraction as well as a method based on half-integer tune resonant extraction.

6.1 PAMELA

The PAMELA design consists of two concentric rings; we focus on the inner ring, which is used for the entire treatment energy range of protons and for low energy carbon ions (the outer ring is used to accelerate carbon ions to higher energies). For reference, we again give a table of the PAMELA lattice parameters (table 6.1); one of the key features of the PAMELA lattice is the large FFAG magnet scaling index, which results in a small orbit excursion in comparison to some other proton FFAG designs and a horizontal phase advance greater than π radians per period.

The layout of the PAMELA ring is shown in Fig. 6.1; the reference axis that is defined in the PAMELA design report is marked by a dashed black line. This reference axis locates the centre of a defocusing magnet at the midpoint of a period, then at

Table 6.1: Table of lattice and cavity parameters for the PAMELA proton ring.

PAMELA lattice	
Energy range	30.95-250 MeV
Number of cells	12
Cell type	FDF
Magnet length	0.3144 m
Magnet scaling index (k)	36.721
r_0	6.251 m
B_0 (F/D)	1.7282/-2.0 T
Horizontal tune	8.64
Vertical tune	3.24
Orbit excursion	0.176 m
rf cavities	
Harmonic number	10
Injection rf frequency	19.2 MHz
Extraction rf frequency	45.6 MHz
Number of cavities	8
Maximum voltage per cavity/turn	20/160 kV
Synchronous phase	$60° - 70°$

the end of a period the reference axis is rotated through 30°. Figure 6.2 shows the closed orbit positions found for several energies by tracking along the reference axis of the PAMELA periodic cell using Zgoubi. As with the proton synchrotron, when investigating extraction we will consider the position of particles in a plane perpendicular to the direction of the closed orbit within a long drift, for this reason it is convenient to define a new reference axis. The new reference axis is shown in Fig. 6.3, the key difference between this and the original axis is that a rotation of 15° is made at a distance of 21.90 cm from the final focusing magnet of the FDF triplet, which is followed by 129.43 cm of drift space and then another 15° rotation. As a consequence of this definition, the reference axis for the 129.43 cm drift length is parallel to the progression of the closed orbits within the drift.

Figure 6.4 shows Poincaré portraits in horizontal phase space for three beam energies that are spread across a large proportion of the PAMELA treatment energy range. In the case of the synchrotron, the design extraction energy was located on the reference axis; we took steps to ensure efficient extraction given a small momentum spread of the bunch, however the extraction process was mainly independent of the extraction energy due to the scaling of the magnetic fields during acceleration. From Fig. 6.4, we see that variable energy extraction from a FFAG may not be so straightforward due to the need to extract from different closed orbit positions. A number of kicker-based options have been considered for the PAMELA lattice, including extracting along the horizontal axis by using either a kicker that covers the full horizontal aperture of the machine or by using a kicker that can be moved across the horizontal aperture. Alternatively, particles may be extracted along the vertical axis. Tracking studies indicate that horizontal extraction

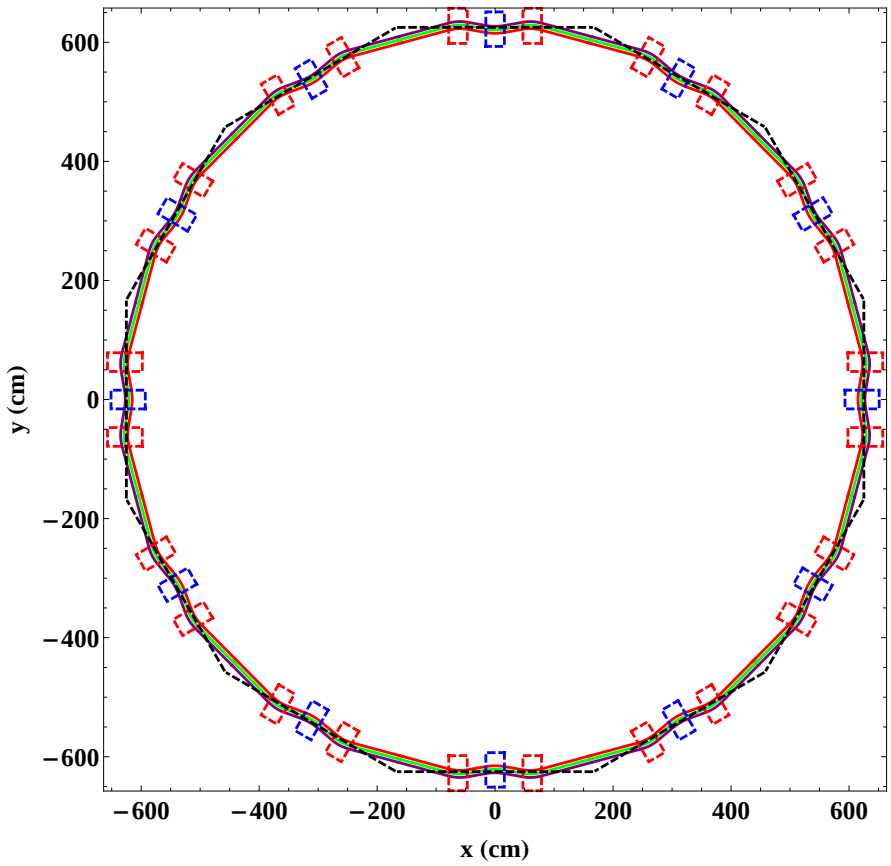

FIGURE 6.1: Diagram of the PAMELA proton ring. Red and blue dashed lines respectively show the focusing and defocusing magnets of the 12 FDF cells. Solid red, green and purple lines show the closed orbit position for 60, 120 and 230 MeV protons respectively. The large scaling index of the PAMELA magnets leads to a small orbit excursion during acceleration.

from the PAMELA lattice is unsuitable due to the challenges it poses for both hardware and beam dynamics; vertical extraction was found to be a viable alternative [29].

6.2 Resonant extraction from the PAMELA lattice

PAMELA is designed to have a repetition rate of 1 kHz. Kicker and resonant extraction methods have been considered for extraction along the vertical axis; here we focus on the resonant extraction method. The truncation of the multipole field components in the PAMELA FFAG magnets leads to a small non-zero chromaticity, furthermore the tune range of the machine can be manipulated by changing the ratio of the focusing and defocusing fields. An F/D ratio can be found which leads to the vertical betatron tune crossing a half-integer value at the desired extraction energy. Resonance crossing has been the subject of a number of FFAG studies [63, 90, 91], in the present study we aim to reproduce tracking results for the half-integer resonant extraction method already investigated for PAMELA and then compare these results with those for the proton synchrotron.

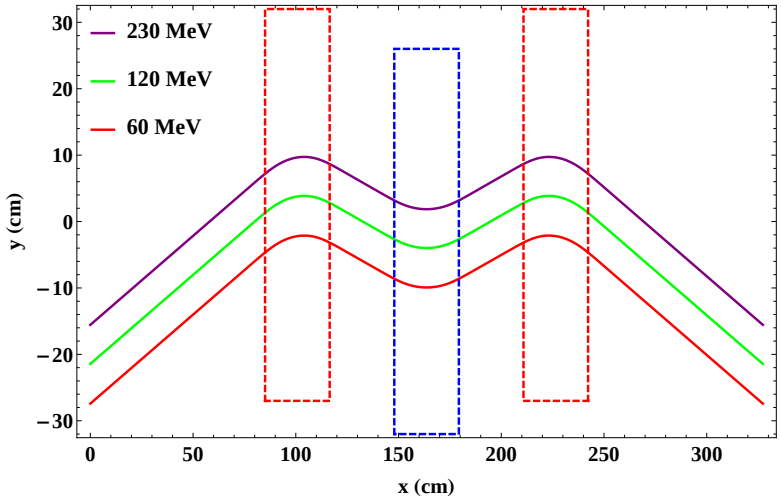

FIGURE 6.2: Plot to show the closed orbit positions (solid lines) in a single PAMELA cell for a range of particle energies. The PAMELA focusing and defocusing magnets are shown by red and blue dashed lines respectively. The design of PAMELA follows some of the principles of a radial sector FFAG, with regions of forward and backward bending; however unlike a radial sector FFAG, the magnets are rectangular. The PAMELA lattice can be constructed with a 12 period regular polygon (Fig. 6.1), with each side consisting of one of the cells shown in the plot. Each cell is ∼335 cm long: the magnets and the short drifts that separate the magnets are each 31.44 cm long, whilst the remainder of the cell is made up of the long drifts at its beginning and end. The black dashed line in Fig. 6.1 defines our initial reference axis for the cell. During the later tracking studies, we adopt a new reference axis that runs parallel to the closed orbits in the long drift section (Fig. 6.3).

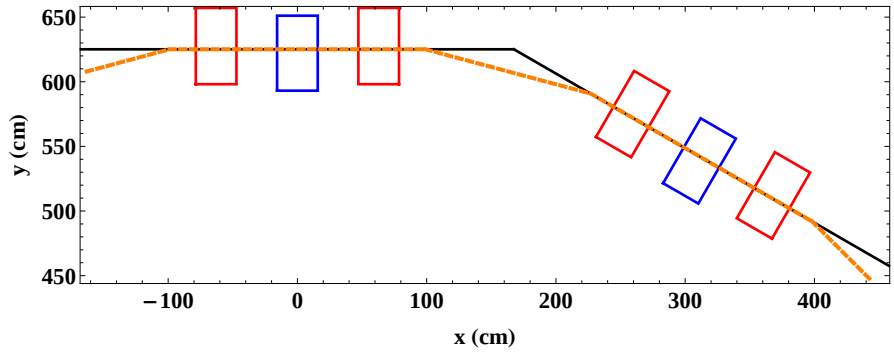

FIGURE 6.3: For the purpose of tracking, we define a new reference axis for a PAMELA period (marked by the dashed orange line). In this case a rotation of the reference axis by 15° is made at a distance of 21.90 cm from the second focusing magnet of the FDF triplet; this is followed by a 129.43 cm drift length and then a second 15° rotation to the reference trajectory. The new reference axis is parallel to the closed orbits through the drift.

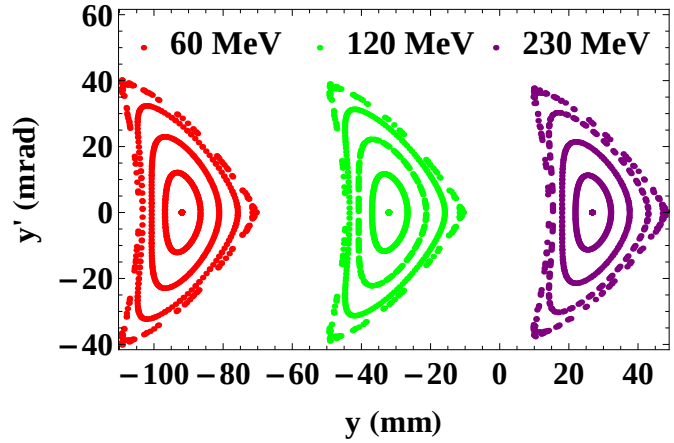

FIGURE 6.4: Poincaré portrait for three beam energies in the PAMELA lattice. Particles were given different starting y coordinates relative to the corresponding closed orbit (up to the limit of stable motion) and tracked through PAMELA using Zgoubi. The portrait is formed at the centre of the long drift that was defined in Fig. 6.3. Variable energy extraction is complicated by the need to extract from the different closed orbit positions.

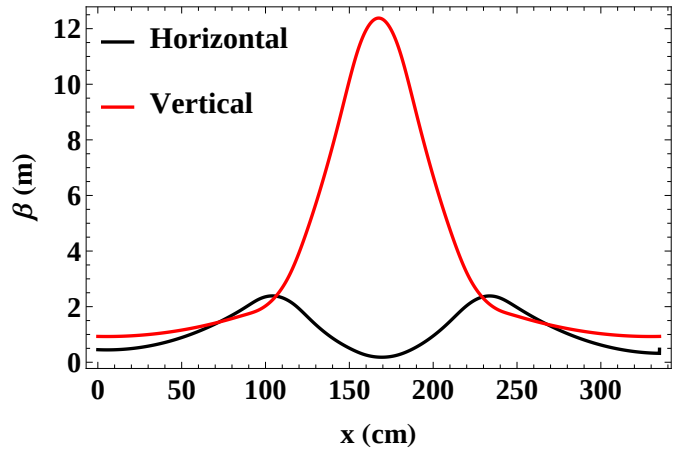

FIGURE 6.5: β functions along the length of a single PAMELA cell (centred on the defocusing magnet of the FDF structure). The near scaling design means that the β functions are independent of energy.

We start with a PAMELA model that contains field components in the magnets up to decapole. An advantage of the PAMELA magnet design is that it allows for the different field components to be controlled individually: we multiply the sextupole strength by a factor of $f_{\text{sex}} = 0.98$, which gives PAMELA a small negative chromaticity (Fig. 6.6); secondly we multiply the decapole component by a factor of $f_{\text{dec}} = 1.1$ to give a more uniform gradient of tune vs. kinetic energy (Fig. 6.7). The kinetic energy at which the vertical tune crosses the half-integer resonance can then be controlled by changing the D/F ratio of the triplet, Fig. 6.8 shows the variation in the vertical tune range when the field in the defocusing magnet at the reference radius, $B_{\text{D},0}$, is varied by $\pm 1\%$. $B_{D,0}$ can be scaled so that the particle bunch crosses a half-integer resonance close to the desired extraction energy; as an example we choose to scale $B_{D,0}$ by 0.99, for which the vertical

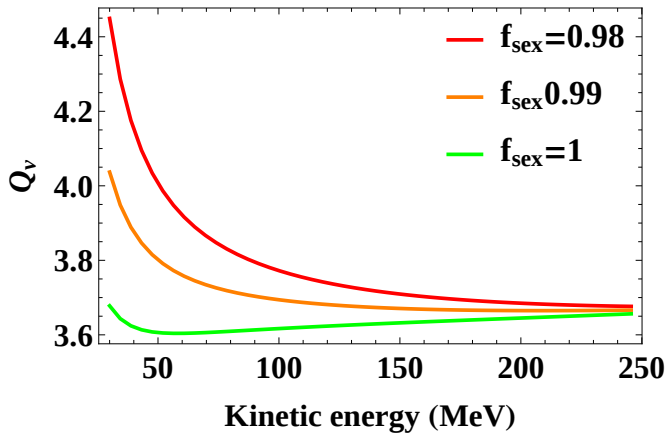

FIGURE 6.6: A small variation in the sextupole strengths of the PAMELA FFAG magnets leads to a significant change in the tune profile. For a resonant extraction method based on crossing the half-integer tune, we look for a small negative chromaticity for the vertical motion; this will allow the lattice to be set so that the half-integer tune can be crossed at an arbitrary energy, without other low order resonances being crossed before the beam reaches the extraction energy.

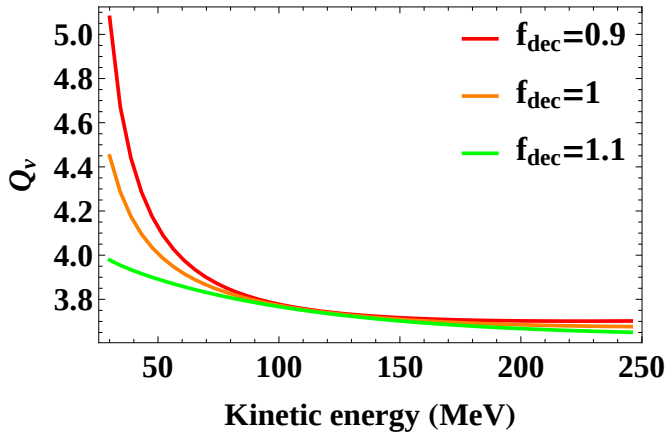

FIGURE 6.7: After multiplying the sextupole component of the Taylor series expansion of the scaling FFAG magnetic field profile by 0.99, the decapole component is then scaled to limit the tune excursion during acceleration. In the above plot we see that increasing the decapole component (where f_{dec} is the scaling factor for the decapole component) reduces the tune excursion.

tune crosses a half-integer when the bunch has kinetic energy of 180 MeV.

6.2.1 A Hamiltonian for describing systems close to resonance

In section 5.3.3 we gave a Hamiltonian that describes the motion of particles in transverse phase space when circulating in a lattice that is close to a third-integer tune and contains sextupole perturbations. The Hamiltonian was later used to optimise the extraction of particles from the lattice. In this section, we give a Hamiltonian [36, 92] that describes the motion of particles that are circulating in a lattice that is close to an arbitrary order resonance and contains multipole perturbations of an appropriate order to drive the resonance. This Hamiltonian will be used to provide an understanding of resonant

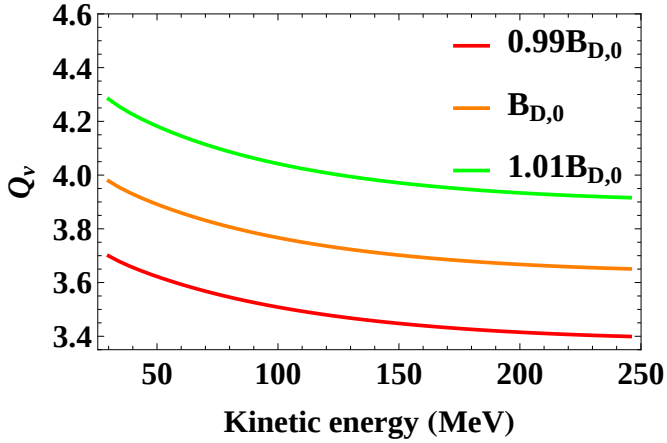

FIGURE 6.8: The above plot gives the vertical betatron tune vs. proton kinetic energy for three settings of the defocusing magnetic field at the reference radius within the PAMELA FDF triplet; $B_{\mathrm{D},0}$ is the nominal field value that is defined in table 6.1. We see that a 1% change in the reference magnetic field for the defocusing magnet leads to a change in the tune of approximately 0.3 for the entire energy range. $B_{D,0}$ can be scaled so that the particle bunch experiences a half-integer resonance at the desired extraction energy.

extraction methods investigated for FFAGs within this thesis. For one degree of freedom, the Hamiltonian that describes the motion of a particle along a transverse axis for a betatron tune that is close to the resonance $aQ_y = n$ (with integer a and n) is written as:

$$H_1 = \Delta J_{1y} + \alpha(J_{1y}) + V_y, \tag{6.1}$$

where J_{1y} and ϕ_{1y} are a transformation of the standard action-angle variables $(J_y, \phi_y) \mapsto (J_{y1}, \phi_{y1})$:

$$J_{1y} = J_y,$$

$$\phi_{1y} = \phi_y - \frac{n\theta}{a}.$$

$\Delta = Q_y - \frac{n}{a}$ gives the difference between the tune of a particle with zero action and the resonance condition for the unperturbed system. The term θ represents the azimuthal angle in the plane of the circular accelerator, and replaces the closed orbit path length, s, as the independent variable for the Hamiltonian. The function α, which is dependent upon the action J_y, is based upon the detuning with amplitude along the given axis. As an initial example, we consider the tracking results produced in chapter 5 when we first moved the synchrotron lattice close to a third-integer resonance and included a single sextupole perturbation (Fig. 5.13), in this case, $\Delta = 1.33 \times 10^{-3}$ and the sextupole perturbation is given by:

$$S_y = \frac{\sqrt{2}}{24\pi}\beta_y^{3/2} K_2 L,$$
$$= 0.335\,\mathrm{m}^{-1/2}.$$

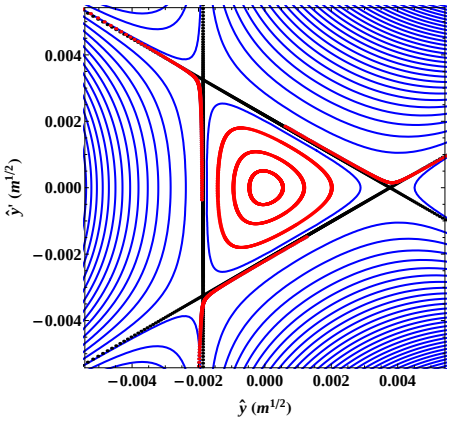

FIGURE 6.9: Plot showing contours of constant Hamiltonian for the system described by Eq. 6.2 (blue lines) in normalised transverse phase space, with the separatrix marked by the black points. We see that the tracking results that were originally presented in Fig. 5.13 (marked in red) are in good agreement with the new Hamiltonian.

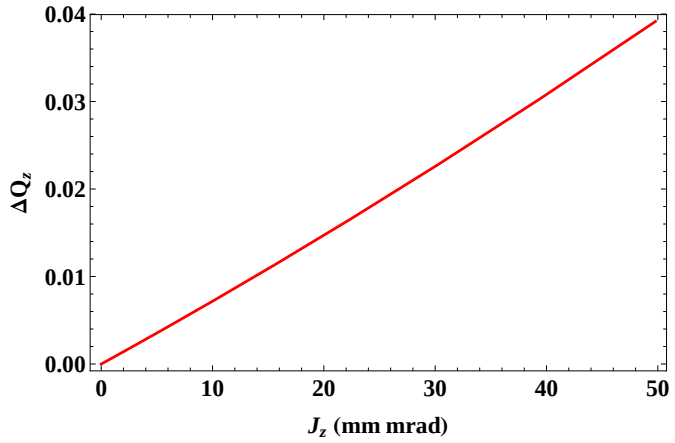

FIGURE 6.10: Vertical tune shift with action for the PAMELA lattice. The non-linear fields within the near-scaling FFAG magnet with a large scaling factor lead to a significant tune shift with amplitude.

As the unperturbed system is linear, we take $\alpha \approx 0$, which gives us the Hamiltonian:

$$H_1 = 1.33 \times 10^{-3} J_{1y} + 0.335 J_{1y}^{3/2} \cos(3\phi_{1y}). \tag{6.2}$$

Plotting the contours of constant H_1 in normalised phase space (Fig. 6.9) shows that we obtain a separatrix that is in good agreement with that found by tracking particles with Zgoubi. In the case of the PAMELA lattice, there is significant amplitude detuning (Fig. 6.10) due to the large FFAG scaling index, k, and the resulting contribution of non-linear fields (mainly of the sextupole and octupole components) to the dynamics of particles within the ring. Through the application of Hamilton's equation to Eq. 6.1, we obtain:

$$\frac{d\phi_{1y}}{d\theta} = \frac{\partial H_1}{\partial J_{1y}} = \Delta + \frac{\partial \alpha(J_{1y})}{\partial J_{1y}} + \frac{\partial V}{\partial J_{1y}}, \tag{6.3}$$

and it follows that $\frac{\partial\alpha(J_{1y})}{\partial J_{1y}}$ describes the tune shift with amplitude for the unperturbed system. If we fit an r^{th} order polynomial to a plot of the difference in tune between a particle of action J_y and a particle at the closed orbit vs. action, then we get:

$$\frac{\partial\alpha(J_{1y})}{\partial J_{1y}} = \alpha_1 J_{1y} + ... + \alpha_r J_{1y}^r. \tag{6.4}$$

The function $\alpha(J_{1y})$ is then found by integrating Eq. 6.4 with respect to J_{1y}. For the PAMELA lattice, we fit a second order polynomial to tracking data in order to give the curve shown in Fig. 6.10.

We will drive the half-integer resonant extraction by including a quadrupole field error within the lattice; the vector potential for a quadrupole magnet is:

$$A_x = \frac{K_1}{2}\left(y^2 - z^2\right), \tag{6.5}$$

by writing y and z in action angle variables, and then expanding the trigonometric terms, we get:

$$A_x = \frac{K_1}{2}(\beta_y J_{1y} + \beta_y J_{1y}\cos(2\phi_{1y}) - \beta_z J_{1z} - \beta_z J_{1z}\cos(2\phi_{1z})). \tag{6.6}$$

We are interested in the terms that vary slowly with increasing turn number as these are the perturbations that add coherently over a large number of turns and can lead to amplitude growth. For motion along the vertical axis when the vertical tune is close to half-integer, this includes both terms in Eq. 6.6 that are dependent on J_{1z}. If we were instead looking at motion along the horizontal axis, which is away from a half-integer tune, it is the single term that is dependent on J_{1y} and independent of ϕ_{1y} that would be of importance (which represents the horizontal tune shift due to the perturbation). When including the relevant perturbation terms, the Hamiltonian is given by:

$$H_1 = \Delta J_{1z} + \frac{1}{2}\alpha_1 J_{1z}^2 + \frac{1}{3}\alpha_2 J_{1z}^3 - \frac{1}{4\pi}\bar{\beta}_z K_1 L J_{1z}(1 + \cos(2\phi_{1z})), \tag{6.7}$$

where $K_1 L$ represents the integrated strength of the quadrupole perturbation. The quadrupole perturbation is included within the PAMELA lattice by scaling the quadrupole term of the Taylor series expansion of the magnetic field around the reference radius for the horizontally defocusing magnet within a single FDF cell, i.e.:

$$B_z = B_{z,0}\left(1 + f_{\text{quad}}\frac{k}{R_0}\Delta R + \frac{(k-1)k}{2R_0^2}\Delta R^2 + ...\right), \tag{6.8}$$

where f_{quad} is the scaling factor; this scaling factor effectively introduces a dipole field that increases with radius, and a fixed quadrupole field. We initially set a scaling factor of $f_{\text{quad}} = 1.003$, which gives a quadrupole error strength of:

$$K_1 = \frac{0.003 B_{z,0} k}{B\rho R_0}. \tag{6.9}$$

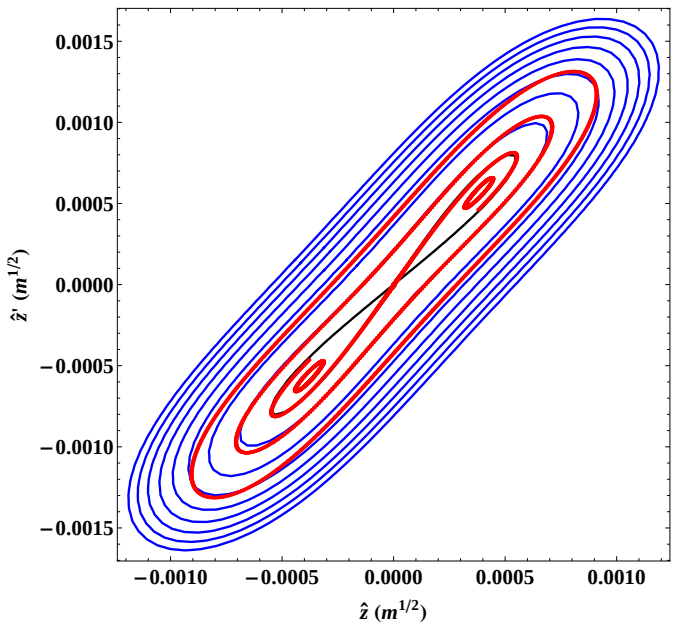

FIGURE 6.11: Poincaré plot showing the vertical phase space trajectories for 180 MeV protons at the mid point of a long drift in the PAMELA lattice, the multipole components of the FFAG magnets have been set so that the betatron tune crosses a half-integer resonance at a proton energy of 180 MeV. The quadrupole component of the defocusing magnet within a single FDF cell has been scaled by $f_{\mathrm{quad}} = 1.003$ in order to perturb the beam. In the above, the trajectories found by tracking particles through the lattice with Zgoubi for 800 turns are marked in red, whilst the blue lines show contours of constant Hamiltonian (Eq. 6.7), and the black line shows the separatrix.

Using Zgoubi tracking data for the unperturbed system, we calculate a mean vertical β function of $\bar{\beta}_z = 11.49\,\mathrm{m}$ in the horizontally defocusing magnet, and then in turn use this to calculate the perturbation strength for a 180 MeV particle.

Figure 6.11 compares the phase space trajectories obtained through tracking particles with Zgoubi and by plotting lines of constant Hamiltonian (Eq. 6.7); we note that for the conditions shown, evaluating the Hamiltonian gives a good qualitative description of the particle motion found through tracking. At this stage we make the assumption that the Hamiltonian will also offer a reasonable description of similar systems (e.g. with different perturbation strengths and differences between the zero action particle tune and the resonant tune). For reference, we include a Poincaré plot (Fig. 6.12) that is for a system that is identical to that of Fig. 6.11 apart from the amplitude detuning, which in this case is set to zero for the unperturbed system ($\alpha(J_{1z}) = 0$). We see that for zero amplitude detuning, the contours do not close within the plot range shown; amplitude detuning limits the oscillation amplitude growth of a particle that is close to the half-integer resonance. In Fig. 6.11 we can identify two stable fixed points; the conditions for a fixed point are:

$$\frac{\partial H_1}{\partial \phi_{1z}} = \frac{1}{2\pi}\bar{\beta}_z K_1 L J_{1z} \sin(2\phi_{1z}) = 0, \tag{6.10}$$

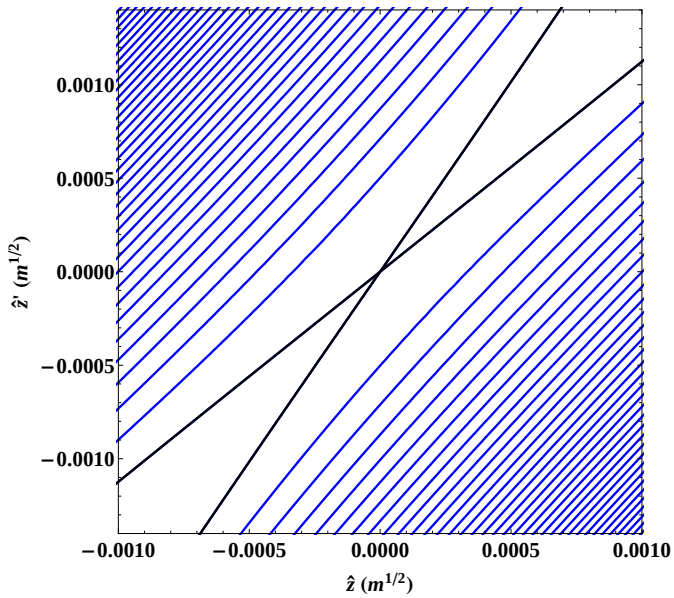

FIGURE 6.12: Plot showing contours of constant Hamiltonian in the case that there is zero amplitude detuning ($\alpha(J_{1z}) = 0$), but with otherwise the same conditions as in Fig. 6.11. Over the range shown, the contours are no longer closed; amplitude detuning limits the total amplitude growth exhibited by a particle that is close to the half-integer resonance.

$$\frac{\partial H_1}{\partial J_{z1}} = \Delta + \frac{\partial \alpha(J_{1z})}{\partial J_{1z}} - \frac{1}{4\pi}\bar{\beta}_z K_1 L(1 + \cos(2\phi_{1z})) = 0. \tag{6.11}$$

From Eqs. 6.10 and 6.11 we see that fixed points may be found at intervals of $\pi/2$ (i.e. $\phi_{1z} = 0, \pi/2, ...$) with:

$$\frac{\partial \alpha(J_{1z})}{\partial J_{1z}} = \frac{1}{4\pi}\bar{\beta}_z K_1 L(1 + \cos(2\phi_{1z})) - \Delta. \tag{6.12}$$

After setting ϕ_{1z} to 0 and $\pi/2$ in Eq. 6.12, we find that the cosine term must equal ± 1. In the case that $\cos(2\phi_{1z}) = -1$, the quadrupole perturbation term in Eq. 6.12 reduces to zero, a fixed point is then found when:

$$\frac{\partial \alpha(J_{1z})}{\partial J_{1z}} = -\Delta\Big|_{\phi_{1z}=n\pi/2}, \tag{6.13}$$

where n is an odd integer. Remembering that we set the unperturbed PAMELA lattice so that it crosses a half-integer tune at a particle energy of approximately 180 MeV, we find through tracking that the exact distance of the tune from resonance is $\Delta \approx -1.81 \times 10^{-4}$ in our initial example. Equation 6.13 is satisfied when there is a small contribution from amplitude detuning. The fixed points found when ϕ_{1z} is an odd multiple of $\pi/2$ correspond to the stable fixed points that are seen in Fig. 6.11. As Δ decreases, the stable fixed points occur at larger action of the particle oscillation; however, given a positive tune shift with amplitude, there are no solutions for the stable fixed points when $\Delta > 0$.

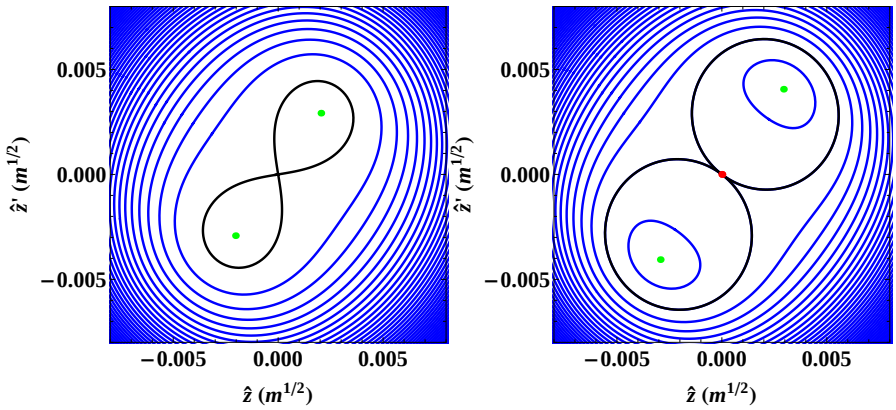

(a) $\Delta = \frac{1}{4\pi}\bar{\beta}_z K_1 L$ (b) $\Delta = \frac{1}{2\pi}\bar{\beta}_z K_1 L$

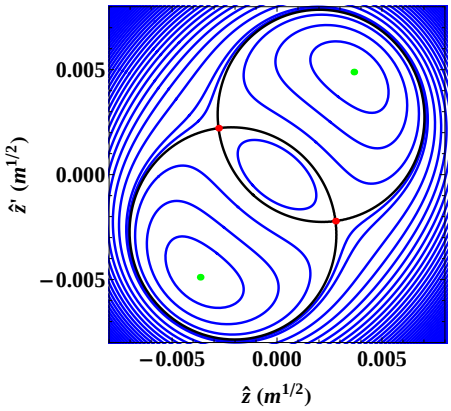

(c) $\Delta = \frac{3}{4\pi}\bar{\beta}_z K_1 L$

FIGURE 6.13: The above Poincaré plots are based upon the tune shift with amplitude $\frac{\partial\alpha(J_{1z})}{\partial J_{1z}}$ being positive, and the distance from resonance Δ and quadrupole strength K_1 being negative. In Fig. (a) $\Delta > \frac{1}{2\pi}\bar{\beta}_z K_1 L$ and there are no solutions for the unstable fixed point (Eq. 6.14). In Fig. (b) we see that if $\Delta = \frac{1}{2\pi}\bar{\beta}_z K_1 L$, then there is an unstable fixed point (marked red) at the origin of the plot. As Δ becomes more negative, then the unstable fixed points are found at greater action, and a new region of stability opens at the centre of the plot (Fig. (c)). The stable fixed points, marked green, are found at greater action as Δ becomes more negative.

When $\cos(2\phi_{1z}) = 1$, the perturbation term from the quadrupole is maximised, giving:

$$\frac{\partial\alpha(J_{1z})}{\partial J_{1z}} = \frac{1}{2\pi}\bar{\beta}_z K_1 L - \Delta\Big|_{\phi_{1z}=n\pi/2}, \tag{6.14}$$

where n is an even integer. Given a positive tune shift with amplitude and a negative value of K_1, then there are no solutions to Eq. 6.14 unless:

$$\Delta \leq \frac{1}{2\pi}\bar{\beta}_z K_1 L. \tag{6.15}$$

In Fig. 6.13 we see the evolution of the Poincaré portraits as Δ becomes more negative; in particular we note that the stable and unstable fixed points occur at larger particle actions and that a new region of stable motion appears.

During the half-integer resonance extraction scheme, the particles oscillate around the two stable fixed points given by Eq. 6.13; at this stage we can begin to identify some further requirements of this specific extraction technique. Firstly, the increase in oscillation amplitude exactly at the resonance condition is limited due to amplitude detuning; extraction requires that the particle bunch continues to accelerate, Δ will become more negative due to the negative lattice chromaticity, and the particles will oscillate around stable fixed points that occur at larger amplitudes with respect to the unperturbed closed orbit. Ultimately, the stable fixed point locations should lead to particle oscillations that cross the wire of an electrostatic septum. Secondly, we saw in Fig. 6.13 that as Δ became more negative, a new stable region of phase space opened at the centre of the Poincaré plot. The value of Δ at which this stable region opens is dependent upon the strength of the quadrupole perturbation; for good extraction efficiency, we want a perturbation strong enough to ensure that the stable region does not open during extraction. In the next section we will look more closely at this extraction method by using the Zgoubi tracking code.

6.3 A Zgoubi model of the half-integer extraction process

In the previous section we found that the action at which stable fixed points are located is dependent upon the rate of amplitude detuning and the distance of an unperturbed particle from the resonance condition (Eq. 6.13). Before this, we had tuned the PAMELA lattice so that particles crossed the half-integer tune at 180 MeV and so that other integer or half-integer resonances would not be crossed during an acceleration cycle for either the horizontal or vertical axis. We will consider an extraction plane at the midpoint of the long drift within the PAMELA cell; given an FDF cell structure, the vertical β function is a minimum at this position within the long drift, which may not be optimal for extraction efficiency. We also note that a non-linear relationship between tune and kinetic energy means that the gradient, $\frac{d\Delta}{dE_k}$, is small at 180 MeV (Fig. 6.14); this, combined with the relatively small β function, will result in the transverse position of the stable fixed points changing slowly with acceleration at the extraction point. For the purpose of our investigation, we select a new scaling coefficient for the magnetic field at the reference radius of the defocusing magnets, which leads to the vertical tune crossing a half-integer value at 50 MeV; the required scaling constant for this is $f_{\mathrm{B_{D,0}}} = 0.990$ (Fig. 6.15 shows the modified tune profile). We will start the tracking study with a scaling error on a single defocusing quadrupole of $f_{\mathrm{quad}} = 1.005$. A total of 216 particles are given an initial kinetic energy of 30 MeV, and are distributed uniformly over a circle of geometric emittance 0.8 mm mrad in the normalised vertical phase space. Each particle is attributed a weighting coefficient based upon a Gaussian distribution of particles along each transverse axis (Fig. 6.16). We will look at how the rate of acceleration affects the resonant extraction of particles from the ring. The acceleration specifications for PAMELA give a maximum cavity voltage of 20 kV (a total voltage of 160 kV per

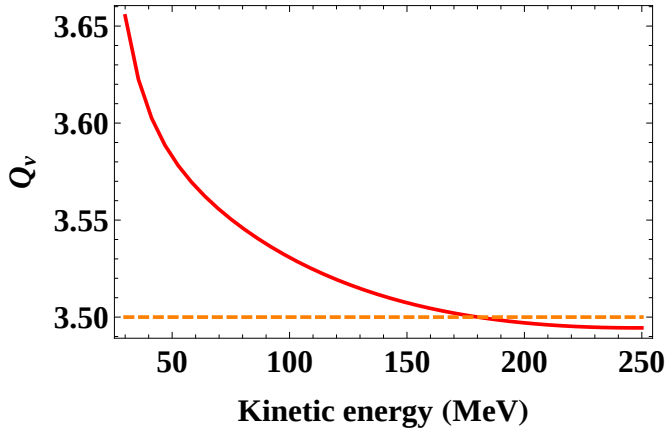

FIGURE 6.14: The PAMELA lattice is initially set so that the vertical betatron tune (red line) crosses a half-integer value (orange dashed line) when the kinetic energy is 180 MeV, and so that no other integer or half-integer resonances are crossed during an acceleration cycle. Achieving a negative chromaticity and limiting the range of the tune profile required scaling of the sextupole and decapole components obtained through the Taylor series expansion of the scaling FFAG law by $f_{\text{sex}} = 0.99$ and $f_{\text{dec}} = 1.05$, respectively. In order to have the tune cross a half-integer at 180 MeV, the field at the reference radius of the defocusing magnet was scaled by $f_{\text{B}_{\text{D},0}} = 0.994$.

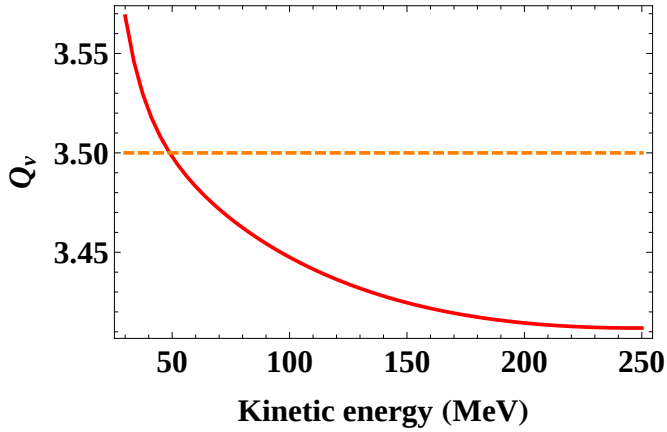

FIGURE 6.15: The PAMELA lattice is modified so that the vertical betatron tune (red line) crosses a half-integer value (orange dashed line) at a particle energy of approximately 50 MeV. To achieve this, the field at the reference radius of the defocusing magnet was scaled by $f_{\text{B}_{\text{D},0}} = 0.990$.

turn), and a synchronous phase of 60-70°. We model acceleration within this study as a fixed increment of the particle energy over each turn, which is dependent on the quoted total voltage per turn and a synchronous phase of 60° (i.e. $\Delta E_k = V_0 \sin(\pi/3)$). The acceleration rates used are based on voltages of 40 kV, 80 kV, 120 kV and 160 kV per turn. For each acceleration rate we consider a range of locations along the z axis for an electrostatic septum, any particle that reaches a transverse position beyond the septum location is defined as having been extracted. Figure 6.17 shows the mean and standard deviation of the extracted particle energies for the four rates of acceleration, we see that the mean and standard deviation are lowest for the 40 kV cavity voltage and when the septum wire position is closest to the reference axis. Ideally, it should

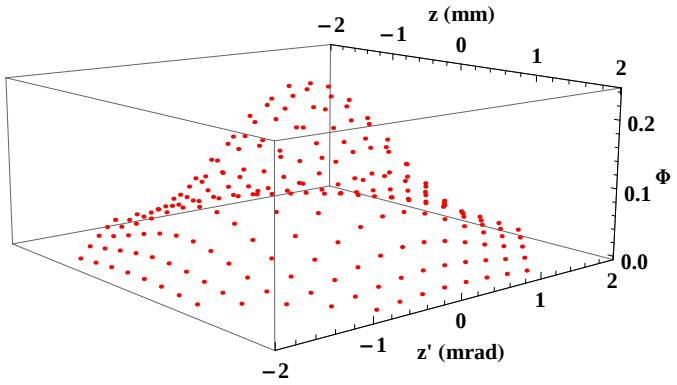

FIGURE 6.16: For the purpose of the PAMELA tracking study particles are initially distributed uniformly within a circle of area $1.5\,\pi$ mm mrad in vertical normalised phase space. A weighting coefficient is attributed to each of the 216 particles, which is based on the particles having a Gaussian distribution along each of the transverse axes. The above plot shows the particle distribution in the non-normalised z, z' coordinates with their weighting coefficients, Φ.

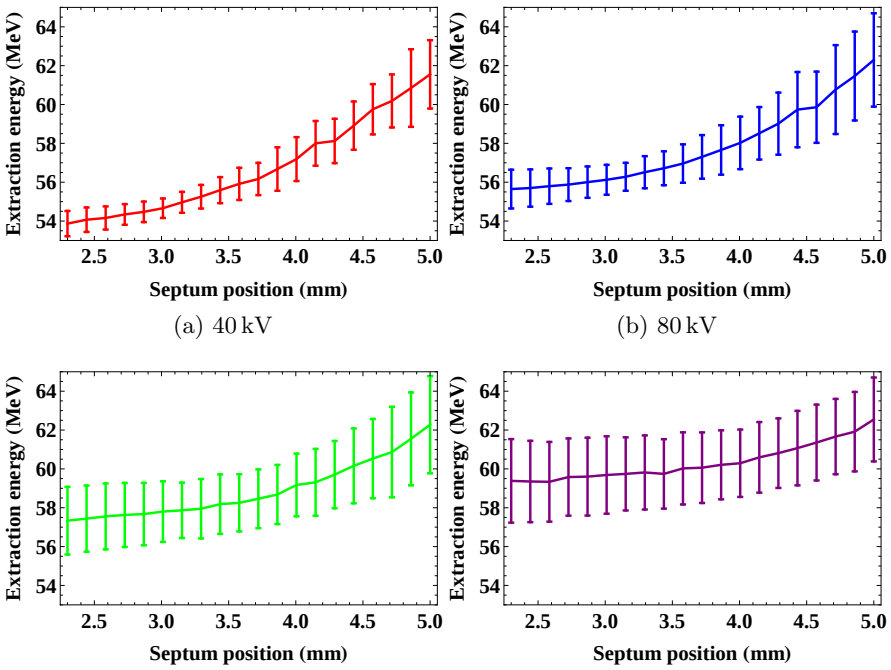

(a) 40 kV (b) 80 kV

(c) 120 kV (d) 160 kV

FIGURE 6.17: Plot to show the mean extraction energy vs. transverse position of the electrostatic septum for total cavity voltages per turn of 40 kV, 80 kV, 120 kV and 160 kV. We note that the mean extraction energy and the standard deviation of the extraction energy for lower septum position values increase with acceleration rate.

be possible to extract a beam for proton therapy that has a fractional energy spread of no more than ±0.1%; in Fig. 6.17 we see a minimum standard deviation of the extracted energy of $\sigma_{E_k} \approx 0.6$ MeV (1.1% of the mean energy), which falls short of the minimum standard. To estimate the efficiency of extraction we look for the proportion of particles that will cross an electrostatic septum wire of 0.1 mm diameter without impacting upon it (i.e. using the same method as was applied to the proton synchrotron

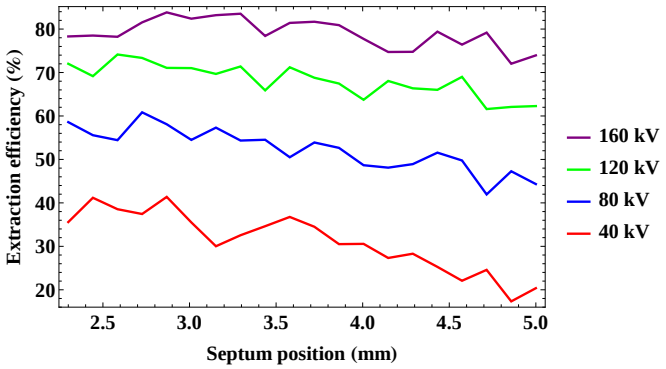

FIGURE 6.18: Extraction efficiency vs. transverse position of the electrostatic septum wire for cavity voltages of 40 kV, 80 kV, 120 kV and 160 kV. The extraction efficiency improves significantly for greater rates of acceleration.

in section 5.6.1). Figure 6.18 shows the extraction efficiency vs. septum wire position along the transverse axis for the four acceleration rates, we note that the efficiency increases with increasing acceleration rate and decreases for greater displacement of the septum wire from the reference axis. In chapter 5, the extraction efficiency for the proton synchrotron was estimated to be up to 98%; the tracking study carried out for extraction from the PAMELA lattice using a half-integer resonance suggests an efficiency of up to 85%, however, this could potentially be improved by further optimisation of the lattice (e.g. by increasing the vertical β function at the extraction location).

Figure 6.19 demonstrates the effect of varying the perturbing field error on the extraction energy and efficiency. We apply perturbations of $\Delta B = 0.005, 0.01$ and 0.02 to simulations in which the acceleration voltage is 80 kV and 120 kV; for the simulations carried out, increasing the perturbation strength generally increases the mean extraction energy, but the standard deviation of the extraction energies and extraction efficiency do not change significantly.

6.4 Dynamics of the half-integer resonant extraction process

Figure 6.20a shows the position of particles along the z axis vs. kinetic energy during tracking when the cavity voltages are 40 kV and 160 kV, and the field error is given by $\Delta B = 0.005$; it can be seen that just beyond the the resonant energy (at ~ 55 MeV), the width of the beam appears larger for the acceleration rate based on a 40 kV accelerating voltage. The situation is demonstrated more clearly in Fig. 6.20b, where we show the mean of the absolute z coordinates for the particles in the bunch ($\langle |z| \rangle$) vs. kinetic energy; this should approximately give the amplitude of the centre of mass of the particles local to one of the phase space islands once the beam has split. As the bunch energy crosses the resonance energy, particles begin to orbit around the new stable fixed points, initially particles roughly follow the separatrix, which itself encloses an increasing area due to the acceleration of the bunch and the subsequent change of Δ; this is seen as the sharp rise

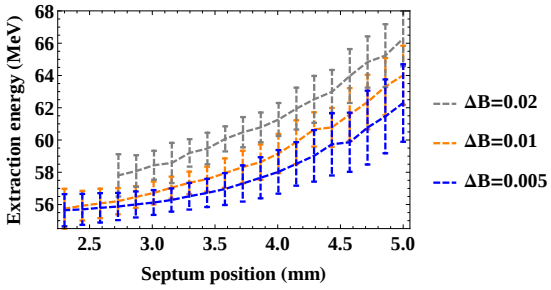

(a) Extraction energy (80 kV)

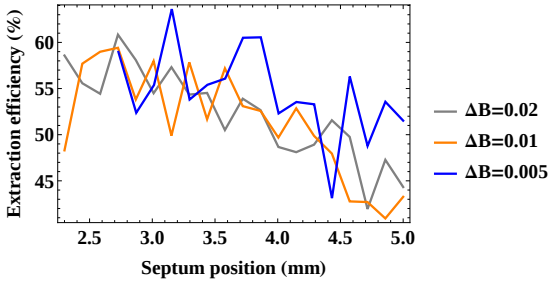

(b) Extraction efficiency (80 kV)

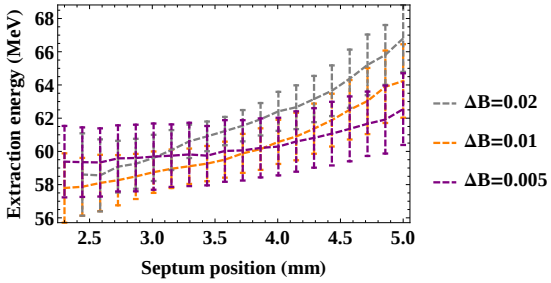

(c) Extraction energy (160 kV)

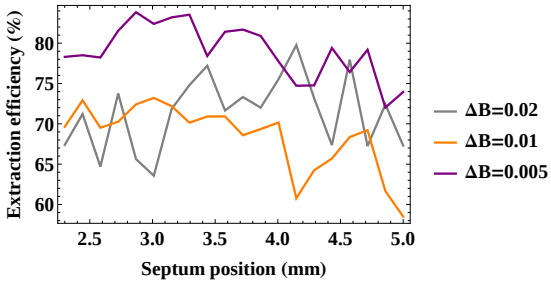

(d) Extraction efficiency (160 kV)

FIGURE 6.19: Extraction energies and efficiencies vs. septum position for total cavity voltages of 80 kV (Figs. (a) and (b)) and 160 kV (Figs. (c) and (d)). For each acceleration rate, data are presented for quadrupole perturbations of $\Delta B = 0.005, 0.01$ and 0.02. A larger perturbation tends to lead to extraction occurring at a higher mean energy, however there is no significant difference in the extraction efficiencies for the perturbations and acceleration rates considered.

of $\langle|z|\rangle$ in Fig. 6.20b shortly after the beam crosses the resonant energy. The acceleration of the beam then causes the particles to become trapped within two islands of phase space which are located around the stable fixed points, these islands move out to greater amplitudes with the stable fixed points as the beam is accelerated further. This is seen in Fig. 6.20b, beyond the initial sharp rise of $\langle|z|\rangle$, as oscillatory motion that approximately follows the curve for the stable fixed point. In particular we note that $\langle|z|\rangle$ follows the stable fixed point more closely from a lower energy for the lower rate of acceleration, which is consistent with the lower mean extraction energy for the lower acceleration voltages. Figure 6.20c shows $\langle|z|\rangle$ vs. turn number, with turn zero defined as the turn at which the bunch crosses the resonant energy, the rate of amplitude growth is greater for the higher acceleration rate; this is consistent with the extraction efficiencies shown in Fig. 6.18.

6.4.1 Flat top acceleration scheme

The lag between the increase in amplitude of the stable fixed point and the initial blow-up of the beam offers a route for reducing the energy spread of the extracted beam; for the 160 kV accelerating voltage in Fig. 6.20b, $\langle|z|\rangle$ has not significantly increased by the time the beam has been accelerated to 55 MeV, however, the amplitude of the stable fixed point has exceeded the transverse position of the septum wire. If the acceleration voltage is initially 160 kV, but is then set to zero once the beam has reached 55 MeV, particles will continue to move close to the separatrix in the vertical transverse phase space and will ultimately encounter the septum wire. In this case, the energy spread of the extracted bunch will depend upon the energy spread of the bunch at injection and the longitudinal dynamics during acceleration, rather than being introduced by the extraction process. We use the same initial distribution of 216 particles as earlier, and track the particles with Zgoubi for the new acceleration regime; the extraction efficiencies found through tracking in this case, $\sim 85\%$ (Fig. 6.21), are comparable to those found when an accelerating voltage of 160 kV is applied continuously.

6.5 Summary and conclusions for the FFAG design study

One of the key differences between the half-integer resonant extraction applied here and the third-integer resonant extraction applied to the proton synchrotron is the duration of extraction. For the tracking study based on the coasting beam, all particles were extracted within 70 turns following the rf being switched off (Fig. 6.22), this corresponds to extraction being completed within $\sim 28\,\mu$s (and with a non-uniform time profile) for the 55 MeV beam, which is short in comparison with a spill duration of several seconds from the synchrotron; the beam extracted from the FFAG using the half-integer resonance is appropriate for use in the same way as the beam extracted by a kicker (i.e. delivering dose to voxels on a shot by shot basis).

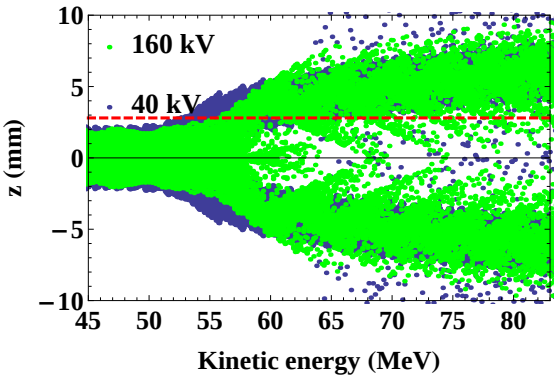

(a) Position of particles along transverse vertical axis vs. kinetic energy.

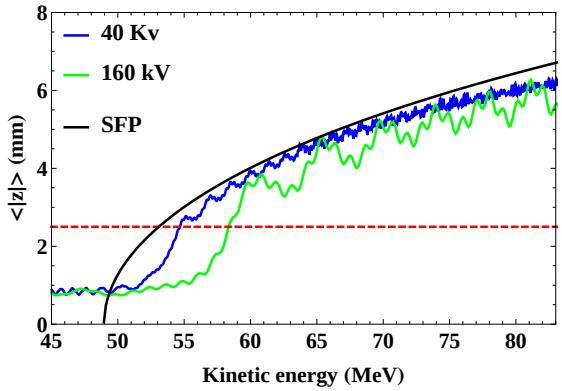

(b) Mean absolute z position vs. kinetic energy.

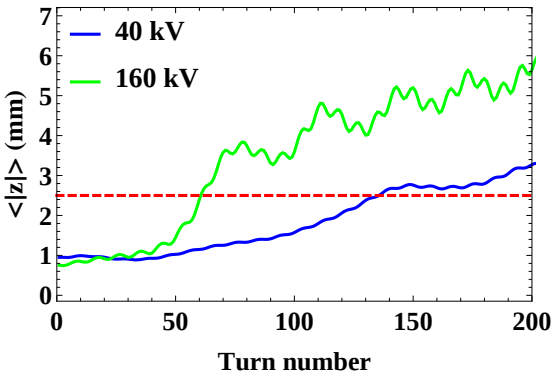

(c) Mean absolute z position vs. turn number.

FIGURE 6.20: In Fig. 6.17 we saw that the mean extraction energy increased with acceleration rate. The above plots show in more detail how the transverse vertical position of particles evolves during resonance crossing. Figure (a) shows the positions of all 216 particles within the tracking study along the transverse vertical (z) axis vs. kinetic energy. Figure (b) shows the mean absolute z position for the 216 particles (given acceleration rates of 40 kV and 160 kV) and for the stable fixed points vs. kinetic energy. It can be seen that, on average, the particles are in closer proximity to the stable fixed points from lower energies for the lower acceleration rate than for the higher acceleration rate. Finally, in Fig. (c), which shows the mean absolute z position of particles vs. turn number (again for the two rates of acceleration), we see that the mean rate of amplitude growth is higher for the greater rate of acceleration. It is the increased rate of amplitude growth that is important for good extraction efficiency. For the greater rate of acceleration, the mean absolute z position of the stable fixed points increases more rapidly than is the case for the lower acceleration rate. The blow up of the bunch initially lags behind the amplitude increase of the stable fixed points, leaving particles at the origin and close to the separatrix of a phase space that has the same features as in Fig. 6.13a. The amplitudes of the particles increase rapidly as they follow the separatrix (shown by the initial rapid rise of $\langle|z|\rangle$ in Fig. (b)); the particles then become trapped around the stable fixed points during the acceleration process, after which their amplitudes increase at approximately the same rate as for the stable fixed points.

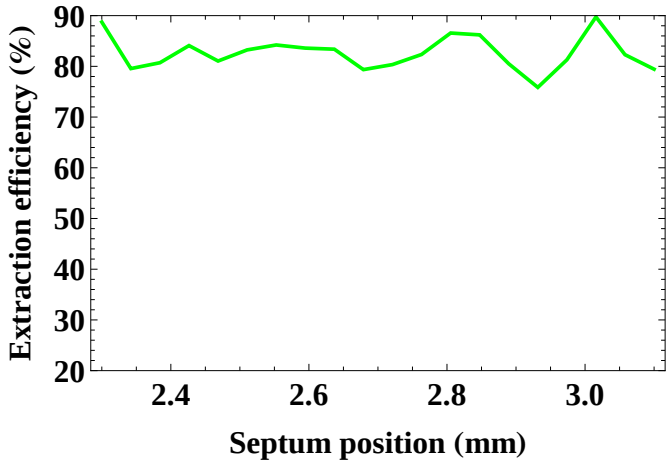

FIGURE 6.21: Extraction efficiency vs. position of electrostatic septum wire for the case where the beam is initially accelerated using a voltage of 160 kV, but the accelerating rf is switched off once the beam has reached a nominal energy of 55 MeV. We find that the extraction efficiency is comparable to that obtained when the 160 kV accelerating voltage remains on during the extraction process (Fig. 6.18).

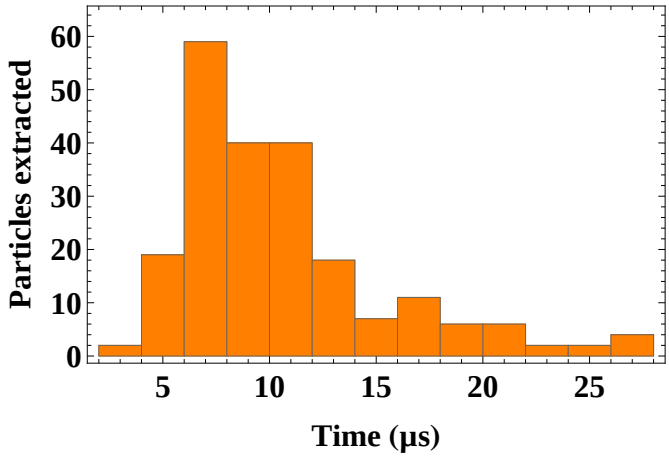

FIGURE 6.22: Histogram showing the extraction profile for half-integer resonance crossing, with turn 0 being the time at which the particles reach 55 MeV. We see that the time profile of the particle extraction is non-uniform and is completed within $\sim 28\,\mu$s (corresponding to 70 turns)

Half-integer resonant extraction has been used in practice to extract beam from a proton therapy accelerator; specifically on the earliest hospital based system, the Loma Linda University Medical Center's weak focusing synchrotron [93], where quadrupoles are used to move the horizontal tune close to a half-integer, and an octupole is used to perturb the beam and give amplitude detuning. During the extraction, the appearance of the horizontal phase space is similar to that shown for the vertical phase space in Fig. 6.13c, with a stable region surrounding the initial closed orbit position that is bound by a separatrix. Beyond the separatrix, the oscillation amplitude of particles grows significantly. To control the rate of extraction, the area contained by the separatrix is reduced by varying the strength of the quadrupoles, bringing the tune closer to half-integer. Using these methods, the Loma Linda accelerator was able to extract particles

over a period of $\sim 0.5\,\mathrm{s}$, however, early reports stated that the time structure of the extracted beam was difficult to control [94].

Chapter 7

Comparisons and Conclusions

7.1 Comparisons between synchrotrons and FFAGs for hadron therapy

The thin pencil-like beam extracted from an accelerator may be applied in several different ways (as discussed in chapter 1), they may be subject to scattering to create a treatment field (passive scattering), or they may be swept across the tumour volume laterally using steering magnets and longitudinally by varying the particle energy (active scanning). Tumours that are in (or close to) the lungs or heart move with the natural cycles of the body, whilst other organs may move non-periodically. If the duration of the irradiation of a tumour is comparable to the period of movement within the body or if a patient moves during treatment, then the dose distribution across the tumour may become distorted; this is known as an interplay effect, and is particularly problematic when using active scanning techniques due to the extra time that these take in comparison to passive scattering [95]. At least one study has concluded that active scanning of moving tumours can significantly affect the distribution of dose throughout a tumour volume, and that faster beam scanning technologies should be developed in order to make active scanning of moving tumours practical [96].

A clinical advantage of a high intensity, variable extraction energy machine such as PAMELA is that it may significantly reduce the time taken to irradiate a tumour, and therefore offer better treatment options in the case of moving tumours than currently possible using existing machines. The key to a successful treatment in this scenario is being able to provide adequate uniformity of dose across the tumour volume, and good conformation of dose to the boundaries of the treatment volume despite the motion of the tumour. Prior to active scanning, the tumour volume is visualised as smaller subvolumes, referred to as voxels, and a radiation dose is prescribed to each voxel. Active beam scanning can follow one of three time profiles: spot scanning, raster scanning or continuous scanning. In all cases, it is common practice to build up dose within the entire tumour volume by delivering dose to iso-energetic planes, starting with the furthermost plane. To improve the homogeneity of the dose throughout the tumour

volume, repainting of the tumour can be used; in this instance, the prescribed dose is not delivered during a single 3D scan of the tumour, but is instead delivered over some number of iterations. Repainting has been shown to offer the best dose homogeneity when the number of repainting iterations is increased, and the dose delivery time per voxel is shorter, however, this comes at the cost of an overall increase in treatment time [97].

For spot scanning, the steering magnets are set so that a voxel will be the focus of the extracted beam, the beam is switched on for the time necessary to supply the required dose to the voxel, the process is then carried out for the remaining voxels within the iso-energetic plane and then for the voxels at different depths after varying the beam energy. In calculating the time taken to deliver a treatment, the factors that have to be considered include the time taken to switch the beam on and off, the beam intensity (and resulting irradiation time), the time taken to move from voxel to voxel and the acceleration time when obtaining different beam energies.

When using a kicker or half integer resonance based extraction technique from PAMELA, the pulsed output of the machine is best suited to active scanning following the spot scanning method. The PAMELA design report [29] quotes eight bunches per voxel as being the average number required to obtain dose uniformity to within 2% due to the machine performance. The report also identifies repainting of a tumour as being a primary method for controlling the total dose delivered to the tumour, by keeping the bunch intensity fixed and then varying the number of times voxels are given a dose, rather than modulating the intensity of the beam source (the report also discusses the concept of accelerating between one and five bunches simultaneously within PAMELA as a method of quantising the intensity).

As a key metric, we will initially consider the time taken for a single iso-energetic plane to be irradiated once, and then for the entire tumour volume. We take a cubic treatment volume, measuring $100 \times 100 \times 100\,\mathrm{mm}^3$, with treatment spots separated by 5 mm, giving 400 voxels in a single plane. For PAMELA, the time taken for an acceleration cycle (Fig. 7.1) will be closely related to the time required per voxel. The time taken to sweep the beam across the plane using steering magnets is also accounted for, we base this calculation on a best case scenario of having steering magnets that can move the beam at a rate of 100 mm/ms and 50 mm/ms along the y and z transverse axes respectively (such magnets have been successfully tested at HIMAC [98]), the total time for beam travel is therefore ~ 24 ms. Finally, there is contribution to the delivery time of dose to a voxel which is dependent upon the extraction time from the machine, for kicker based extraction we assume a pulse length to the order of $0.1\,\mu$s, which we treat as negligible (if we instead selected the half integer resonance method, then the additional time due to extraction would be $28\,\mu$s per voxel). The time to irradiate a single plane is therefore ~ 424 ms (based on a 1 kHz repetition rate), whilst the time to irradiate the entire volume is the time taken to irradiate a single plane multiplied by the total number of planes that make up the volume.

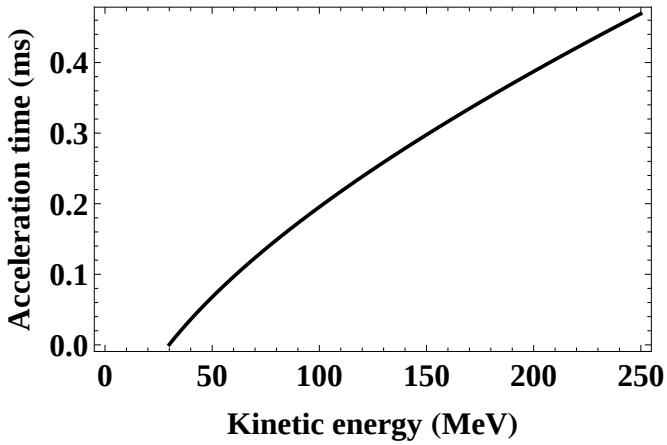

FIGURE 7.1: Plot to show the time taken for a proton to reach a given extraction kinetic energy in PAMELA. The acceleration rate is based on a total cavity voltage of 160 kV per turn and a synchronous phase of 60°. Protons reach a maximum energy of 250 MeV after ~ 0.44 ms, which is in keeping with the PAMELA repetition rate of 1 kHz.

For a proton synchrotron with beam being extracted using third integer resonance, the time taken to pause/resume extraction of the particle beam will be of importance, this can be expected to take tens of μs [99]. The pause time of a proton synchrotron was simulated in chapter 5, where we found that the time taken to pause extraction (once the transverse rf had been switched off) was dependent upon the number of turns needed for particles that have already crossed the separatrix to intercept the septum wire. In this scenario, a proportion of the bunch particles will continue to follow stable orbits close to the separatrix; upon resuming the transverse rf perturbation, particles will begin to cross the separatrix, with the number of turns required for particles to travel from the separatrix to the septum wire determining the resume time. It is the pause time of 60 μs (found through simulation) that we will use as the pause/resume time for the beam in this example. The pause/resume time of the beam and the time taken to sweep the beam across the plane determine the minimum time for irradiating a plane once, which for our example, we can calculate to be ~ 72 ms. The time taken to irradiate each voxel can be the dominant factor, with typical durations of between 1 and 10 ms [100–102], in which case the time taken to irradiate the example plane will take between 472 and 4072 ms. A single acceleration cycle is usually adequate to irradiate an iso-energetic plane, however, it is commonplace to have an acceleration cycle for each extraction energy. A slow cycling synchrotron typically has a repetition rate of ~ 1 Hz, making beam acceleration a significant overhead in the treatment time, whereas a rapid cycling synchrotron has a repetition rate that is to the order of 10 Hz (there are design studies for rapid cycling synchrotrons featuring both single turn [103] and slow [104] extraction).

For raster scanning, the beam is not switched off between the voxels on an iso-energetic plane, which eliminates the dead time experienced during spot scanning when moving between voxels and waiting for the beam to resume after being paused. Raster

scanning is the method used at HIT, where it was found that the method could produce a dose distribution that was almost identical to that obtained with spot scanning [105]. During raster scanning, the dose delivered to the tumour volume as the beam is being moved should be accounted for; given the stability of beam extraction during rf knockout, this can be done during the treatment planning stage [106]. The extra dose delivered during beam transition is proportional to the beam intensity and inversely proportional to the beam transition time; accounting for this extra dose whilst planning a treatment allows for the intensity of extracted beam to be increased without leading to an error in the dose given. Testing of a fast raster scanning system has been carried out at the HIMAC treatment facility [107]. Fast scanning relies upon the ratio of voxel dose delivered whilst the beam is stationary to that delivered during beam transition being minimised; for a given prescribed dose to the tumour volume, fast scanning will require either a higher intensity beam or a greater number of rescans than a slower scanning system. Testing was carried out by researchers at HIMAC on a spherical target of 60 mm diameter, with the spots within the volume separated by 3 mm, irradiation was carried out for 18 s; two average rates of scanning were applied, 1.23 and 0.154 ms per spot per scan. Within the 18 s, the tumour volume is scanned once at the lower scanning rate and eight times at the faster rate. Scanning takes place during a single spill of the synchrotron for both acceleration rates, with the spread out Bragg peak produced by using a range shifter (8 s of the total 18 s irradiation time is taken by changing the total range shifter thickness). Analysis of the measured dose distributions shows that both scanning rates gave a distribution that was in good agreement with the treatment plan and that the dose homogeneity through the irradiated volumes was within $\pm 3\%$. The differences between the two measurements was less than 1%.

For continuous scanning, the transverse motion of the beam is not paused at individual voxels for each iso-energetic plane. The total dose delivered to a location during a single scan can be controlled by varying the intensity of the extracted beam, or by changing the rate of scanning. Such methods have been tested at the cyclotron based treatment facility at PSI [108], where the beam scanning magnets can sweep the beam at rates of up to 20 mm/ms and 5 mm/ms along the y and z transverse axes. The intensity of the treatment beam is varied by controlling the extraction efficiency from the cyclotron, with a change from zero output to maximum intensity taking approximately 1 ms, whilst the range of protons within a patient is varied using a pair of carbon wedges, with a change of ~ 5 mm taking less than 50 ms. Testing at PSI has shown that continuous scanning can offer a dose distribution that is comparable to that obtained through spot scanning as well as reducing the treatment time (to less than a third of that for spot scanning for the experimental set up at PSI).

When we first considered the irradiation of the cubic volume using a spot scanning method, we saw that the beam extracted slowly from a synchrotron could be used to irradiate a single iso-energetic plane in approximately the same amount of time as would be taken with a FFAG operating on a single cycle per voxel basis, however, the longer

acceleration cycles of the synchrotron mean that the FFAG would take considerably less time to irradiate the entire tumour volume. Conventionally, passive scattering is used to treat moving tumours as it provides a treatment field that almost simultaneously irradiates an entire tumour volume, giving a uniform dose distribution throughout the tumour volume, but at the cost of the relatively poor conformation of dose to some boundaries of a tumour volume that is expected from passive scattering methods.

Minimising the time taken to irradiate a tumour using active scanning techniques has the potential to offer improved dose conformation to and distribution through a tumour volume. The benefits of raster and continuous scanning have been investigated at existing synchrotron and cyclotron based facilities, and have shown promising results, however, such scanning methods are not applicable to the existing PAMELA designs due to the short spill length of beam extraction using a kicker or half integer resonance crossing. If we consider a spot staying time of 0.154 ms and the beam being swept across a plane within the cubic volume at the maximum speed allowed by the steering magnets, then beam extracted from a synchrotron could irradiate the plane in $\sim$ 86 ms, which is approximately a fifth of the time taken by PAMELA. A range shifter is applied at HIMAC to avoid having acceleration cycles between irradiation of iso-energtic planes; the time taken to change the total thickness of the range shifter is approximately 500 ms, which still represents a significant overhead. Based on such a system, the time taken by a synchrotron to irradiate the example volume once will be comparable to that taken by PAMELA. In the case of both PAMELA and the synchrotron, multi-painting of the volume will improve the dose distribution; if rescanning occurs for iso-energetic planes independently, then the contribution from range shifter transitions to the total treatment time will be minimised. In order to offer further improvements to the dose distribution, methods for varying the extraction energy of particles from the synchrotron during a single spill (rather than having to use range shifters) are being developed for the HIMAC synchrotron [109–111]. For these new schemes, the beam is initially accelerated to the maximum energy required for treatment; after irradiating each iso-energetic plane, the beam is then decelerated. Using this method a 2 mm range shift can be achieved in $\sim$ 20 ms.

7.2 Conclusions

In the past 5 years the number of hadron therapy centres worldwide has almost doubled, with 55 currently operating. In the next 3 years, at least a further 31 centres are expected to open [21]. As physicists and clinicians become more experienced in delivering hadron therapy, the techniques used to give the best dose distributions across tumours have become more advanced. However, as was the case with photon therapy (which is reliant upon a 3 GHz electron linac), hadron therapy has so far been dependent on some fundamental technologies; these are accelerators of cyclotron and synchrotron type designs. Techniques have been developed for both types of accelerator that allow for

good conformation of dose to a tumour (particularly in comparison with photon based treatments), but there are advantages and disadvantages to using either accelerator type and there are strong motivations for developing new approaches to delivering hadrons for therapy. One such motivation is the size of accelerator required for hadron therapy: delivery of photon therapy using a linac is realised within a single, relatively small, room (a hospital may have several of these linacs, all working independently), whereas the size of the accelerator required for hadron therapy means that several treatment rooms are served by a single accelerator. Furthermore, the rigidity of the beams used for treatment means that a gantry weighing up to 600 tons [38] may be required to guide the beam around a patient; this, combined with a facility layout that is influenced by the accelerator size, makes hadron therapy more costly to implement and treatment scheduling more challenging than is the case for photon therapies. Research to develop technologies that could in future see hadron therapy being delivered as a single room solution is ongoing, and includes work on dielectric wall and plasma wakefield accelerators. These technologies were beyond the scope of this thesis. Instead, the focus of this thesis has been on the potential benefits of FFAG type accelerators, which are often considered to have the potential to offer hadron therapy in an interim period between current day cyclotron and synchrotron based facilities and a future of small linear accelerator based facilities. Fixed field alternating gradient accelerators do not necessarily mean a significant difference from current accelerators in terms of size (although FFAG designs could lead to gantries of reduced weight [112]), however it has been thought that FFAG could meet a second motivation for developing new accelerator technologies: the improvement of treatment standards.

For a synchrotron, the extraction energy is easily varied and third integer resonant extraction prolongs beam spill to a duration of seconds. These features make synchrotron accelerators well suited to active scanning techniques, which can be associated with the best conformation of dose to a tumour boundary and good uniformity of dose throughout the tumour volume. Active scanning is less effective when a tumour moves during irradiation (for example, when a tumour is located within a lung), as in these cases the dose distribution through the tumour may become distorted. This distortion may be minimised if a tumour is scanned rapidly, however the time taken to change the particle extraction energy within a synchrotron (due to repetition rates typically to the order of $\sim 1\,\mathrm{Hz}$) has limited the speed of active scanning. Without the requirement of ramping magnetic fields, fixed field alternating gradient accelerators can have much higher repetition rates than synchrotrons ($\sim 1\,\mathrm{kHz}$), whilst the relatively small radial excursion (in comparison to a cyclotron) makes variable extraction energy possible. The main aim of this thesis has been to consider whether FFAG accelerators may lead to improved dose distributions in comparison to those possible through using a synchrotron.

In chapters 1 and 2, we introduce aspects that are important to this thesis. Specifically, chapter 1 gives a brief history of the use of radiation for treating cancer, details why it is desirable to use protons for treating cancer and gives general descriptions of

the current accelerators used for proton therapy as well as FFAG. Chapter 2 introduced the magnetic fields that are used to guide and focus charged particle beams within an accelerator and how these fields may be applied to effectively control the beam dynamics. We then gave four case examples of particle accelerator designs; one synchrotron (which is active in delivering hadron therapy) and three FFAGs (each of a different design philosophy: scaling radial sector; non-scaling; scaling spiral sector).

Within this thesis, particle tracking studies play an important role in defining the dynamical properties of different accelerator lattices. The aim of chapter 3 was to introduce some possible approaches to particle tracking, to identify which approach is most appropriate for the systems considered in this thesis and to verify the use of the chosen approach. The Zgoubi tracking code was selected as it offers an appropriate level of accuracy when tracking non-paraxial rays as well as the flexibility to add new components. Zgoubi has been developed over a number of years, which offers both advantages and disadvantages; on one hand, it has many features that were useful for the tracking studies carried out, on the other hand, it is a very complex code and a user must be aware of the potential for nuances in the component definitions. A number of experienced local users were able to offer support for Zgoubi, and the pyZgoubi interface [113] is a helpful addition; however, it was decided early on that verifying the use of Zgoubi was important. The arc method of particle tracking was conceived as a simple tracking method, which could be used to check the results returned by Zgoubi, or potentially as an alternative to using Zgoubi. For the test lattice considered, the results from both tracking codes converged to a common solution. As may be expected for a numerical integrator, the time taken for Zgoubi to track particles through the accelerator lattices within this thesis has been a drawback, and has led to a reduction in the scope of some aspects of the work carried out. For example, in chapter 5, the slow extraction study is limited to tracking particles that are distributed uniformly in a region close to the separatrix. Ideally, a more complete study would have particles initially located in phase space according to a realistic bunch distribution and tracking should take place over a much greater number of turns than the 600 used (giving more information about the time structure of the extracted beam). Although not discussed in chapter 3, some time had been spent developing a program that could track particles in parallel (using the CUDA GPGPU (general-purpose computing on graphics processing units) environment). This code used the arc method as numerical integrator, but was not pursued due to the relatively poor rate of convergence for the arc algorithm, and the amount of time that would be required to develop the code sufficiently. Given more time, a GPGPU implementation of a tracking code based on an improved version of the arc method, or an alternative numerical integrator, could have allowed for a more tracking study with a wider scope.

Chapter 4 offers further verification of the Zgoubi tracking code, this time by comparing the results obtained by tracking through hard edge and field map models of the EMMA lattice with experimental results. An important aspect of this chapter is the

evaluation of whether hard edge modelling offers a sufficient approximation of an FFAG when carrying out tracking studies within the design stage of FFAG development. The large transverse apertures and short lengths of the EMMA magnets, combined with the short drift lengths that separate the magnets within the EMMA lattice, make EMMA a demanding test of hard edge modelling. The field maps that were calculated using Opera 3D, formed an intermediate step between hard edge modelling and experimental data. Once the calibrations made by Tesla Engineering during the manufacture of the EMMA magnets had been incorporated into the Opera design, we found that tracking through the field map model provided a good description of the beam dynamics that had been obtained experimentally. We also identified that the close proximity of the two magnets within an EMMA cell meant that the yoke from the defocusing quadrupole would influence the field for focusing quadrupole and vice-versa. We did not carry out tracking studies that included alignment errors (which could offer greater insight into the significance of the small differences between the dynamics observed in the comparison of field map and experimental data), however, this may be a challenge as the field for a cell may not be described simply by the superposition of fields from the individual quadrupoles. When comparing hard edge and field map models, we used the differences in values calculated for orbital period, closed orbit position and betatron tune as metrics for the level of agreement between the two models. We found that due to the large fringe field contribution within an EMMA cell, and the different transverse offsets of the defocusing and focusing quadrupoles, it was not possible to find a hard edge model that simultaneously minimised each of the metrics by varying the field strengths alone. Although we were unable to find close agreement, the hard edge model provided enough of the general features of the beam dynamics found experimentally (including a parabolic plot of orbital period vs. momentum and the approximate tune range) in order for hard edge modelling to a very useful tool during the design stage. As part of the comparison of field map and experimental data, we looked at methods for characterising the EMMA lattice experimentally; as part of this process, we developed new methods for reconstructing the momentum distribution of a bunch based on measurements made using a pair of BPMs and for calculating the Courant-Snyder parameters at the location of a single BPM. These new methods where used when calculating the tune and orbital period dependence on transverse amplitude, which again were compared with results of field map tracking.

EMMA is a unique machine, with several features that pose a significant challenge when calculating beam dynamics through hard edge modelling; any FFAG built for investigating the application of proton therapy is likely to pose its own challenges. The comparison between hard edge modelling, field map and experimental data carried out for EMMA cannot provide a definitive statement on the adequacy of hard edge modelling for FFAG design. However, we use the results obtained to give an appreciation of the level of difference between models that can perhaps be expected, and an indication as to when field map modelling may be appropriate.

In chapter 5, we looked at the design of a proton synchrotron with the aim of providing a benchmark for FFAG based proton therapy. The basis of the synchrotron is six DOFO cells, and it is anticipated that the machine would accelerate particles between 60 MeV and 230 MeV, and have a repetition rate of ~ 1 Hz. We investigated the properties of beam extraction using a third integer resonance; to achieve this we followed methods described for the HIMAC synchrotron, which involved setting the quadrupoles so that the horizontal tune is close to a third integer and including two pairs of sextupoles within the lattice. The sextupoles within a pair are separated by a betatron phase advance of πQ_y, and their strengths set so that a sextupole pair contributes a geometric perturbation, but does not affect the chromaticity of the lattice. When close to third integer resonance, the geometric perturbation leads to a separatrix being formed in phase space. Before the extraction process begins, particles follow stable orbits within the area contained by the separatrix. During the extraction process, particles are pushed over the separatrix, and the perturbation leads to particles reaching the transverse position of an electrostatic septum after some number of turns. We saw that by controlling the ratio of strength of the two sextupole pairs, we were able to rotate the separatrix in transverse phase space; this allowed us to rotate the separatrix to optimise the extraction efficiency. The particles within a bunch are expected to have some range of momenta, as a result of this, the separatrices for different momenta particles are offset from one another in phase space. This offset can reduce the efficiency of extraction due to the finite length of a septum wire that is positioned at an angle that matches the on momentum separatrix. To align a single arm of the separatrices for a range of momenta, we sought to satisfy the Hardt condition by controlling the chromaticity of the lattice after including an additional pair of sextupoles (this time the strengths were set so that the pair affected the chromaticity and not the geometric perturbation).

We investigated the extraction of particles by using a transverse rf kicker to push particles over the separatrix, which is a current gold standard for synchrotron-based hadron therapy delivery. We saw that the tune difference between the origin of phase space ($Q_y = 1.672$) and one side of the separatrix ($Q_y = 5/3$) meant that a range of transverse rf frequencies would be required to extract all of the particles following stable orbits at the start of the extraction process. As discussed in the comments on chapter 3, a limitation of the extraction study carried out was that particles tracked during the simulation had starting positions close to the separatrix. This was due to the CPU time that would have been required for a more complete simulation, but it also meant that it was possible to extract the particles using a single transverse rf frequency. This is in keeping with the practical experiences of rf knock-out extraction for hadron therapy, where a mixed transverse rf frequency is used to diffuse particles towards the separatrix and a single rf frequency is used to push particles over the separatrix. The result of the approximation made is that whilst the tracking studies can be used to indicate features such as the extraction efficiency and the emittance of the extracted beam, we are not able to estimate the time structure of the extracted beam whilst the rf perturbation is

active. However, we were able to estimate the time taken to pause beam extraction once the rf kicker has been switched off, as this is based on the time taken for particles to move from a position just beyond the separatrix to the amplitude of the electrostatic septum. Through the tracking study, we found that the rf knock-out procedure for the synchrotron design led to a high extraction efficiency and a horizontal emittance of the extracted beam that was smaller than required.

In chapter 6, we looked at a half integer resonance based extraction method for the PAMELA lattice. A half integer extraction method has previously been applied for the synchrotron at the Loma Linda Medical Center; this is similar to the third integer extraction method discussed in chapter 5 in that it relies upon geometric perturbations to limit the stable area of phase space, and then uses methods to control the rate at which particles cross the separatrix. The method applied to the PAMELA lattice differs from this approach, and instead relies upon the rapid blow up of beam shortly after a half integer resonance is crossed. We also consider the extraction process for the vertical rather than horizontal axis due to the large transverse aperture and significant contribution of non-linear fields within the PAMELA lattice. We aim to cross the half integer resonance with a negative chromaticity, as this has previously been associated with a smaller energy spread of extracted beam than is the case with a positive chromaticity (due to amplitude detuning). To achieve this, we first altered the sextupole and decapole field components to give a small negative chromaticity and then varied the field at the reference radius for the defocusing magnets so that the tune crosses a half integer close to the desired extraction energy. The perturbation that leads to beam extraction is included by modifying the quadrupole component of a single defocusing FFAG magnet. As key metrics, we consider the extraction efficiency and energy range of the extracted beam when different rates of acceleration and quadrupole perturbations are included. Initially, acceleration of the beam continues throughout the extraction process, and we find that the acceleration rate is the most important factor for both the extraction efficiency and the energy spread of extracted beam. Specifically, we find that a more rapid rate of acceleration leads to improved extraction efficiency but at the cost of an increased energy spread. This problem can be overcome by stopping acceleration shortly after the beam has crossed the resonant energy, in which case good extraction efficiency is preserved for high acceleration rates, but there is no energy spread introduced by the extraction process. A better understanding of the extraction process was obtained by considering the phase space portraits for the system, as well as analysis of tracking data. We found that the initial blow up of the beam lags behind the evolution of the contours of constant Hamiltonian during the acceleration process. If the two stable fixed points that occur once the betatron tune is less than half integer separate rapidly in phase space (as is the case for a greater rate of acceleration), then the initial blow up of the beam will encompass a greater range of particle amplitudes, and it is possible to obtain good extraction efficiency.

This chapter begins with a comparison of the extraction techniques demonstrated

for synchrotron and FFAG. In chapter 6, we found that the half integer based method applied to PAMELA provides an alternative to kicker-based extraction (with beam extracted as a short pulse), but that the method does not lead to slow extraction. The slow extraction techniques that have been developed for the HIMAC synchrotron allow for a treatment volume to be scanned by a particle beam in a very short period of time; minimising the time taken to scan a volume is associated with improved dose distributions across moving tumours. Although the fixed fields of an FFAG make high repetition rates possible, the time that is required to accelerate a beam may still lead to volume irradiation times that are longer for an FFAG operating in a pulsed extraction mode than is possible with a synchrotron. For the scenario considered, there is no obvious and significant benefit to using an FFAG rather than a synchrotron, however it is important to remember that synchrotron based treatment delivery is much more developed than is the case for FFAG. One possible avenue for further research would be to investigate slow extraction techniques for FFAG, whilst a second could be to look at minimising the number of voxels required to treat a tumour volume (by being able to vary the spot size during irradiation).

Appendix A

Approximate model of an RF cavity

Given a time dependent longitudinal electric field, $E(t) = E_0 \sin(\omega t + \phi_0)$, within an rf cavity of length, L, the change in the kinectic energy of a particle is described below. The Lorentz force on a charged particle that's travelling through an electric field is given by

$$\frac{d\mathbf{p}}{dt} = q\mathbf{E}(t). \tag{A.1}$$

Given the relation

$$dE_k = vd\mathbf{p}, \tag{A.2}$$

then,

$$\Delta E_k \approx q \int_L \mathbf{E}(t)dx, \tag{A.3}$$

if the velocity of particle does not change much as it passes through the rf cavity. Considering the case where a particle passes the midpoint of the cavity at $t = 0$ gives

$$\Delta E_k = qE_0 \int_{-L/2}^{L/2} \sin(\omega t + \phi_0)dx \tag{A.4}$$

which is equivalent to

$$\Delta E_k = qE_0 \int_{-L/2}^{L/2} (\sin(\omega t)\cos(\phi_0) + \cos(\omega t)\sin(\phi_0))dx. \tag{A.5}$$

As sin is odd, then,

$$\Delta E_k = qE_0 \int_{-L/2}^{L/2} (\cos(\omega t)\sin(\phi_0))dx \tag{A.6}$$

Assuming constant velocity as particle makes a single pass throught the cavity, then,

$$t = \frac{L}{v} \tag{A.7}$$

and

$$\Delta E_k = qE_0 \sin\phi_0 \int_{-L/2}^{L/2} \cos\left(\frac{\omega x}{v}\right) dx \tag{A.8}$$

$$\Delta E_k = qE_0 \sin\phi_0 \frac{v}{\omega} \left[\sin\left(\frac{\omega x}{v}\right)\right]_{-L/2}^{L/2} \tag{A.9}$$

Given that

$$\omega = \frac{2\pi c}{\lambda}, \tag{A.10}$$

then

$$\Delta E_k = qE_0 \sin\phi_0 \frac{v}{\omega} \left[\sin\left(\frac{\omega x}{v}\right)\right]_{-L/2}^{L/2} \tag{A.11}$$

$$\Delta E_k = qV_0 \sin\phi_0 T, \tag{A.12}$$

where T is the transit time factor, and is given by

$$T = \frac{\sin\left(\frac{L\pi}{\beta\lambda}\right)}{\frac{L\pi}{\beta\lambda}}. \tag{A.13}$$

Appendix B

EMMA modelling data

B.1 Quadrupole calibration

Each of EMMA's 84 quadrupoles have been calibrated so as to reduce the variation in integrated magnetic field between each magnet of a given type (focusing or defocusing). The calibration, which was carried out by Tesla Engineering, involved changing the effective length of each magnet by moving the field clamp position longitudinally along the

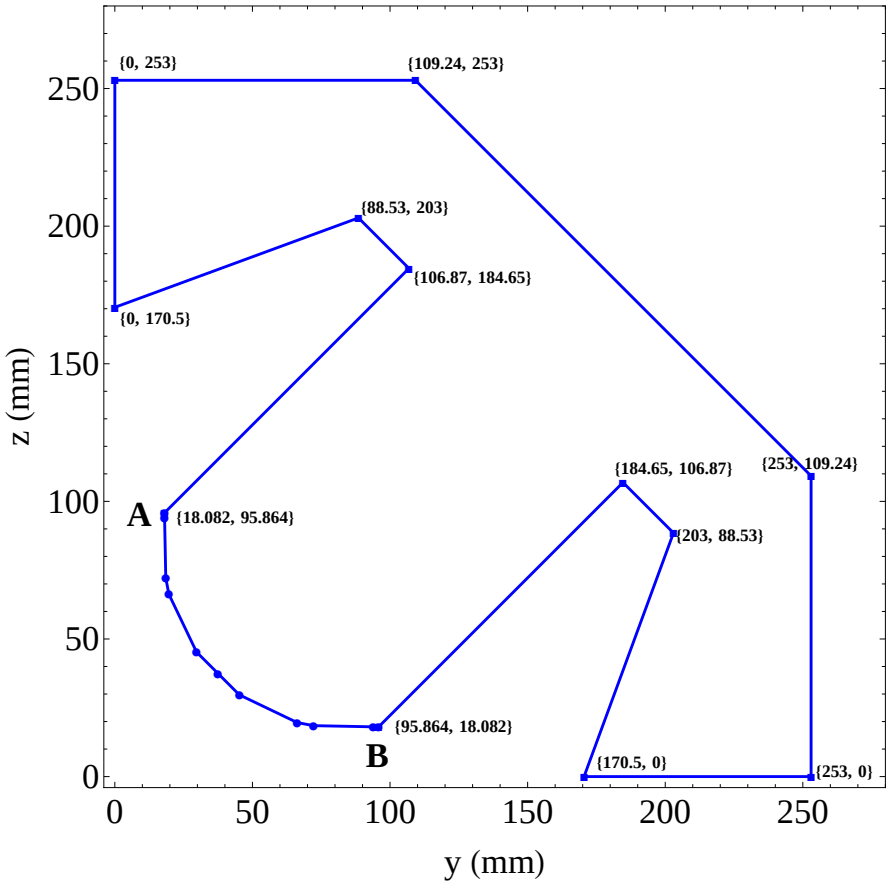

FIGURE B.1: Cross section of a quadrant of the defocusing quadrupole. The profile of the pole tip (from point A to B) is provided in table B.1.The defocussing quadrupole iron has a thickness of 65 mm.

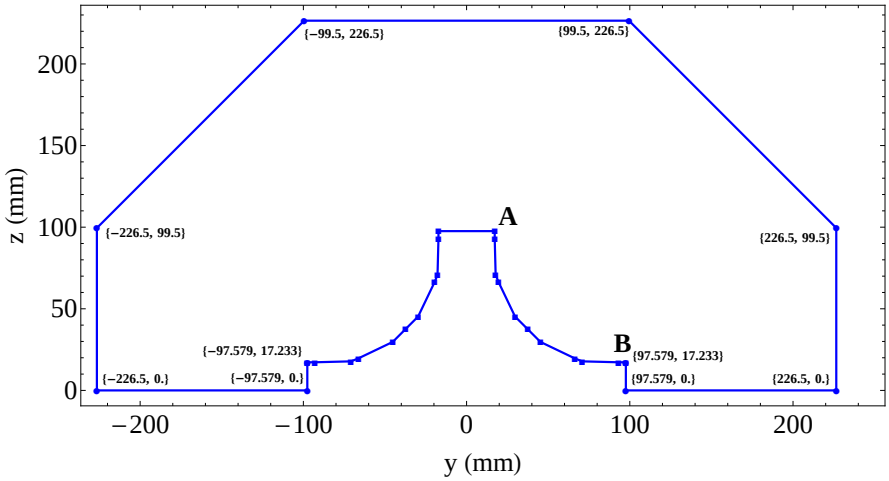

FIGURE B.2: Cross section of half of the defocusing quadrupole field clamp. The profile of the pole tip (from point A to B) is provided in table B.1. The clamp has a thickness of 8 mm (8 × 1 mm laminations of 1200-100A steel). The original cell design specifies a distance of 17 mm between the D yoke and field clamp.

TABLE B.1: Profile of the defocusing pole tip from point A to B of Fig. B.1.

Point	$y\,(mm)$	$z\,(mm)$
1 (A)	17.233	95.015
2	17.233	93.015
3	17.799	70.832
4	19.622	66.291
5	29.698	45.255
6	37.477	37.477
7	45.255	29.698
8	66.291	19.622
9	70.832	17.799
10	93.015	17.233
11 (B)	95.015	17.233

TABLE B.2: Profile of the defocusing quadrupole field clamp from point A to B of Fig. B.2.

Point	$y\,(mm)$	$z\,(mm)$
1 (A)	18.082	97.58
2	18.082	93.864
3	18.506	72.246
4	19.622	66.291
5	29.698	45.255
6	37.477	37.477
7	45.255	29.698
8	66.291	19.622
9	72.246	18.506
10	93.864	18.082
11 (B)	97.58	18.082

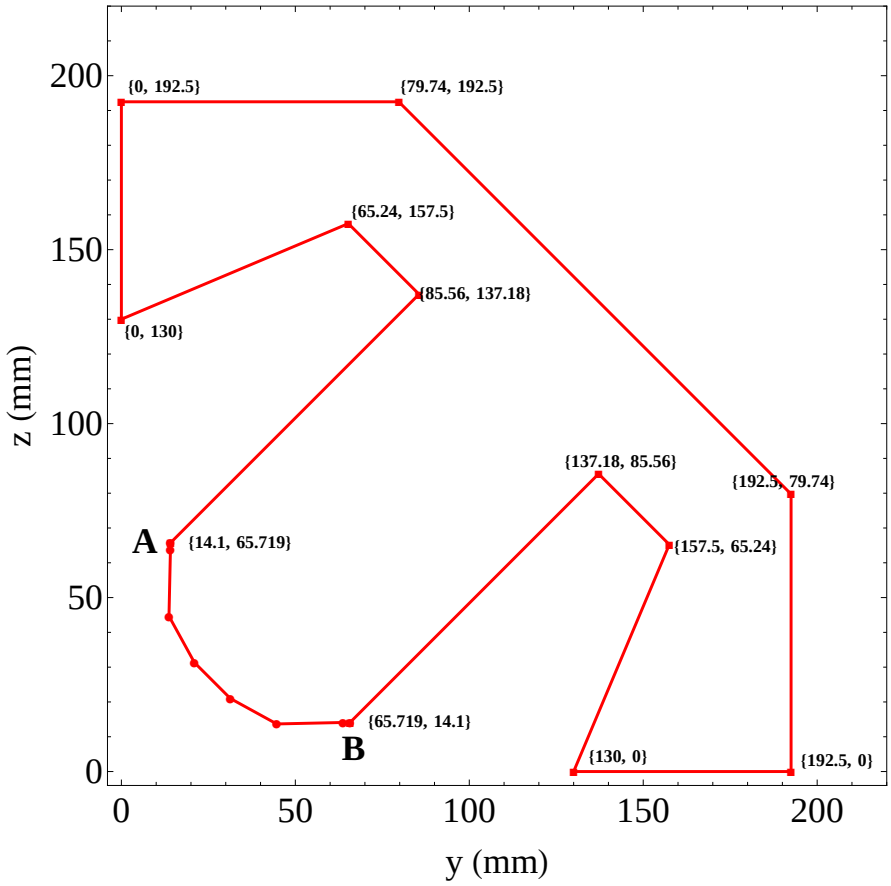

FIGURE B.3: Cross section of a quadrant of the focusing quadrupole. The profile of the pole tip (from point A to B) is provided in table B.3. The defocussing quadrupole iron has a thickness of 55 mm.

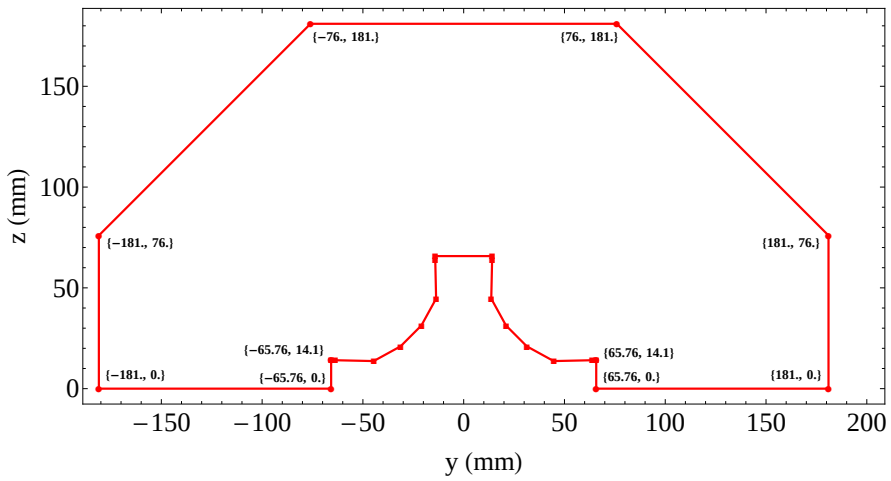

FIGURE B.4: Cross section of half of the focusing quadrupole field clamp. The profile of the pole tip (from point A to B) is provided in table B.3. The clamp has a thickness of 5 mm (5 × 1 mm laminations of 1200-100A steel.)

TABLE B.3: Profile of the focusing pole tip from point A to B of Fig. B.3.

Point	$y\,(mm)$	$z\,(mm)$
1 (A)	14.100	65.719
2	14.100	63.719
3	13.680	44.650
4	21.000	31.320
5	31.320	21.000
6	44.650	13.680
7	63.719	14.100
8 (B)	65.719	14.100

TABLE B.4: Saturation curves for 1200-100A and 1006.

(A) 1200-100A

$H\,(A/m)$	$B\,(T)$
0	0
10	0.016
20	0.0212
30	0.0477
40	0.0849
50	0.127
60	0.202
70	0.286
80	0.371
90	0.451
100	0.53
200	1.01
300	1.2
400	1.3
500	1.36
600	1.4
700	1.44
800	1.46
900	1.48
1000	1.5
2000	1.6
3000	1.65
4000	1.69
5000	1.71
6000	1.75
7000	1.77
8000	1.79
9000	1.81
10000	1.83
20000	1.98
30000	2.07
40000	2.12

(B) 1006

$H\,(A/m)$	$B\,(T)$
0	0
56	0.1
101	0.2
144	0.3
186	0.4
229	0.5
275	0.6
324	0.7
376	0.8
433	0.9
498	1
570	1.1
667	1.2
823	1.3
1167	1.4
1875	1.5
2909	1.6
5152	1.7
9474	1.8
15833	1.9

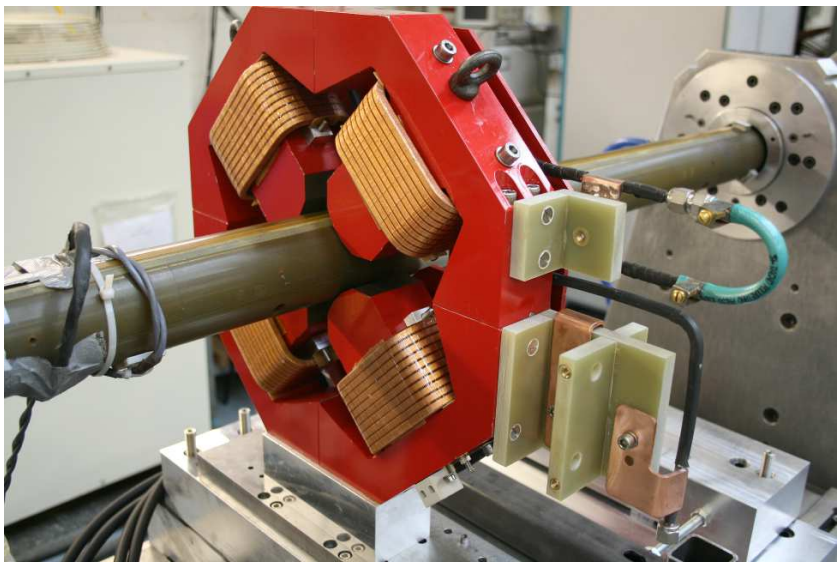

FIGURE B.5: A focusing quadrupole on the rotating coil test bench. The field clamp positions were adjusted in order to give the same integrated field (measured at 35 mm) for all magnets of the given type.

magnet axis until the desired integrated magnetic field was achieved, with the measurement of the integrated field carried out using a rotating coil (Fig B.5). For both focusing and defocusing quadrupoles the integrated field was measured at a radius of 35 mm, with the coil currents set to 376 A and 364 A for the focusing and defocusing magnets respectively. Post calibration, the mean integrated field was (0.55078 $\pm$ 0.00004) T for the focusing magnets and (0.47240 $\pm$ 0.00003) T for the defocusing magnets. Figure B.7 shows the positions of the field clamps following calibration. The mean positions for the field clamps (measured from the end of the corresponding magnet yoke to the start of the clamp) are 30.2 $\pm$ 0.6 mm and 18.1 $\pm$ 0.8 mm for the focusing and defocusing magnet clamps, respectively. Within the Opera model of the EMMA cell, the field clamps have been located according to the mean positions stated above (rather than the design position). Before modelling the complete cell, the EMMA magnets (with field clamps) were modelled independently with the coil currents set as for the calibration method used by Tesla Engineering. The quadrupole currents were then scaled in order to match the integrated fields calculated with Opera (at an offset of 35 mm from the quadrupole centre) to those measured with rotating coil.

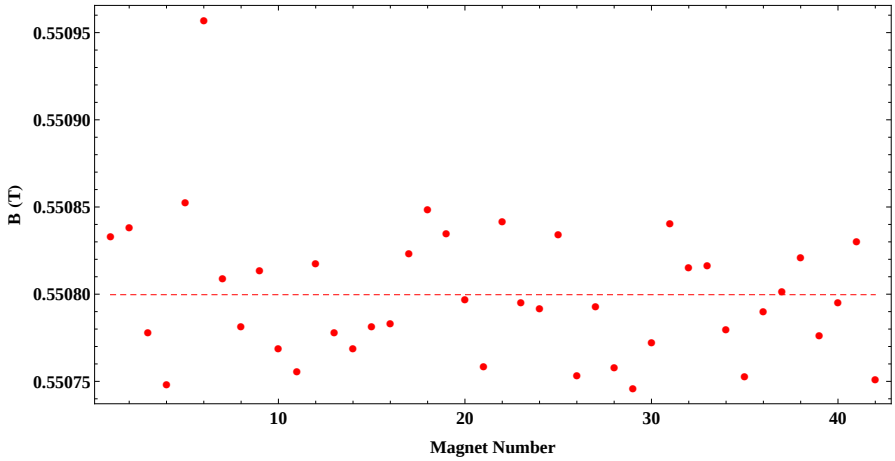

(a) Focusing quadrupole

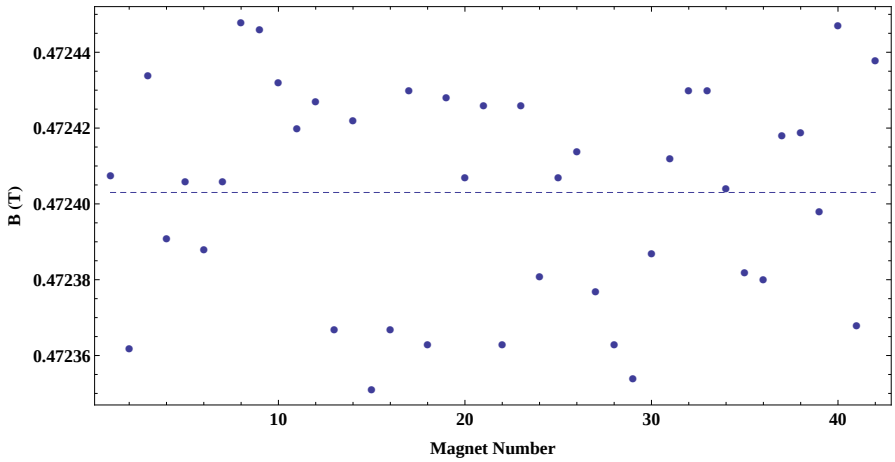

(b) Defocusing quadrupole

FIGURE B.6: Measured integrated fields for each of the EMMA quadrupoles following calibration.

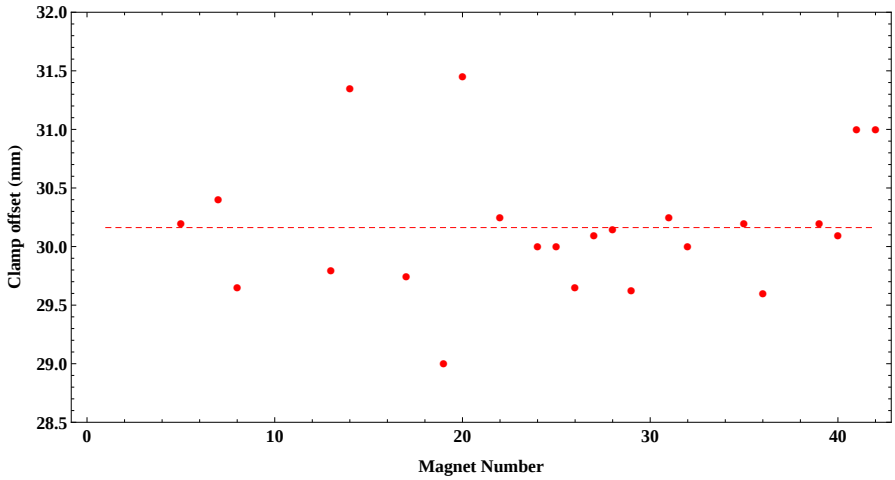

(a) Focusing quadrupole

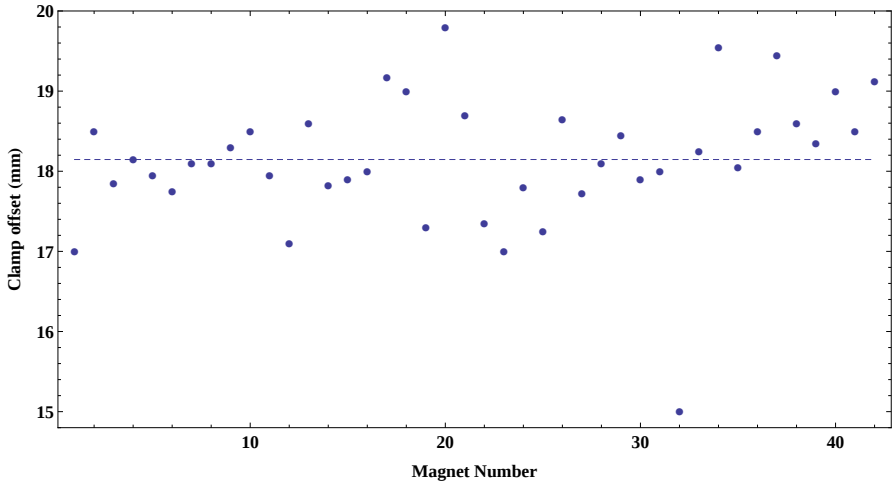

(b) Defocusing quadrupole

FIGURE B.7: Offset of field clamp from quadrupole iron following calibration, with dashed lines to show the mean position. Clamp position data was unavailable for 19 of the focusing quadrupoles.

Bibliography

[1] M. Baker. Medical linear accelerator celebrates 50 years of treating cancer. *Stanford University*, 2007.

[2] R.R. Wilson. Radiological use of fast protons. *Radiology*, 47(5):487–491, 1946. PMID: 20274616.

[3] J.S. Loeffler and M. Durante. Charged particle therapy–optimization, challenges and future directions. *Nature Reviews Clinical Oncology*, 10(7):411–24, 2013.

[4] P. Symonds, C. Deehan, J.A. Mills, and C. Meredith, editors. *Walter and Miller's Textbook of Radiotherapy*, chapter 3. Elsevier, 2012.

[5] P. Symonds, C. Deehan, J.A. Mills, and C. Meredith, editors. *Walter and Miller's Textbook of Radiotherapy*, chapter 10. Elsevier, 2012.

[6] H. Paganetti, editor. *Proton Therapy Physics*. CRC Press, 2012.

[7] `https://commons.wikimedia.org/wiki/User:HPaul`. [Online; accessed 03-November-2015].

[8] T. Bortfeld. An analytical approximation of the Bragg curve for therapeutic proton beams. *Medical Physics*, 24(12):2024–2033, 1997.

[9] J.H. Lawrence, C.A. Tobias, J.L. Born, C.C. Wang, and J.H. Linfoot. Heavy-particle irradiation in neoplastic and neurologic disease. *Journal of Neurosurgery*, 19(9):717–722, 1962. PMID: 14462992.

[10] C.C. Ma and T. Lomax, editors. *Proton and Carbon Ion Therapy*. CRC Press, 2013.

[11] H. Owen, R. MacKay, K. Peach, and S. Smith. Hadron accelerators for radiotherapy. *CONTEMPORARY PHYSICS*, 55(2):55–74, APR 3 2014.

[12] A. M. Koehler, R. J. Schneider, and J. M. Sisterson. Flattening of proton dose distributions for largefield radiotherapy. *Medical Physics*, 4(4):297–301, 1977.

[13] E. Grusell, A. Montelius, A. Brahme, G. Rikner, and K. Russell. A general solution to charged particle beam flattening using an optimized dual-scattering-foil

technique, with application to proton therapy beams. *Physics in Medicine and Biology*, 39(12):2201, 1994.

[14] D.J. Brenner, C.D. Elliston, E.J. Hall, and H. Paganetti. Reduction of the secondary neutron dose in passively scattered proton radiotherapy, using an optimized pre-collimator/collimator. *Physics in Medicine and Biology*, 54(20):6065, 2009.

[15] D. Jette and W. Chen. Creating a spread-out Bragg peak in proton beams. *Physics in Medicine and Biology*, 56(11):N131, 2011.

[16] A.M. Koehler, R.J. Schneider, and J.M. Sisterson. Range modulators for protons and heavy ions. *Nucl.Instrum.Meth.*, 131:437–440, 1975.

[17] `https://en.wikipedia.org/wiki/User:MarkFilipak`. [Online; accessed 03-November-2015].

[18] T. Kanai, K. Kawachi, Y. Kumamoto, H. Ogawa, T. Yamada, H. Matsuzawa, and T. Inada. Spot scanning system for proton radiotherapy. *Medical Physics*, 7(4):365–369, 1980.

[19] S. Scheib and E. Pedroni. Dose calculation and optimization for 3D conformal voxel scanning. *Radiation and Environmental Biophysics*, 31(3):251–256, 1992.

[20] A. Knopf, T.S. Hong, and A. Lomax. Scanned proton radiotherapy for mobile targetsthe effectiveness of re-scanning in the context of different treatment planning approaches and for different motion characteristics. *Physics in Medicine and Biology*, 56(22):7257, 2011.

[21] Particle Therapy Co-Operative Group. `http://ptcog.ch/`. [Online; accessed 29-August-2014].

[22] U. Amaldi, P. Berra, K. Crandall, D. Toet, M. Weiss, R. Zennaro, E. Rosso, B. Szeless, M. Vretenar, C. Cicardi, C. De Martinis, D. Giove, D. Davino, M.R. Masullo, and V. Vaccaro. LIBOa linac-booster for protontherapy: construction and tests of a prototype. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 521(23):512 – 529, 2004.

[23] G.J. Caporaso, T.R. Mackie, S. Sampayan, Y.J. Chen, D. Blackfield, J. Harris, S. Hawkins, C. Holmes, S. Nelson, A. Paul, B. Poole, M. Rhodes, D. Sanders, J. Sullivan, L. Wang, J. Watson, P.J. Reckwerdt, R. Schmidt, D. Pearson, R.W. Flynn, D. Matthews, and J. Purdy. A compact linac for intensity modulated proton therapy based on a dielectric wall accelerator. *Physica Medica*, 24(2):98 – 101, 2008. Highlights of the EFOMP Congress - First European Conference on Medical Physics Il Ciocco, Italy, 20-22 September.

[24] T. Tajima, D. Habs, and X. Yan. Laser acceleration of ions for radiation therapy. *Reviews of Accelerator Science and Technology*, 02(01):201–228, 2009.

[25] CONFORM. `http://www.conform.ac.uk`. [Online; accessed 29-August-2014].

[26] C. Kleffner, D. Ondreka, and U. Weinrich. The Heidelberg ion therapy (HIT) accelerator coming into operation. *AIP Conference Proceedings*, 1099(1):426–428, 2009.

[27] TASSA. `http://www.tassausa.org`. [Online; accessed 03-November-2015].

[28] J. Fourrier, J. Pasternak, and F. Méot. Acceleration in spiral FFAG using field map data. In *Proceedings of EPAC 2008*, Genoa,Italy, 2008.

[29] The PAMELA conceptual design report. Technical report, 2012. [Online; accessed 22-May-2014].

[30] D. Trbojevic and V.S. Morozov. Non-scaling Fixed Field Alternating Gradient Permanent Magnet Cancer Therapy Accelerator. In *Proceedings of IPAC 2011*, San Sebastián, Spain, 2011.

[31] K. R. Symon, D. W. Kerst, L. W. Jones, L. J. Laslett, and K. M. Terwilliger. Fixed-field alternating-gradient particle accelerators. *Phys. Rev.*, 103:1837–1859, 1956.

[32] J.B. Rosenzweig. *Fundamentals of Beam Physics*, chapter 2. Oxford University Press, 2003.

[33] S.Y. Lee. *Accelerator Physics*, page 48. World Scientific Publishing, 2 edition, 2004.

[34] H. Wiedemann. *Particle Accelerator Physics*. Springer, 2 edition, 1998.

[35] D.A. Edwards and M.J. Syphers. *An Introduction to the Physics of High Energy Accelerators*. Wiley-VCH Verlag GmbH, 2008.

[36] A. Wolski. *Beam Dynamics in High Energy Particle Accelerators*. Imperial College Press, 2014.

[37] K.Y. Ng. *Handbook of Accelerator Physics and Engineering*, page 108. World Scientific Publishing, 1999.

[38] M. Galonska. Heidelberg gantry experience, 2014. presentation at the 'Modern Hadron Therapy Gantry Developments' workshop, Cockcroft Institute, UK.

[39] M.J. Barnes, J. Borburgh, B. Goddard, and C. Houricam. Injection and extraction magnets: septa.

[40] S. Machida *et al.* Commissioning of 150 MeV FFAG synchrotron. In *Proceedings of EPAC 2004*, Lucerne, Switzerland, 2004.

[41] Y. Yonemura *et al.* Development of FFAG accelerator at KEK. In *Proceedings of PAC 2005*, Knoxville, Tennessee, 2005.

[42] Kyushu University. `http://www.cabas.kyushu-u.ac.jp/home/ja/ffag`. [Online; accessed 03-November-2015].

[43] M. Aiba and F. Méot. Determination of KEK 150 MeV FFAG parameters from ray-tracing in TOSCA field maps. Technical Report CERN-NUFACT-NOTE-140. CARE-Note-2004-030-BENE. CEA-DAPNIA-2004-188. CERN-NEUTRINO-FACTORY-NOTE-140, CERN, Geneva, October 2004.

[44] M. Aiba *et al.* Beam injection and extraction in 150 MeV FFAG. In *Proceedings of EPAC 2002*, Paris, France, 2002.

[45] S. Machida. Scaling fixed-field alternating gradient accelerators with a small orbit excursion. *Phys. Rev. Lett.*, 103:164801, Oct 2009.

[46] S.L. Sheehy, K.J. Peach, H. Witte, D.J. Kelliher, and S. Machida. Fixed field alternating gradient accelerator with small orbit shift and tune excursion. *Phys. Rev. ST Accel. Beams*, 13:040101, Apr 2010.

[47] T. Yokoi, K. Peach, and H. Witte. Beam extraction in PAMELA NS-FFAG. In *Proceedings of IPAC 2010*, Kyoto, Japan, 2010.

[48] T. Yokoi. Kicker-less FFAG: An approach to high intensity FFAG. volume 26, pages 1873–1886, 2011.

[49] J. Pasternak, J. Fourrier, E. Froidefond, B. Autin, F. Méot, D. Neuveglise, and T. Planche. Spiral FFAG for proton therapy. In *Particle Accelerator Conference, 2007. PAC. IEEE*, pages 1404–1406, 2007.

[50] F. Méot. The ray-tracing code zgoubi status. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 767:112 – 125, 2014.

[51] MAD. `http://mad.web.cern.ch/mad/`. [Online; accessed 29-August-2014].

[52] E. Keil, J.S. Berg, and A.M. Sessler. Electron model of an FFAG muon accelerator. In *Proceedings of EPAC 2004*, Lucerne, Switzerland, 2004.

[53] C. Johnstone and S. Koscielniak. Tune-stabilized linear-field FFAG for carbon therapy. In *Proceedings of EPAC 2006*, Edinburgh, Scotland, 2006.

[54] B.D. Muratori, J.K. Jones, and H.L. Owen. An EMMA racetrack. In *Proceedings of IPAC 2012*, New Orleans, Louisiana, USA, 2012.

[55] M. Berz and K. Makino. *COSY Infinity 9.1 Beam Physics Manual*. Michigan State University, 6 2011.

[56] M. Berz. Differential Algebraic Description of Beam Dynamics to Very High Orders. *Part.Accel.*, 24:109–124, 1989.

[57] M. Berz. Arbitrary order description of arbitrary particle optical systems. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 298(13):426 – 440, 1990.

[58] K.M. Hock and C.S. Edmonds. Tracking protons in an FFAG and a synchrotron with tiny arcs. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 664(1):140 – 147, 2012.

[59] S. Machida. Resonance crossing and dynamic aperture in nonscaling fixed field alternating gradient accelerators. *Phys. Rev. ST Accel. Beams*, 11:094003, Sep 2008.

[60] S. Machida *et al.* Acceleration in the linear non-scaling fixed-field alternating-gradient accelerator EMMA. *Nature Physics*, 8(3):243–247, March 2012.

[61] ASTeC. EMMA design drawings. `http://www.astec.ac.uk/emmafiles/magnets/`, 2008. [Online; accessed 14-February-2014].

[62] Cobham plc. Opera-3D design software. `http://www.cobham.com/opera`, 2014. [Online; accessed 10-February-2014].

[63] J.M Garland. *Beam Dynamics Studies of the EMMA Linear Non-Scaling FFAG*. PhD thesis, University of Manchester, 2013.

[64] F. Jackson *et al.* The status of the ALICE R&D facility at STFC daresbury laboratory. In *Proceedings of IPAC 2011*, pages 934–936, San Sabástian, Spain, 2011.

[65] J.M. Garland *et al.* Characterisation of the ALICE accelerator as an injector for the EMMA NS-FFAG. In *Proceedings of IPAC 2010*, pages 4343–4345, Kyoto, Japan, 2010.

[66] J. Laskar, C. Froeschlé, and A. Celletti. The measure of chaos by the numerical analysis of the fundamental frequencies. application to the standard mapping. *Physica D Nonlinear Phenomena*, 56:253–269, May 1992.

[67] Wolfram Research Inc. *Mathematica*. Wolfram Research Inc., Champaign, Illinois, version 9.0 edition, 2012.

[68] R.E. Meller *et al.* Decoherence of kicked beams. Technical Report SSC-N-360, 1987.

[69] S.Y. Lee. Decoherence of the kicked beams II. Technical Report SSCL-N-749, 1991.

[70] J. Shi and S. Ohnuma. Decoherence and recoherence of beam in phase space. In *Proceedings of the 1993 Particle accelerator conference*, volume 5, pages 3603–3605, 1993.

[71] M.G. Minty, A.W. Chao, and W.L. Spence. Emittance growth due to decoherence and wake fields. In *Proceedings of the 1993 Particle accelerator conference*, volume 5, pages 3037–3039, 1995.

[72] G. P. Jackson. The effect of chromatic decoherence on transverse injection oscillation damping. In *Proceedings of Technical Workshop on Feedback Control of Multi-Bunch Instabilities in Proton Colliders at the Highest Energies and Luminosities*, pages 117–130, 1992.

[73] K. Manukyan *et al.* Energy spread extraction and processing from the kicked beam centroid motion in electron storage rings. *Journal of Instrumentation*, 6(10):T10003, 2011.

[74] K. Astrelina and A. Petrenko. Measurement of beam parameters in the VEPP-5 damping ring using betatron oscillations decoherence. In *Proceedings of RUPAC 2012*, pages 385–387, 2012.

[75] B. Jia *et al.* Electron beam energy spread measurements using optical klystron radiation. *Phys. Rev. ST Accel. Beams*, 13:080702, 2010.

[76] E. Tarazona and P. Elleaume. Measurement of the absolute energy and energy spread of the ESRF electron beam using undulator radiation. *Rev. Sci. Instrum.*, 67(9):3368–3368, 1996.

[77] S. Kashiwagi *et al.* Photo-electron beam longitudinal phase space tomography studies at the ATF. In *XX International Linac Conference*, pages 149–151, 2000.

[78] K.N. Ricci, T.I. Smith, and E.R. Crosson. Sub-picosecond electron bunch profile measurement using magnetic longitudinal dispersion and off-phase RF acceleration. In *The eighth workshop on advanced accelerator concepts. AIP Conf. Proc.*, volume 472, pages 735–744, 1999.

[79] E. R. Crosson *et al.* The determination of an electron beam's longitudinal phase space distribution through the use of phase-energy measurements. *Nuclear Instruments and Methods in Physics Research A*, 375:87–90, 1996.

[80] G. Rumolo and R. Tomas. Decoherence of a longitudinally kicked beam with chromaticity. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 528(3):670 – 676, 2004.

[81] B. D. Muratori *et al.* Preparations for EMMA commissioning. In *Proceedings of IPAC 2010*, pages 4337–4339, Kyoto, Japan, 2010.

[82] M. G. Ibison *et al.* ALICE tomography section: Measurements and analysis. *Journal of Instrumentation*, 7(04):P04016, 2012.

[83] P.B. Wilson. Transient beam loading in electron - positron storage rings. Technical Report PEP-0276, CERN-ISR-TH-78-23, 1978.

[84] S. Machida. Longitudinal emittance blowup in fixed field alternating gradient muon accelerators. *Phys. Rev. ST Accel. Beams*, 9:104002, Oct 2006.

[85] S. Fukumoto, K. Endo, K. Muto, T. Kitagawa, T. Inada, H. Tsujii, Y. Takada, A. Maruhasi, Y. Hayakawa, J. Tada, and S. Sato. Design of a synchrotron for proton therapy. In *Particle Accelerator Conference, 1989. Accelerator Science and Technology., Proceedings of the 1989 IEEE*, pages 700–702 vol.1, March 1989.

[86] K. Noda, M. Kanazawa, A. Itano, E. Takada, M. Torikoshi, N. Araki, J. Yoshizawa, K. Sato, S. Yamada, H. Ogawa, H. Itoh, A. Noda, M. Tomizawa, and M. Yoshizawa. Slow beam extraction by a transverse RF field with AM and FM. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 374(2):269 – 277, 1996.

[87] L. Badano, M. Benedikt, P.J. Bryant, M. Crescenti, P. Holy, A.T. Maier, M. Pullia, S. Rossi, and P. Knaus. Proton-ion medical machine study (PIMMS), 1. (CERN-PS-99-010-DI), 1999.

[88] T. Furukawa, M. Kanazawa, M. Muramatsu, K. Noda, S. Shibuya, E. Takada, and S. Yamada. Progress of RF-knockout extraction for ion therapy. In *Proceedings of EPAC 2002*, Paris, France, 2002.

[89] Particle Therapy Co-Operative Group. Facilities in operation. `http://www.ptcog.ch`, 2014. [Online; accessed 01-December-2014].

[90] M. Aiba, S. Machida, Y. Mori, and S. Ohnuma. Resonance crossing experiment at a proof of principle fixed field alternating gradient accelerator. *Phys. Rev. ST Accel. Beams*, 9:084001, 2006.

[91] T. Yokoi, J. Cobb, K. Peach, and S. Sheehy. Beam acceleration studies of proton NS-FFAG. In *Proceedings of EPAC 2008*, Genoa, Italy, 2008.

[92] R.D. Ruth. Single particle dynamics and nonlinear resonances in circular accelerators. 1985.

[93] F.E. Mills. Technical assessment of the Loma Linda University proton therapy accelerator. 1989.

[94] J.R. Alonso. Proton therapy construction projects in the United States. 1992.

[95] M.H. Phillips, E. Pedroni, H. Blattmann, T. Boehringer, A. Coray, and S. Scheib. Effects of respiratory motion on dose uniformity with a charged particle scanning method. *Physics in Medicine and Biology*, 37(1):223, 1992.

[96] J. Lambert, N. Suchowerska, D.R. McKenzie, and M. Jackson. Intrafractional motion during proton beam scanning. *Physics in Medicine and Biology*, 50(20):4853, 2005.

[97] S.M. Zenklusen, E. Pedroni, and D. Meer. A study on repainting strategies for treating moderately moving targets with proton pencil beam scanning at the new gantry 2 at psi. *Physics in Medicine and Biology*, 55(17):5103, 2010.

[98] T. Furukawa, T. Inaniwa, S. Sato, T. Shirai, Y. Takei, E. Takeshita, K. Mizushima, Y. Iwata, T. Himukai, S. Mori, S. Fukuda, S. Minohara, E. Takada, T. Murakami, and K. Noda. Performance of the NIRS fast scanning system for heavy-ion radiotherapy. *Medical Physics*, 37(11):5672–5682, 2010.

[99] S. Sato, T. Furukawa, and K. Noda. Dynamic intensity control system with rf-knockout slow-extraction in the {HIMAC} synchrotron. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 574(2):226 – 231, 2007.

[100] Y. Li, X. Li, H. Li, W. Cao, R.X. Zhu, N. Sahoo, M. Gillin, X. Zhang, L. Kardar, L. Liao, G. Lim, J.Y. Chang, Z. Liao, R. Komaki, and J.D. Cox. On the interplay effects with proton scanning beams in stage III lung cancer. *Medical Physics*, 41(2), 2014.

[101] Y. Tsunashima. *Verification of the clinical implementation of the respiratory gated beam delivery technique with synchrotron-based proton irradiation.* PhD thesis, The University of Texas Graduate School of Biomedical Sciences, 2012.

[102] C. Bert, S.O. Grözinger, and E. Rietzel. Quantification of interplay effects of scanned particle beams and moving targets. *Physics in Medicine and Biology*, 53(9):2253, 2008.

[103] D. Trbojevic *et al.* Lattice design of a rapid cycling medical synchrotron for carbon/proton therapy. In *Proceedings of IPAC 2011*, San Sabástian, Spain, 2011.

[104] L. Honghuan. *Design of a Compact Medical Synchrotron and 3D Dose Delivery of Rapid Cycling Beam.* PhD thesis, Indiana University, 2013.

[105] T. Haberer, W. Becher, D. Schardt, and G. Kraft. Magnetic scanning system for heavy ion therapy. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, A330:296–305, 1993.

[106] T. Inaniwa, T. Furukawa, T. Tomitani, S. Sato, K. Noda, and T. Kanai. Optimization for fast-scanning irradiation in particle therapy. *Medical Physics*, 34(8):3302–3311, 2007.

[107] T. Furukawa, T. Inaniwa, S. Sato, T. Shirai, Y. Takei, E. Takeshita, K. Mizushima, Y. Iwata, T. Himukai, S. Mori, S. Fukuda, S. Minohara, E. Takada, T. Murakami, and K. Noda. Performance of the NIRS fast scanning system for heavy-ion radiotherapy. *Medical Physics*, 37(11):5672–5682, 2010.

[108] A. Schätti, D. Meer, and A.J. Lomax. First experimental results of motion mitigation by continuous line scanning of protons. *Physics in Medicine and Biology*, 59(19):5707, 2014.

[109] Y. Iwata, T. Kadowaki, H. Uchiyama, T. Fujimoto, E. Takada, T. Shirai, T. Furukawa, K. Mizushima, E. Takeshita, K. Katagiri, S. Sato, Y. Sano, and K. Noda. Multiple-energy operation with extended flattops at HIMAC. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 624(1):33 – 38, 2010.

[110] K. Noda, T. Furukawa, T. Fujimoto, Y. Hara, T. Inaniwa, Y. Iwata, K. Katagiri, N. Kanematsu, K. Mizushima, T. Miyoshi, S. Mori, T. Murakami, Y. Sano, S. Sato, T. Shirai, E. Takada, Y. Takei, and S. Yonai. Recent progress of HIMAC for sophisticated heavy-ion cancer radiotherapy. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 331(0):6 – 9, 2014. 11th European Conference on Accelerators in Applied Research and Technology.

[111] K. Mizushima, K. Katagiri, Y. Iwata, T. Furukawa, T. Fujimoto, S. Sato, Y. Hara, T. Shirai, and K. Noda. Experimental studies of systematic multiple-energy operation at HIMAC synchrotron. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 331(0):243 – 247, 2014. 11th European Conference on Accelerators in Applied Research and Technology.

[112] D. Trbojevic, R. Gupta, B. Parker, E. Keil, and A.M. Sessler. Superconducting non-scaling FFAG gantry for carbon/proton cancer therapy. In *Particle Accelerator Conference, 2007. PAC. IEEE*, pages 3199–3201, June 2007.

[113] S. Tygier, R.B. Appleby, J.M. Garland, K. Hock, H. Owen, D.J. Kelliher, and S.L. Sheehy. The pyzgoubi framework and the simulation of dynamic aperture in fixed-field alternating-gradient accelerators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 775(0):15 – 26, 2015.