Imperial College London

Laser-hybrid Accelerator for Radiobiological Applications (LhARA)

John Adams Institute Accelerator Design Project 2024

April 10, 2024

G. Passarelli[1], M.Pereira[1], C. Jolly[2], S. Leadley[2], C. Lehmann[2], S. Preston[2], L. Bradley[3], G. Christian[3], J. Hills[3] and L. Kennedy[3]

[1] *John Adams Institute for Accelerator Science, Royal Holloway, University of London, Egham, United Kingdom*

[2] *John Adams Institute for Accelerator Science, University of Oxford, Oxford, United Kingdom*

[3] *John Adams Institute for Accelerator Science, Imperial College London, London,United Kingdom*

Abstract

An evaluation of Stage 1 of the proposed LhARA beamline and a series of design suggestions is presented. An investigation of lattice design with the aim of minimising end station spot size and beam losses was made using MADX and BDSIM; by varying Gabor lens and quadrupole strength a minimum spot size of just under 1 cm is achievable. Loss studies in BDSIM were carried out with the baseline 3 cm beam both to inform the design of the beamline, and to validate the effectiveness of elements such as the beam pipe aperture and collimators, in order to calculate end station dose rates and quantify their impact. The dose rate for a 3 cm beam through the optimised design was calculated at $15.12\,\mathrm{Gy\,s^{-1}}$. A single octupole was inserted before the arc in simulation to assess how effectively bunch uniformity could be improved at the end station, where improvement was shown with an increase in the correlation coefficient to $r_p = 0.60$ on a QQ plot, versus $r_p = 0.45$ in the unperturbed case. A comparison between solenoid models and Gabor lenses demonstrates not only the efficacy but also the comparative efficiency of Gabor lenses.

A design for the two RF bunching cavities was proposed via a geometry optimisation algorithm informed by longitudinal phase-space simulations. The best achievable bunching was found to reduce the beam length from 4 ns to 1.18 ns, reducing the energy spread from the nominal 2% used for the simulation to 0.3%. The bunching was found to be less effective for carbon ions. A 3D model and simulation was created in CST, with the resulting field map used for a more effective simulation in BDSIM.

A study into magnet architecture to meet the parameters set out by the aforementioned MADX optimisations yielded designs for a dipole, a quadrupole and an octupole that met or exceeded the required magnetic field strengths at 99.8% purity or above. A design for a permanent magnet quadrupole was also developed to increase beam capture from the initial laser-target interaction.

Acknowledgements

This report represents the capstone project of the John Adams Institute coursework and lecture series for 1st year PhD and DPhil students. It is the culmination of 3 months of work by us students, and many more months of effort from others behind the scenes. We would like to express our deepest gratitude to all those who gave their time and energy to helping us through this project, both for the advice and the management they have provided us. In particular we would like to highlight the major ongoing contributions from:

- Emmanuel Tsesmelis
  University of Oxford, John Adams Institute for Accelerator Science
- William Shields
  Royal Holloway University of London, John Adams Institute for Accelerator Science
- Kenneth Richard Long
  Imperial College London
- Ciprian Plostinar
  European Spallation Source, John Adams Institute for Accelerator Science
- Hossein Ghasem
  Diamond Light Source

Thanks should also go to those academics who have generously provided us their knowledge and expertise: Attilio Milanese, Jeremie Bauche, Jaroslaw Pasternak, and all the lecturers for the JAI accelerators course. We are also grateful to Philip Burrows, for his management of the JAI program, without whom the project could not have happened.

Lastly, it would be remiss of us not to mention our PhD and DPhil supervisors, some of whom have already been listed above, but all of whom deserve our thanks for their ongoing support and understanding in this project.

Contents

List of Figures

List of Tables

Chapter 1

Motivation

Over half of all cancer patients require radiation therapy (RT), however access to RT is not universal and varies between different income counties. Access is especially limited for hadron therapy; by the end of 2021 only 20 countries had hadron therapy facilities in clinical operation. This lack of accessibility can be attributed to the high cost of construction and running of these facilities. Additionally, ion beam radiobiology is currently under-explored and our understanding of optimal dosage between ion species is still underdeveloped. Further to this there are novel regimes of proton beam delivery which show promise, but are yet unproven in a clinical setting. These include beam delivery at FLASH dose rates doses and delivery of multiple spaced out microbeams, both of which have evidence of therapeutic benefits but require further study to understand their effects.

1.1 Radiobiological Study

LhARA presents an opportunity to investigate the Relative Biological Effectiveness (RBE) of hadron therapy and develop our Radiobiological Models (RBMs). RBE is the ratio of a given dose, delivered by photons, compared to the required dose of another particle species for the same radiobiological effect. At present the RBE value for protons widely used in clinical treatment planning is 1.1 [1]: this value has been clinically effective, but its accuracy is

contentious as RBE values exceeding 1.5 and even approaching 3.0 have been calculated from *in-vivo* studies [2]. This spectrum of values comes hand-in-hand with discussion of the exact effects of individual beam parameters on the final RBE value, and debate on the ideal RBE for treatment plans [3]–[5].

The delivered dose governs the likelihood of damage to the tumour in addition to the impact on healthy tissue around the tumour site. An inaccurate RBE can lead to an overestimation of the dose, negatively impacting healthy tissue and possibly causing further mutation of normal cells. An underestimate of dose insufficiently damages the tumour and allows for further proliferation. These concerns are valid for proton beam therapy, but even more so when it comes to other ion species such as carbon, which has demonstrated RBE values across the range of 2.5 to 5.0 in animal cell studies [6], [7]. Such RBE values display promise for more effective beam therapy modalities, and are of great interest for eventual clinical translation, however they lack widely available data to an even greater extent that proton therapy [7], [8].

Improving radiobiological models (RBMs) is fundamental to planning out treatment, as RBMs are used to forecast radiobiological damage caused by ionising radiation and determine the biological dose received by the patient [9], [10]. The RBM takes into account biological factors such as oxygenation and proliferation characteristics of the tumour, as well as the resistance of different tumours and normal tissues to radiation. There is still much to be investigated concerning the exact mechanisms governing the interactions between ions and the biological environment. How these factors influence hadron therapy effectiveness across a broad variety of beam profiles is deeply important to future treatment planning, and so our understanding of RBMs must be fed, again, with varied data from a range of different beam parameters.

1.2 Novel Regimes

Irradiation in ultra-high dose rate regimes (known as FLASH radiotherapy) using conventional X-ray photons, electrons and protons has demonstrated positive outcomes for both tumour control and healthy tissue survival [11], [12]. Electrons have been used to treat a human

patient in this manner with results comparable to conventional treatment, but across a much shorter treatment time [13]. An additional technique beginning to be explored is micro-beam treatment, where an even spread of smaller beams has also demonstrated clinical efficacy.

1.3 Laser-hybrid Beams for Radiobiology and Clinical Applications

Recent research efforts have pushed thinking towards laser-driven particle sources for radiobiological applications [14], [15]. Facilities such as SCAPA [16] in Scotland have several beamlines dedicated to achieving a greater understanding of laser-plasma interactions at the core of these new ion sources, and the ELI facility in Czechia is in the process of establishing ELIMED (ELI MEDical and multidisciplinary applications) [17], a dedicated section investigating cell irradiation from laser-driven hadron sources [18]–[20].

Laser-driven sources enable particle capture at energies above the space-charge limit, evading emittance growth effects. As such, this technology is of particular interest for radiobiological and clinical applications. However there are still challenges in producing clinically viable laser-driven sources, one example being the high divergence of the beam coming from the source. Such design challenges will need to be addressed for a research facility to be established.

Chapter 2

Introduction

2.1 LhARA overview

The Laser-hybrid Accelerator for Radiobiological Applications (LhARA) is a proposed accelerator facility with the aims of providing a flexible dedicated proton beam-line for the investigation of particle-beam therapy. Through the use of high power lasers in generating protons, LhARA aims to provide beams for hadron therapy of reasonable scale and cost. LhARA will make use of novel acceleration techniques in a flexible beam line that grants access novel irradiation regimes and allow for variation in delivered dosage, ion species, energy, time profile and spatial profiles. LhARA distinguishes itself from other facilities and from other previous laser-driven cell irradiation experiments as being a facility entirely dedicated to radiobiological research. For example, while previous irradiation experiments have been 'single shot' in their nature [21], LhARA is designed to produce a shot-to-shot stability in the particle flux which is more appropriate for radiotherapy and enables more accurate dosimetry predictions and calculations to advance radiobiological research.

Alongside further study of new treatment modalities already under investigation, LhARA will also facilitate a breadth of non-conventional beams to be produced with a variety of spatial, temporal and spectral characteristics, with present simulations estimating LhARA as capable of particle rates on the order of 10^9 particles per shot corresponding to average dose rates of

up to and exceeding $120\,\mathrm{Gy\,s^{-1}}$ for protons. Studies at the LhARA facility would allow for comparison of the radiobiological response to unique dose rates and fractionation schemes for many ion species and on a variety of tissue types. Not only to achieve the previously motivated exhaustive examinations of RBE at endpoints other than standard cell survival curves, but also to assess hypotheses surrounding the unknowns of radiobiological mechanisms within these experimental regimes.

This proposed radiobiological facility has been designed as a laser-hybrid accelerator. It is composed of a laser-driven proton and ion source, coupled onto a conventional particle accelerator lattice. It seeks to utilise new technologies such as the laser-driven target and a solenoid-like focusing component named the Gabor lens to achieve greater control over beam emittance, timing and energy than is achievable from a conventional linear accelerator. The control over the particles delivered to the end stations will be key to the quality of radiobiological research conducted at the facility.

2.2 Laser Target Interactions

LhARA has been designed to use a laser-driven proton and ion source. A short pulse laser hitting the tape drive target generates a proton beam of energy approximately 15 MeV. A key advantage to this relatively new technology includes the high energy of the particles coming from the source. Traditional linear accelerators see significant beam emittance growth in accelerating particles from low energies, but generating particles immediately at this energy bypasses this consideration. The compact nature of the source is also a key advantage.

The target at LhARA will utilise target normal sheath acceleration, in which a short intense laser pulse draws electrons out of an over-dense target, generating a separation of charge and an opposing field. This intense field accelerates electrons through the target generating a "sheath" field on the target rear. This field generates energetic ions and protons, with energy scaling as $E_p \propto I^{\frac{1}{2}}$. To facilitate the production of beams at around 15 MeV and to reduce flux variability a 100 TW laser system with 2.5 J pulses and pulse duration's below 25 fs will be used. A tape

drive spools a thin metal (aluminium or steel) foil to replenish the target surface. A tape drive for this system was developed at Imperial with reliable operation at up to 10 Hz with targets as thin as 5 µm.

A drawback of this technique is that particles emerge from the target with high divergence and with a large energy spread. To combat this, focusing and longitudinal phase-space manipulation are necessary. For the LhARA facility, it is planned to use Gabor lenses to provide the level of focusing needed to capture the particles from this source.

2.3 Gabor lenses

Gabor lenses, as shown in [22], have been proposed as a cost-effective alternative to superconducting solenoid magnets for the beam capture and energy matching section of the lattice. Instead of using strong magnetic fields as would be the case for solenoid magnets, Gabor lenses provide focusing through the electric fields generated by a trapped non-neutral plasma. They have the potential to provide a level of focusing equivalent to a solenoid, but with a significant reduction in cost and power consumption. A schematic of a Gabor lens can be seen in Figure 2.1.

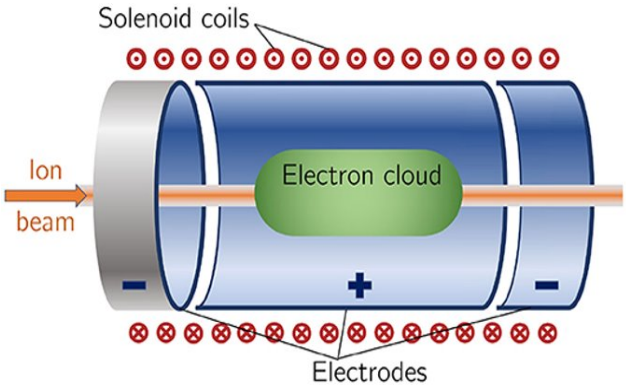

Figure 2.1: Schematic of a Gabor lens.

The operating principle of the Gabor lens is derived from applying Gauss' law for a cylindrical surface to the charge contained within the plasma, as shown in Equation 2.1

$$E_r = \frac{-en_e r}{2\epsilon_0} \tag{2.1}$$

Confinement of a uniform charge distribution will provide radially symmetric focusing for a particle bunch of the opposite sign, with a focal length given by Equation 2.2 [22], [23]. As such, for a Gabor lens the focusing comes from the plasma, not from the solenoid field.

$$\frac{1}{f} = \frac{e^2 n_e}{4\epsilon_0 U} l \tag{2.2}$$

Where e is the electric charge, n_e is the the number density of the electron cloud, U is the kinetic energy of the proton beam, and l is the effective Gabor length.

Electrodes are configured to provide axial confinement of the plasma with a series of solenoid wires providing radial confinement. The solenoid field required to confine the plasma is given by Equation 2.3 [23].

$$B_{\mathrm{GBL}} = B_{\mathrm{sol}} \sqrt{Z \frac{m_e}{m_p}} \tag{2.3}$$

In the case of LhARA, $\sim$0.03 T is the field required to contain the plasma, significantly reducing the overall magnetic field strengths and power required to focus the beam as compared to a traditional solenoid. The LhARA design intends to use Gabor lenses extensively to focus the beam coming from the laser-driven ion source and control the beam size delivered to the end stations.

2.4 Facility description

The LhARA facility in its current design iteration can be separated into two distinct stages, which facilitate the delivery of beam to three end stations. This has been designed in collaboration with radiobiology researchers to meet the beam energies and facility structure requirements required for testing. An overview of the facility and the two stages can be seen in Figure 2.2.

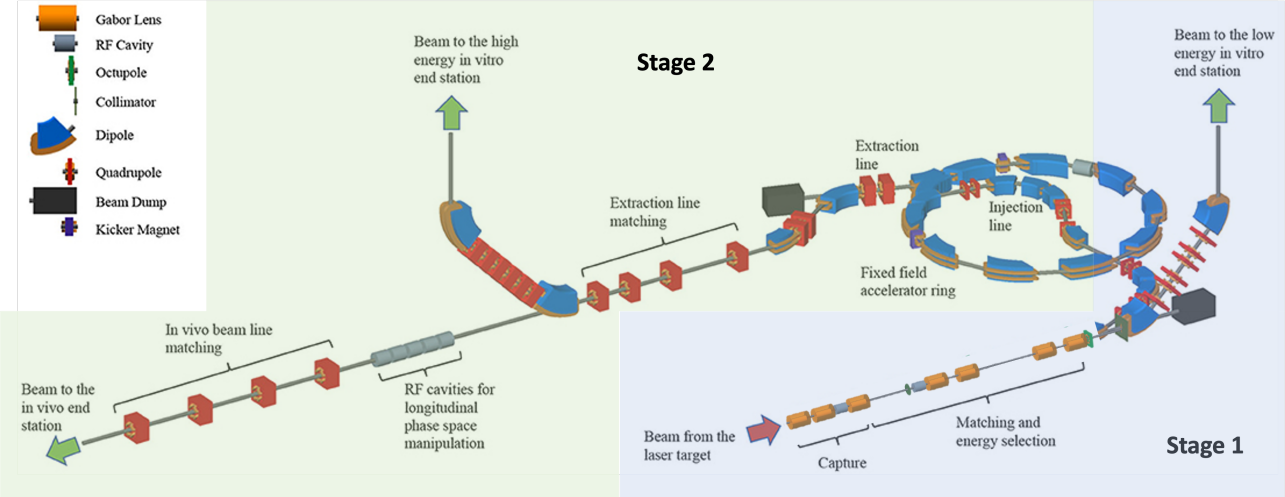

Figure 2.2: LhARA Stages 1 and 2 Layout.

2.4.1 LhARA Stage 1 Description

LhARA Stage 1 can be broken down into five distinct sections. Section 1 has the purpose of capturing the bunches emitted from the ion and proton source and focusing them into a stable bunch. It achieves this with two Gabor lenses, which minimise the transverse momentum of the particle beam. The second section is for matching and energy selection. Within this section another Gabor lens focuses the beam down to a small spot size. This section also has two well-spaced RF cavities. These are used to maintain the longitudinal bunching, and also perform manipulations on the longitudinal phase-space. Four more Gabor lenses bring the beam to parallel before entering the next section. The third section is for beam shaping. An octupole will provide third-order focusing to perturb the first-order focusing of the Gabor lenses. It is positioned directly after Gabor lenses where the beam is large, so as to have maximum effect. Since the Octupole causes a rectangular transverse distribution in the beam, a collimator is

then used to once again match the beam to the circular aperture.

We then have the first switching dipole, which can optionally bend the beam horizontally towards the fixed field accelerator ring of Stage 2, or alternatively continue towards the low energy *in-vitro* end station. A second switching dipole can optionally bend the beam vertically towards the low energy *in-vitro* end station or allow it to continue down an abort line.

The final stage is the vertical matching arc sending the beam towards the low energy *in-vitro* end station.This section contains six quadrupoles form an achromat matching section with transparent optics. By choosing an achromat structure, dispersive effects are cancelled out. It ends with a final dipole to turn an additional 45°, completing the full 90° vertical arc. A 2 m long drift tube then penetrates the end station floor to deliver the beam at bench height.

2.4.2 LhARA Stage 2 Description

Within Stage 2, the focused 15 MeV beam is injected into the fixed field accelerator ring. This then provides further acceleration to a proton energy of 127 MeV. On extraction, bunches are then transported to both the high energy *in-vitro* end station, and to the *in-vivo* end station.

2.4.3 End Stations

The low energy *in-vitro* exit station of stage 1 requires a low uniform dose distribution (<5%), whilst minimising beam losses for machine protection and maximising dose delivery [24]. The proton beam *in-vitro* exit has a maximum spot size of 1-3 cm, with the transverse proton momentum following a Gaussian distribution with energy spread at full width ≤4% (low-energy end station) and ≤ 1% (high-energy end station). Protons with kinetic energy less than 10 MeV don't reach the cell layer, whilst protons with 12 MeV have a Bragg peak at the cell layer and 15 MeV proton beams have a Bragg peak past the cell layer. Similarly, the *in-vivo* biological end station requires a maximum input beam diameter of 1-3 cm, full width energy spread ≤1%, and an input beam uniformity of <5%. The final required maximum dose rates are given in Table 2.1 [24].

	12 MeV protons	15 MeV protons	127 MeV protons
Dose per pulse	7.1 Gy	12.8 Gy	15.6 Gy
Instantaneous dose rate	1.0×10^9 Gy s^{-1}	1.8×10^9 Gy s^{-1}	3.8×10^8 Gy s^{-1}
Average dose rate	71 Gy s^{-1}	128 Gy s^{-1}	156 Gy s^{-1}

Table 2.1: Dose rates required by LhARA.

2.5 Report scope

The purpose of this report is to provide further details and optimisation for stage 1 of the LhARA facility design. The design for stage 2, along with details of the laser target ion source, are out of the scope of this report and will be assessed at a later date. The investigative work has been divided into 3 sections: lattice design, RF cavity design, and magnet design. These are respectively outlined in Sections 3, 4, and 5 of this report.

Chapter 3

Lattice Design

Significant work has already been undertaken in lattice design for LhARA, but many parts have not been optimised. The scope for a flexible LhARA beamline extends to having variable beam sizes at the end stations. To achieve smaller beam sizes requires optimisation of the Gabor lenses and quadrupoles in the lattice. Analysing the effects of the changes made to the beam line in terms of beam loss through the lattice, alongside tracking and understanding the behaviour of different beamline components (including characterising the impact of the octopole on bunch flatness), was a priority for this study. In addition, the final dosage produced by this beamline was important to characterise. Justification of the model's use of solenoids as a substitute for Gabor lenses was required, as well as determining the relative magnetic field required to reproduce the simulated beam using Gabor lenses instead of the model's solenoids.

3.1 Methods

3.1.1 MAD-X

The base model for lattice studies was constructed in the Methodical Accelerator Design tool (MAD-X) [25]. MAD-X is a general purpose accelerator design tool with a focus on beam dynamics and optics optimisation, and facilitated studies of the overall optics and the behaviour

of the Twiss parameters along LhARA stage one with relative ease-of-use and short simulation times. As these two factors were the only ones under examination in the MAD-X studies, only those elements designed to manipulate them were modelled fully. Elements such as the RF cavities, collimators and octupole were modelled as drifts only accurate in overall length.

Aside from constructing a model of the lattice, MAD-X's matching module was of particular interest, finding solutions for the parameters of specified elements to satisfy user defined constraints on the beam parameters. In this case, these constraints not only allowed a particular beam size at the end of the sequence to be targeted, but also ensured the beam was prepared at various points along the lattice for the unmodelled elements that will be present in reality.

Lattice Design Constraints

The first constraint was for the beam to be focused to a minimal beam size after Gabor lens 3, at the position of the first collimator. Focusing to this minimum for all configurations ensures the fixed aperture of the energy cleaning collimator is consistently effective for all configurations of the beamline. This constraint is entirely satisfied by fixing the strength values, K_S, of Gabor lenses 1-3 for all configurations.

To limit any impact of the first RF Cavity on beam optics, Twiss α must be zero between Gabor lenses 2 and 3. $\alpha = 0$ represents a beam that is neither growing nor shrinking in overall size, often referred to as being 'parallel'. A parallel beam should only have momentum in the longitudinal axis, and so any kick received by a particle in the RF field should not effect the transverse properties of the beam.

The third requirement on the lattice is twofold, and these two constraints are paired as they are a means to the same end. At the midpoint of the vertical arc, dispersion must be high and Twiss β low to ensure the efficacy of the momentum cleaning collimator positioned at this point in the lattice. As dispersion rises, the transverse profile of the beam stretches in the dispersive direction (in this case y). A low Twiss β corresponds to a small beam size, and is required here for similar reasons as a small beam size is required between Gabor lenses 2 and 3.

The collimator aperture here is designed to be some size $n\sigma$ of the beam (without considering dispersive terms), hence those particles dispersed due to their momentum are cleaned away.

It is intended that the vertical arc has 'transparent' optics, meaning that the profile of the beam that enters should be the same as the profile of the beam that leaves. To ensure this, the final two constraint are put in place. Twiss β must be at its targeted value, corresponding to the desired beam size, when it enters the arc, and to ensure that the beam is not diverging from this β value, Twiss α must be zero before and after the arc.

The initial state of the MAD-X model was configured for a 3.0 cm spot size (2σ diameter). This baseline configuration satisfied all constraints (Figure 3.1) and produced end station spot sizes of $\sigma = 3.0\,\text{cm}$.

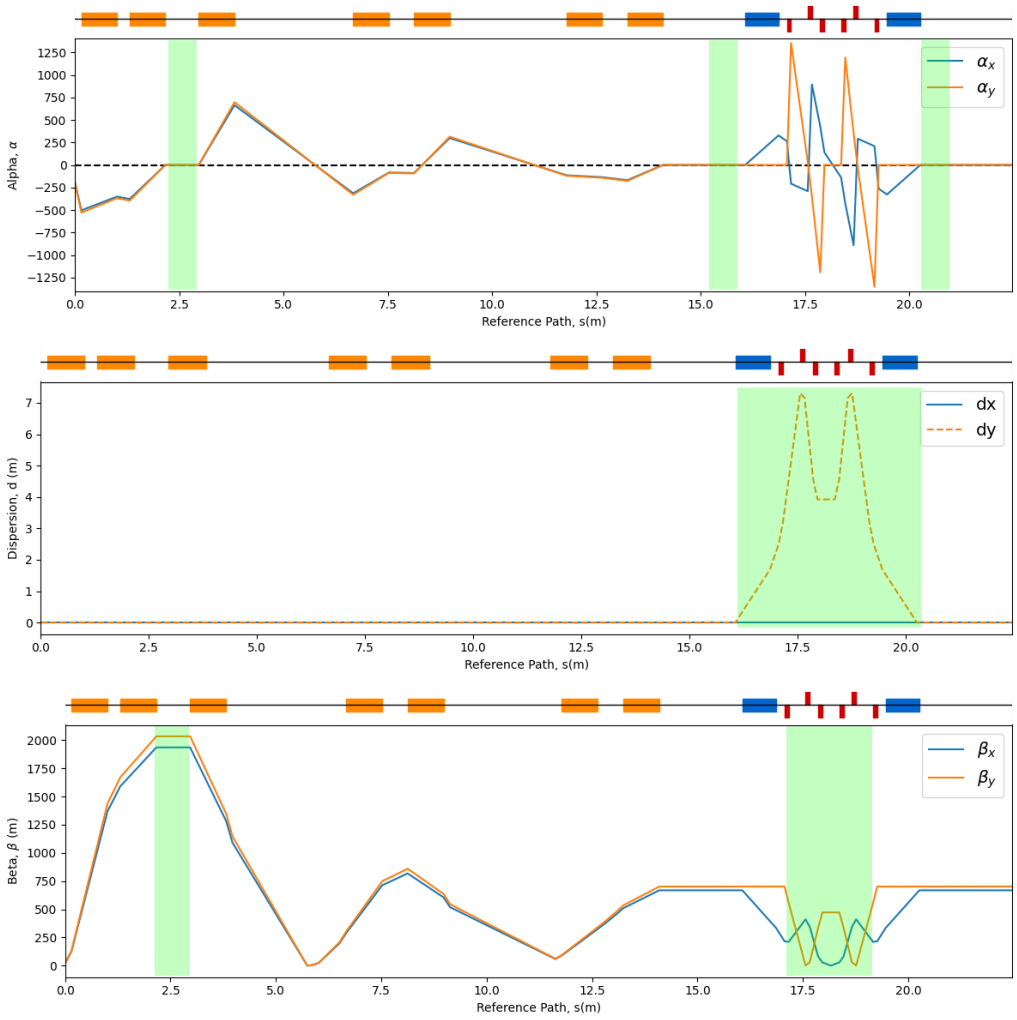

Figure 3.1: Variations of Twiss α, dispersion and Twiss β, respectively, along the MAD-X lattice. Green highlights demonstrate the satisfaction of the previously established constraints.

Matching Approach

With the constraints established and a set of baseline solenoid strength values for a 3.0 cm beam, MAD-X's 'match' module could be utilised to achieve smaller spot sizes. The matching routine for spot size optimisation utilised the same MAD-X lattice as the baseline configuration, with the same beam conditions; the only thing that was varied was the strengths of the solenoid elements representing the Gabor lenses. As mentioned previously, Gabor lenses 1-3 were fixed for all configurations and thus for the matching routine in MAD-X only Gabor lenses 4-7 were varied. Limits must be set for the varied parameters, and so a maximum K_S of 2.49 was calculated for 15 MeV protons, with a maximum solenoid fields of 1.4 T.

MAD-X's match module is capable of using different methods to match the parameters to the target outcome - for all MAD-X optimisation studies here the fast gradient minimisation method was used. Fast gradient minimisation seeks to minimise the sum of squares of the constraint functions using their numerical derivatives over the course of a user defined number of calls, and to within some user-defined tolerance. It is the fastest method available in MAD-X, which was an advantage to this design project due to its finite time scale. As this method produced positive results, and preliminary tests of other methods did not yield any further results beyond the limits of optimisation found in these studies, it was the only method used.

Gabor lens and quadrupole strengths were varied to satisfy constraints 1-5. To achieve the first constraint for all configurations, the first three Gabor lenses had to be treated as a separate system. For the initial round of matching only the latter three solenoids were varied, therefore only constraints 3-5 were required. The Gabor lenses were initially approximated and modelled as solenoids at the project outset, with the Gabor lenses being modelled and imported to BDSIM later. A comparison study between the solenoids and Gabor lenses were made in BDSIM, showing comparable field strengths to achieve the same optics.

3.2 Results

3.2.1 MAD-X Matching

MAD-X Gabor Lens (Solenoid) Optimisation

Varying solenoid strengths using the matching module allowed access to a variety of spot sizes from 3.0 cm down to 1.0 cm, with no limits observed anywhere in the region between those two values.

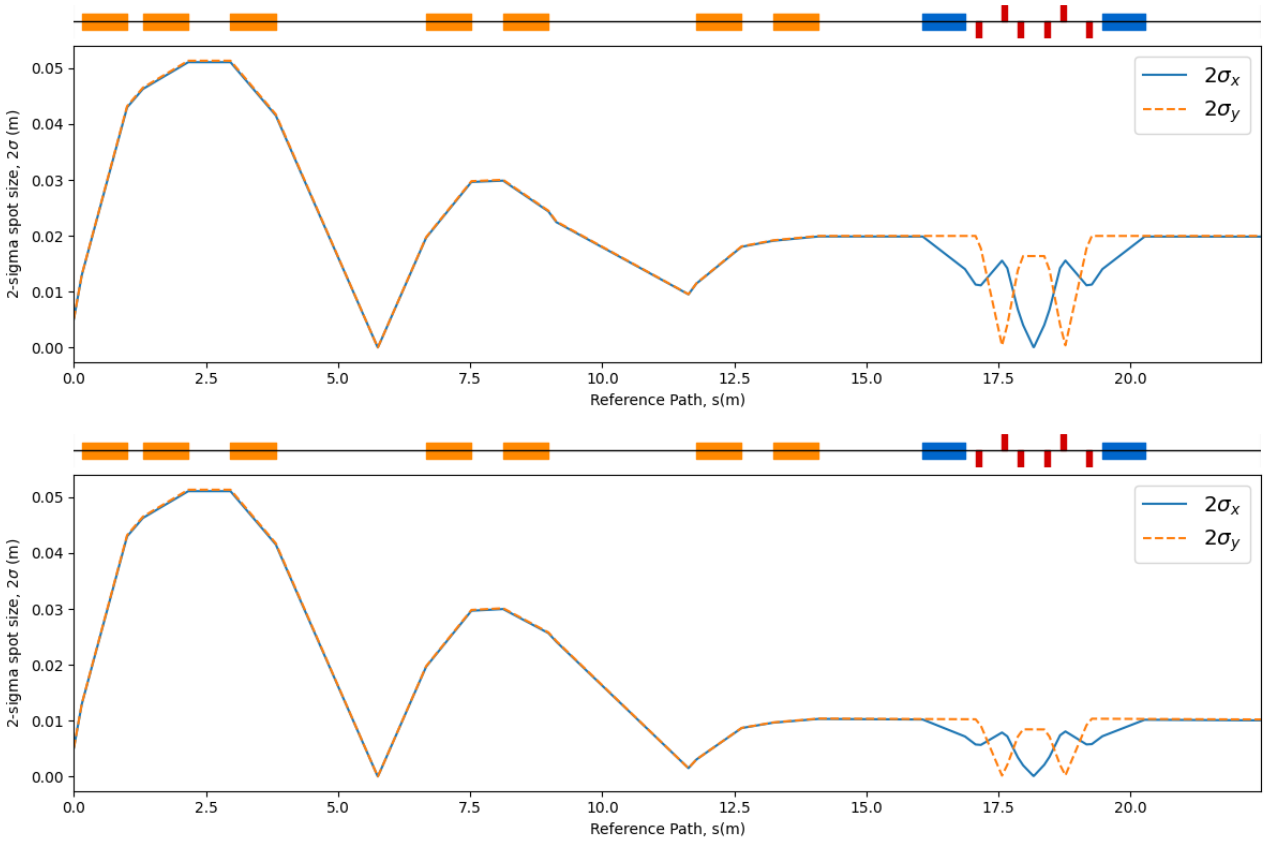

Figure 3.2: Spot size (2σ-diameter) against s for the 2.0 cm (top) and 1.0 cm (bottom) configurations of the MAD-X lattice. Both configurations achieved their target spot sizes whilst satisfying all the lattice constraints. The K_S values for each configuration can be found below in Table 3.1.

	Solenoid Strength, K_S (m^{-1})			
Spot Size (cm)	Lens 4	Lens 5	Lens 6	Lens 7
3.0	1.80	1.61	1.24	1.91
2.0	1.94	1.48	1.82	0.65
1.0	1.93	1.33	2.49	0.88

Table 3.1: Table of solenoid strength values, K_S, for lenses 4-7 in the MAD-X lattice for each spot size configuration (3.0 cm, 2.0 cm and 1.0 cm 2σ diameter).

While spot sizes from 3.0 cm to 1.0 cm could be achieved using the matching method, optimisations for spot sizes below 1.0 cm produced inaccurate values for 2σ diameter, or could only produce the desired values when certain constraints (majorly the constraints on Twiss α and dispersion before and after the arc) were not met or had their given tolerances set to a value too large to be satisfactory.

MAD-X Quadrupole Strength Optimisation

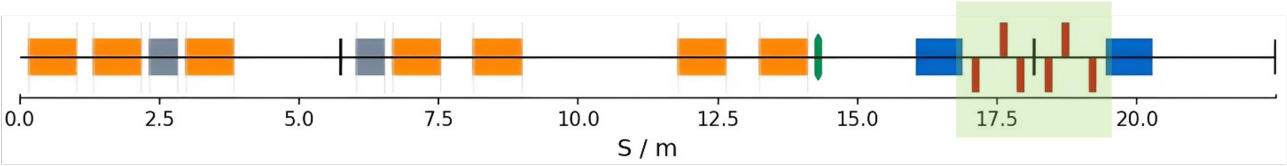

Figure 3.3: MAD-X lattice design with highlighted quadrupoles used to produce reduced beam sizes.

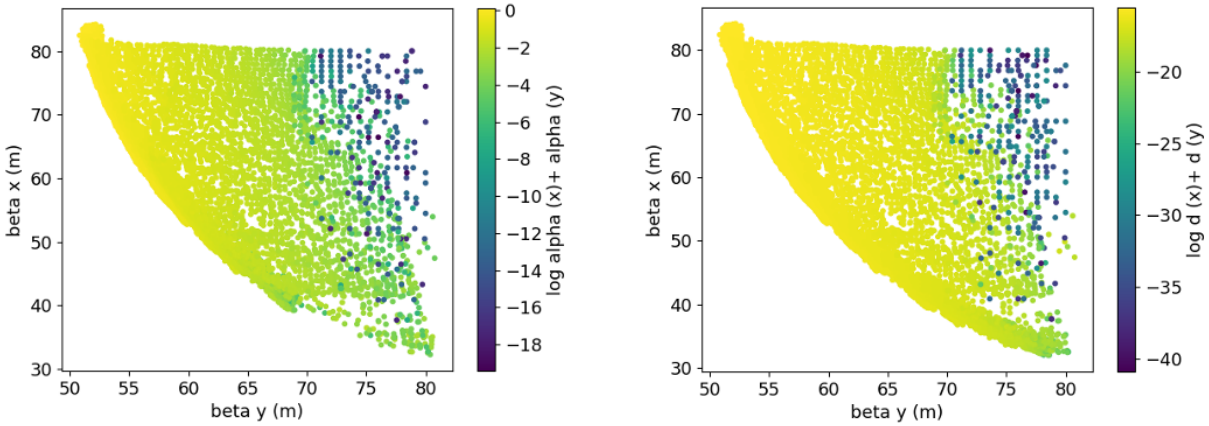

(a) End station Twiss α map.

(b) End station dispersion map.

Figure 3.4: Scan of achievable Twiss β values and MAD-X's matching abilities, with a map of dispersion and Twiss α values.

It was found that MAD-X matching struggled to match other constraints for Twiss β values below 70. It was also unable to achieve Twiss β values below 50 in general. A bias in the x direction was discovered, with the matching function able to satisfy more of the constraints whilst also reducing final spot size.

Quadrupole optimisation was able to produce further reductions to final spot size. A final spot size with Twiss β in both x and y of $\sim$75 being achievable with a suppressed dispersion of 10^{-7} and Twiss α of 10^{-8}.

Twiss β (x)	Twiss β (y)	End Station Twiss α (x)
75.7	74.4	2×10^{-7}
End Station Twiss α (y)	End Station Dispersion (x)	End Station Dispersion (y)
-1×10^{-6}	-2×10^{-8}	-2×10^{-26}
Quad 1 K_S (m^{-1})	Quad 2 K_S (m^{-1})	Quad 3 K_S (m^{-1})
-22.2	31.2	-31.9
Quad 4 K_S (m^{-1})	Quad 5 K_S (m^{-1})	Quad 6 K_S (m^{-1})
31.4	-22.7	-31.3

Table 3.2: Table of example quadrupole strength values, K_S, for Quadrupoles 1-6 in the MAD-X lattice for each spot size configuration optimal balance of constraints.

Further lattice studies were carried out in BDSIM (Beam Delivery Simulation) [26], a program that utilises the Geant4 physics libraries to simulate the transport of a particle beam through a 3D model of an accelerator with realistic physics processes.

3.3 BDSIM Studies

Through the associated python library, Pybdsim[27], the MAD-X model was directly converted to a model in BDSIM, and from there was be edited to facilitate further study. The overall aim of studying the lattice in BDSIM was to observe the effects of introducing more realistic physics. Firstly, where MAD-X simulates only a reference synchronous particle, the beam can be modelled more closely to the Gaussian distribution (expected from the laser source), and by making adjustments to both the physical parameters of the beamline and the way in which physical processes are modelled, several studies were carried out. These investigated:

energy loss and deposition with the introduction of a realistic beampipe aperture; modelling of, and dose rate calculations at, the end station; beam uniformity through the octupole; and a performance study comparing Gabor lenses to solenoids.

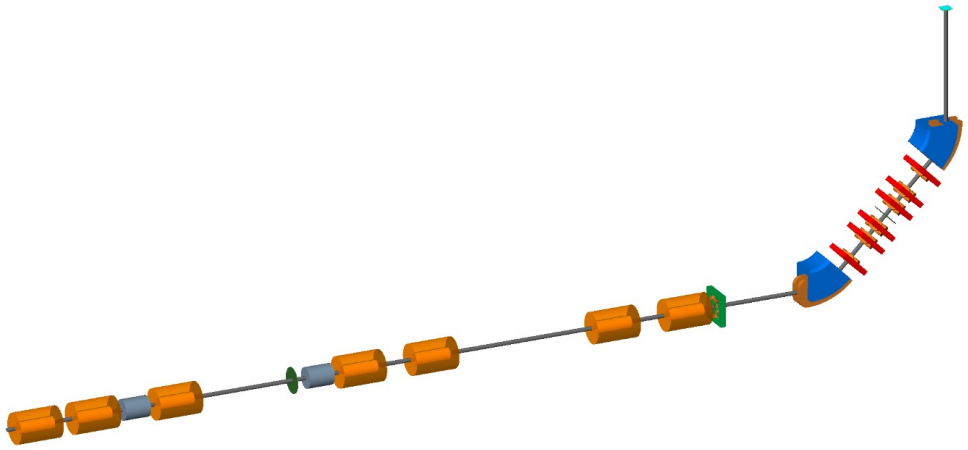

Figure 3.5: The model of the LhARA Stage 1 lattice in the BDSIM visualiser.

Each study required changes from the initial MAD-X converted baseline, including the introduction of some elements previously modelled as drifts, however one aspect that was kept the same for all studies was the previously investigated solenoid strengths. The BDSIM lattice was configured for a 3.0 cm spot size for all studies. The reason for this is twofold. Firstly, as 3.0 cm is the maximum spot size expected for LhARA stage 1, any beam loss studies should exhibit the maximum losses possible. Secondly, neither MAD-X nor BDSIM models space charge effects, which will be more significant at smaller beam sizes. Therefore, by using the largest beam, the simulations are the most accurate to any expected experimental data where space charge effects will be present.

3.3.1 Energy Loss and Deposition

The first studies carried out on the BDSIM lattice examined losses of particles along the beamline with the step-by-step introduction of realistic modelling for particular beamline elements, particularly those that were expected to introduce beam losses (e.g. apertures and collimators). Data for these studies was produced by simulations of 10^7 protons through the lattice in BDSIM.

The 'g4QGSP_BIC_EMZ' Geant4 physics list was used for these simulations and all further BDSIM studies, as it is most common for handling physics for radiobiology/medical applications. For this study in particular, secondary particles were modelled so that the complete nature of energy deposition by lost particles could be understood.

Aperture

In the LhARA lattice, we chose a global aperture size (aper1) within BDSIM to reduce losses. A parameter scan using a solenoid model with 10^4 protons — excluding collimators — determined that a global aperture radius of 3.65 cm produced the best results in minimising beam losses.

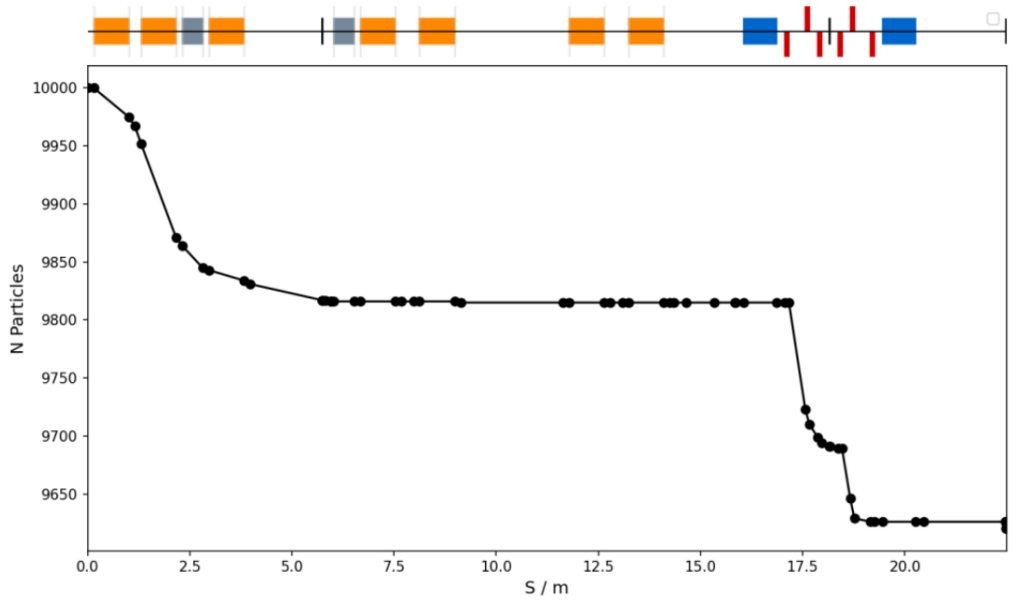

Figure 3.6: Beam loss (N particles) against distance along the lattice, s, with a 10^4 proton solenoid model and global aperture size of 3.65 cm (excluding collimators).

Following Figure 3.6 and subsequent energy deposition studies, the beam loss along the lattice can be understood below.

Examining Figure 3.7 two regions of significant particle loss and energy deposition can be identified, the first of which is at the very beginning of the beamline where the divergent beam from the source enters the energy capture lenses, and is focused for the first time. Approximately 1 in every 10^5 particles will make contact with the lattice for the first time in this region, and will be lost. After the initial hits and losses in this region, there is a small shower of additional

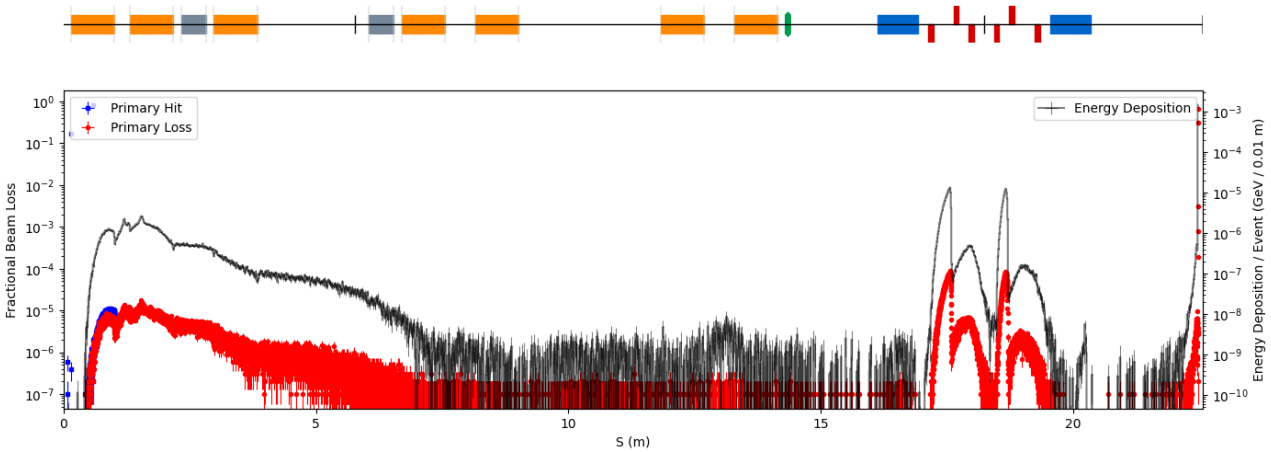

Figure 3.7: Fractional beam loss against position along the reference path, s, using the final 3.65 cm beam pipe aperture. Major energy deposition comes from the initial focusing of the beam by the first solenoids, and collisions with elements of the arc due to high dispersion in that region.

hits and losses, which declines as the beam enters the matching and energy selection section of the beamline (Gabor lenses 4-7). A series of continuous small losses occur (around one in a million particles) in this region.

The next region of significant loss and energy deposition comes as the beam enters the vertical matching arc, where high dispersion induced by the dipoles causes off-momentum particles travelling into the first defocusing quadrupole to take trajectories causing collisions with the lattice; this effect is seen again as the beam enters the second 'cell' of quadrupoles in the arc, creating large peaks in both deposition and loss (approx. 1 in 10^4 particles), followed by a secondary peak as a result of secondary particle showers from proton-lattice collisions.

Collimators

After simulating the lattice with the chosen aperture, the next step was to introduce the two collimators previously modelled as drifts. The first collimator is positioned at the focus of the third Gabor lens, and provides energy cleaning where the beam is at its smallest. It is modelled as a circular aperture with 1.8 mm radius, which is approximately the 2σ beam radius at the focus. The off-energy particles towards the edge of the phase-space are cleaned away, reducing the energy spread of the the distribution entering the first RF cavity through phase-space

manipulation.

The second collimator serves the purpose of momentum cleaning, and so is placed at the midpoint of the vertical matching arc. The high dispersion in the arc facilitates the removal of off-momentum particles by the collimator in the dispersive y-direction. As such, this collimator is modelled with an elliptical aperture, staying well clear of the beam in the x-axis (2.0 cm) but remaining around a 2σ radius in the y-axis (1.2 cm) to maximise momentum cleaning but reduce unnecessary losses in the non-dispersive axis.

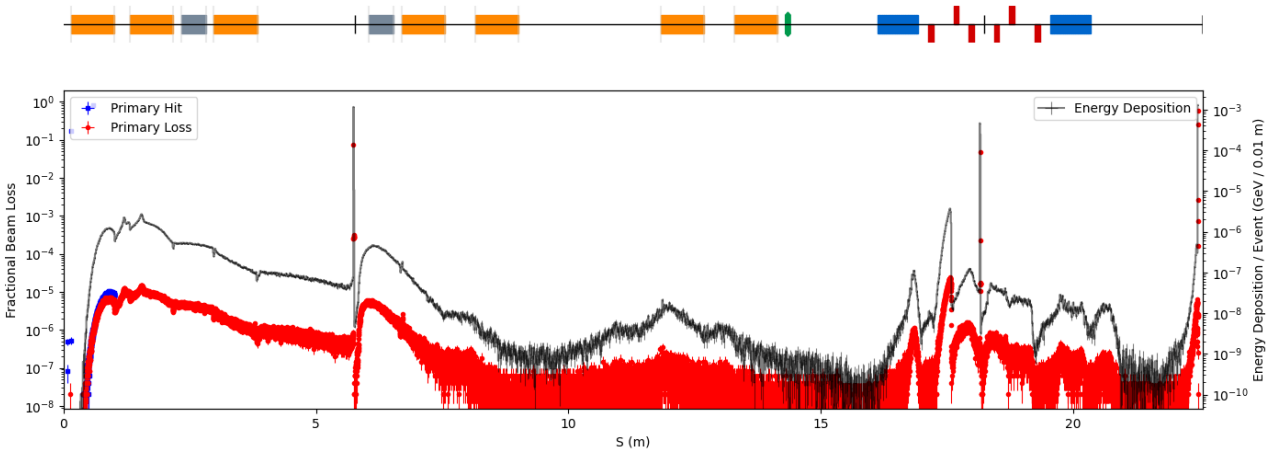

Figure 3.8: Fractional beam loss against position along the reference path, s, using the final 3.65 cm beam pipe aperture and including collimators.

Initially, Figure 3.8 shows similar losses at the first focus as Figure 3.7, however as the beam reaches the first collimator there is a prominent peak in loss and energy deposition at the collimator itself. This is followed by a smaller secondary peak leading into a much more sparsely populated region of continuous small losses (now around one in every ten million particles) with a small rise and fall again as the beam passes through lenses 6 and 7.

As the beam then enters the vertical matching arc, a small initial peak of loss and deposition is present at the dipole, more visible now that the previous region has less noise. This small peak is created by particles that enter the dipole at too shallow an angle, or are off energy/momentum enough to be bent into an unwanted trajectory and strike the dipole itself.

The first quadrupole peak observed in the aperture loss study is still present in the arc, however the second peak is entirely removed and replaced by a narrower but taller peak at the collimator.

The off-momentum particles that collide with the beam pipe as a result of the second quadrupole cell within the arc have been cleaned away by the collimator.

3.4 Octupole

3.4.1 Objectives

An objective of the LhARA study was to demonstrate the feasibility of uniform beams at the end station for controlled *in-vitro* experiments. It was demonstrated theoretically in 1983 [28] that the third-order focusing fields of octupoles were generally sufficient to transform high-energy ion beams from a Gaussian to a uniform distribution. Currently this is explored in simulation using the BDSIM tracking code by inserting an octupole of strength $k_3 = 30\,000\,\mathrm{m}^{-4}$ into the stage 1 beamline shortly before the arc, and firing a test bunch through the lattice to measure the uniformity post-arc at the end station.

A secondary benefit from this work is the applicability of the methodology developed for studying beam uniformity in directly comparing two different bunch profiles for shot-to-shot stability studies. It is particularly helpful to have a quantitative measure of this behaviour, combined with a larger-picture qualitative view, a method for which is described in Section 3.4.2.

3.4.2 How to Describe Uniformity

A given distribution has many moments associated with it, characterising aspects such as peak prominence, tailedness, and so on. This is typically captured by a moment generating function (MGF) with terms:

$$M_n = \frac{\mathbb{E}\left[(\mathbf{X} - \mu)^n\right]}{\sigma^n} \tag{3.1}$$

Kurtosis is the fourth moment ($n = 4$) of a distribution and describes how flat (uniform) a

data distribution is, giving a useful as a measure of bunch uniformity. Additionally, it was also realised that directly comparing quantiles gave an excellent qualitative view of uniformity without losing information.

3.4.3 Quantile Plots

A comprehensive statistical tool exists which can capture these relationships simultaneously, and it is known as a quantile-quantile (QQ) plot. A set of n quantiles on a set may be defined as the collection of lines that generates $n + 1$ equally sized partitions. Repeating this process with 99 quantiles gives you a set's percentiles. A QQ plot is thus the two-dimensional plot of one set of quantiles versus another as seen in Figure 3.9. A straight line with a fixed gradient is observed when two distributions are identical, or a scaled version of each other. When comparing a sampled distribution to a theoretical Gaussian, a normal probability plot is obtained, and it provides an excellent overview of the beam's likeness.

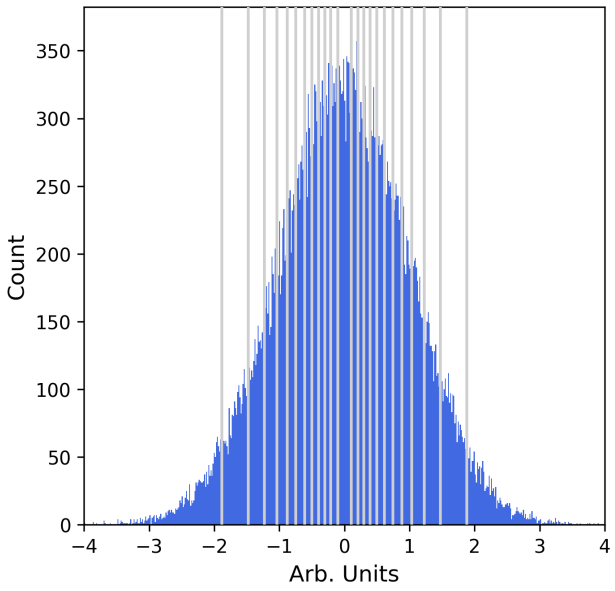

(a) Quantiles found by solving the cumulative density function over an interval.

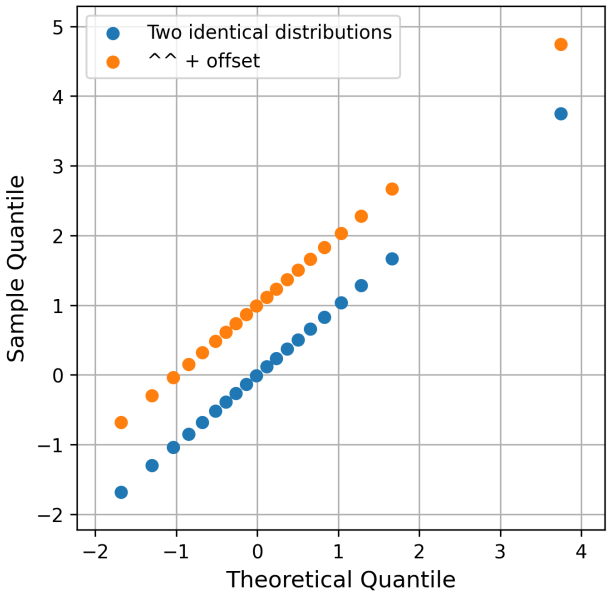

(b) Comparison of two identical distributions (blue), and two with a constant offset (orange).

Figure 3.9: (a) Sampled Gaussian partitioned by quantiles. (b) QQ plot generated from plotting it against itself, and another identical distribution with an offset.

This method was extended to study the uniformity of the proton bunch in a qualitative fashion, before passing through the octupole, and after exiting the bending arc. Combining plots with

beam images like Figure 3.10 reveals a richer structure than might be seen otherwise. In summary, kurtosis is a direct measurement of the smoothness of a distribution, while the QQ plot gives qualitative information on the bunch that could be missed otherwise.

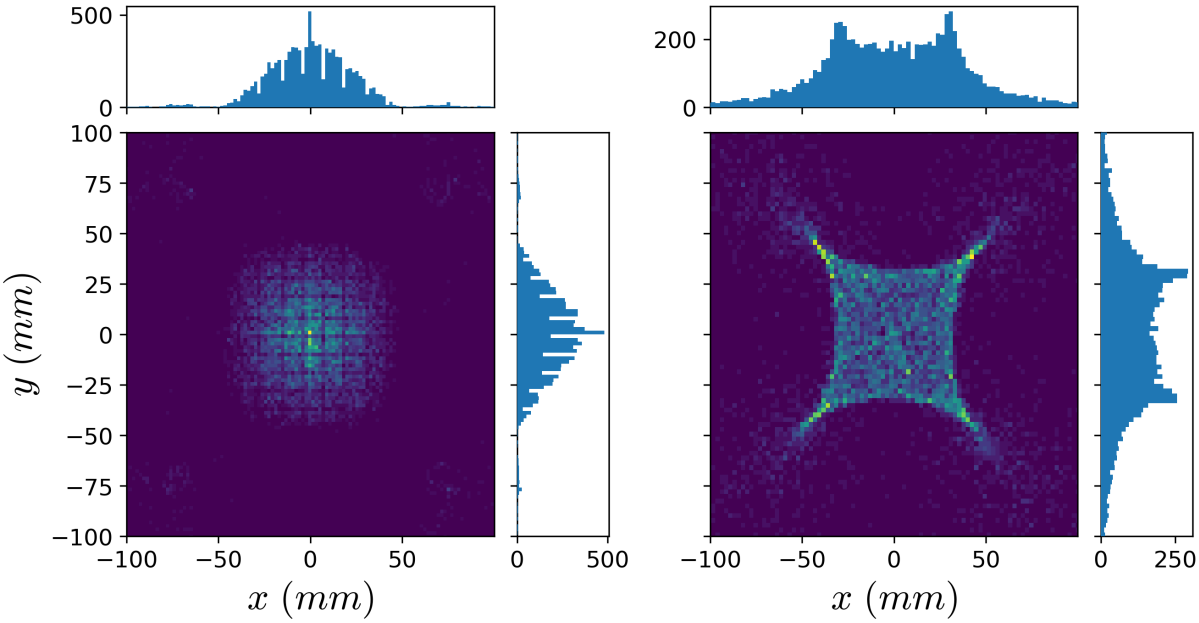

Figure 3.10: (Left) Snapshot of the bunch immediately before entering the octupole with a concentrated central bulge. (Right) Snapshot of the proton bunch exiting the bending arc. A 23° counter-clockwise rotation has been applied to the phase-space plot for easy calculation of kurtosis within the central region ($r = \pm 30\,\text{mm}$).

There is a noticeable heavy-tailedness to the phase-space distribution of the bunch post-arc, and the marginals highlight this. While some or all of each tail may survive to the end station depending on the final settled-upon radius of the vacuum pipe and octupole strength, work is needed to see if more particles can be transported to the centre of the bunch from the tails.

Figure 3.11 does a good job of indicating the improvement in uniformity (pink) achieved with the octupole from just a few simulation tests, despite the worrisome outlook from the snapshot alone. Uniformity can be inferred from a QQ plot by how well a distribution's quantiles match up with the uniform straight line (green). The Pearson Product Moment Correlation Coefficient (PPMCC) measures the linearity of a dataset. It takes on the values of ±1 when data is fully linear, 0 when fully uncorrelated, and expressed as a measure of the covariance

$$r_p = \frac{\sum(x_i - \overline{x})(y_i - \overline{y})}{\sqrt{\sum (x_i - \overline{x})^2 \sum (y_i - \overline{y})^2}} \tag{3.2}$$

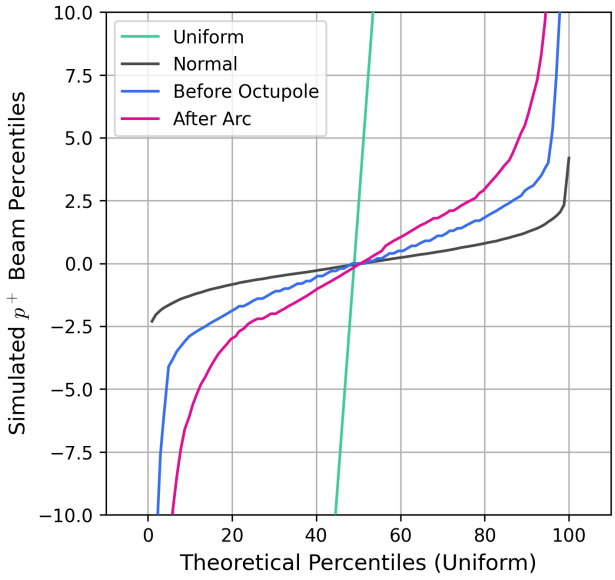

Figure 3.11: QQ plot of uniform (green), standard normal (grey), pre-octupole (blue) and post-arc (pink) distributions. Gaussian is characteristic of an unperturbed bunch, in contrast to an ideal uniform bunch. Each curve is formed from 100 data points.

PPMCC was used in these studies to quantify the beam uniformity on a QQ plot. When the quantiles of a sample distribution are plotted against a theoretical uniform distribution's quantiles, a straight line indicates uniformity. It was found that $p_r = 0.45$ before the octupole and $p_r = 0.60$ after exiting the arc, both rounded to 2 significant figures. This quantitatively shows an improvement in the beam uniformity which was an objective of the test. Equally important is the applicability of this tool in comparing any two bunch distributions, which could be very beneficial for shot-to-shot stability studies. Ultimately, the QQ plot will be a useful tool in providing a qualitative picture of the beam, while QQ PPMCCs combined with distribution moment measurements like kurtosis give a complimentary quantitative picture. Overall, the octupole successfully improved uniformity, with mostly fine-tuning left to do.

3.5 Gabor Lens Comparison

As previously mentioned, the lattice studies were carried out using solenoid elements as an approximation of the proposed Gabor lenses. As a sanity check, Gabor lenses were modelled in BDSIM using field maps generated using 2.1, i.e. for a cylindrical distribution of electrons. For

simplicity the resulting field map is attached to a drift space element in BDSIM, with all fringe fields neglected due to their minimal contribution. The resulting focusing fields are shown in 3.12.

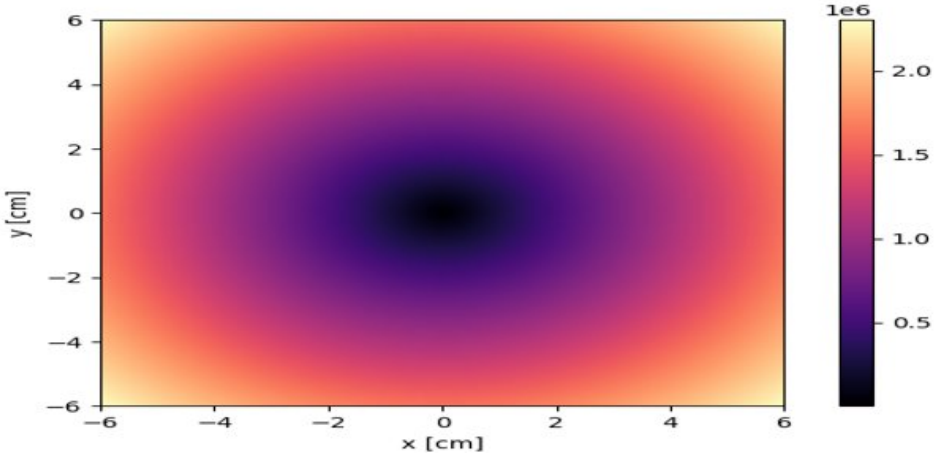

Figure 3.12: Gabor lens electric field map.

When initially imported, there were discrepancies between the optics resulting from Gabor lens and solenoid simulations. Field strengths were fine tuned using the COBYLA algorithm [29], which is a numerical optimisation routine used for objective function minimisation. Here, the objective function is the mean-squared error between the two sets of optics with the focusing strengths as function inputs. Optics results are shown in 3.13, showing an almost exact match.

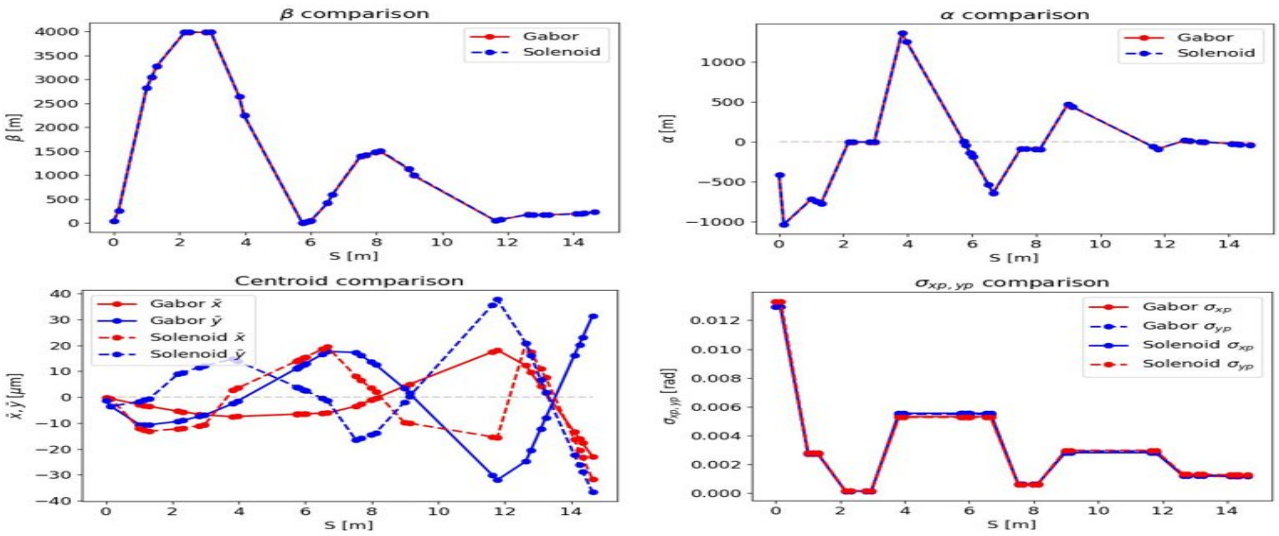

Figure 3.13: Optics summary comparison between Gabor lenses and solenoid magnets.

The field strengths required to achieve the desired optics varied by <1% of the reference strengths given by the solenoids in 3.1, indicating the initial Gabor lens model works as expected

with only slight variation in strengths.

3.6 End Station and Dose Calculation

While observations of the position and prominence of particle losses across the lattice are important for design decisions, a deposition peak that is very important to quantify for LhARA is the final peak at the end station. The total end station energy deposition translates to a dose delivered by LhARA for radiobiological studies; given these studies are the primary purpose LhARA, it is important for it to be measured accurately in simulation, requiring an accurate model for the end station.

The model, seen below in Figure 3.14, was constructed (top to bottom) of; a 75 µm Mylar vacuum window, 250 µm of polystyrene representing a scintillation fibre, a 5 mm air gap, 1.3 mm of polystyrene, and finally 2.4 mm of water to represent a container holding the scoring volume. This design intends to score dose similarly to a Markus ion chamber, and so the actual scoring volume is a 2.65 mm radius cylinder positioned at the centre of the water volume.

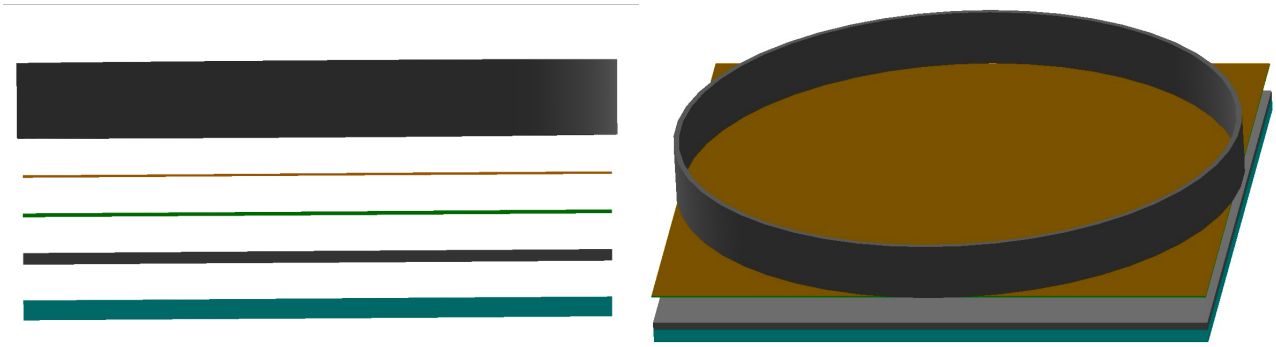

Figure 3.14: The end station model in the BDSIM visualiser, exploded (left) and as modelled (right). The model was constructed to provide similar dose scoring to a Markus ion chamber, recording the dose in a 2.65 mm cylindrical scorer in the final water layer.

BDSIM will record the dose deposited per proton as a histogram, which can be extracted and summed through pybdsim and scaled by a factor of 10^{10} to represent the expected 10^9 particles per shot, and the 10 Hz repetition rate of the laser. The result is the dose rate ($\mathrm{Gy\,s^{-1}}$) provided by the stage 1 lattice, and is a figure of great interest in quantifying the achievements

of LhARA with respect to both the provision of flexible beams for radiobiological studies, and the effectiveness of the emerging technologies being used in its construction.

Before implementing this model in the full stage 1 lattice model, it was important to verify that it was a close enough representation of the LhARA stage 1 end station to provide accurate dose rate information. A simulation was run for the isolated model with a 1 cm (2σ) beam of $10,000$ protons directly into the model, yielding a calculated dose rate of $122.63 \pm 1.42\,\mathrm{Gy\,s^{-1}}$. This is close to LhARA's theoretical dose rate for similar beam conditions of approximately $120\,\mathrm{Gy\,s^{-1}}$ [24].

Dose rate calculations were made for each stage of the loss and deposition study to compare the effects of the introduced elements on the final end station dose. Firstly, as a reference dose, the stage 1 lattice in its initial state (the reference lattice) with a wide 10 cm aperture was calculated, followed by the implementation of the chosen 3.65 cm aperture and then simulations of the introduction of first and second collimator respectively.

All calculations were made for simulations of a 3 cm beam of ten million protons using the 'g4QGSP_BIC_EMZ' Geant4 physics list, modelling secondary particles; the resulting dose rates can be found in Table 3.3.

	Dose Rate ($\mathrm{Gy\,s^{-1}}$)	Change w.r.t Reference
Reference	17.42 ± 0.01	0.0
3.65 cm Aperture	17.15 ± 0.02	-0.27
w/ Collimator 1	16.73 ± 0.03	-0.69
w/ Collimator 2	15.12 ± 0.01	-2.3

Table 3.3: Table of calculated dose rates from BDSIM simulations with the end station model. The changes in dose rate are small when aperture and first collimator are introduced, followed by a more significant change from the second collimator.

The significant impact of the second, momentum cleaning, collimator on the dose rate not only validates the design choices made for the collimators but also continues to validate the motivation for LhARA's aim for smaller beam sizes. Smaller beams will experience less loss in that second collimator and therefore correlate to a higher dose at the end station.

3.7 Further Work

Beyond the lattice studies undertaken and presented here, there is both room for improvement of the modelling of the lattice elements within BDSIM, and a need to further combine the studies into a fully optimised lattice both in terms of dose delivered and overall lattice optics. The next step of simulation is to implement field maps from the magnet designs presented in Chapter 5 and model these components with the exact field strengths and geometries of the design. This work can be followed by another stage of loss studies and dose rate calculations, however these studies themselves would have to be preceded by an optics optimisation taking into account components that were not present in the initial MAD-X model (octupole and RF cavities), to ensure that all the lattice constraints defined in the initial matching are still satisfied.

More work remains to be done on reducing the tailedness of the beam as it leaves the arc to improve the efficacy of the accelerator in its R&D. This could be explored by relocating the existing octupole, although initial studies have found this to have minimal impact on the beam profile out of the arc. There is a small amount of extra space available further down the beamline, but it is not currently known if this would or could accommodate a second lens for fine-tuning. Further comparisons of shot-to-shot performance for different bunch seeds with other tracking codes using QQ plots would be another interesting avenue to explore.

The ultimate aim of all these further steps is to produce a fully functional lattice that both conforms to the desired optics, and delivers a uniform beam (and therefore a uniform dose) to the end station with minimal losses along the beam line besides those intended to ensure both spatial and spectral uniformity.

Chapter 4

RF Cavity Design

4.1 Cavity Design Considerations

The energy and mass of the particle species are important to consider for the cavity design. For LhARA, the design scope currently includes both 15 MeV protons and 48 MeV carbon ions. Ideally the design of Stage 1 will accommodate the acceleration (bunching) of both species. This can prove difficult, since the spacing of multi-cell accelerating structures are suitable only for a particular particle energy and mass. As such, to accommodate the different ion species, the best option is to have independently operating single-cell cavities. It is also sensible to use the same design for each cavity, unless there is clear performance advantages not to.

The energy spectrum of particles from the laser-hybrid target is not currently well defined. For the purposes of this design phase we have made assumptions that the target produces particles with a nominal energy standard deviation of 2% for both protons and carbon ions.

The current Stage 1 design has allocated two RF bunching cavities to the lattice; one immediately after the first two Gabor lenses, and one before the quadrupole arc. Each has been allocated 50 cm of length in the lattice. Due to the significant energy spread of particles coming from the target, the bunch length quickly diverges. Placing the first cavity early is essential for keeping the particles in a single bunch. The second cavity is important for fine-tuning

the longitudinal phase-space to achieve appropriate bunching at the end. Therefore, introducing RF cavities at this stage is necessary to control beam length and manipulate longitudinal phase-space.

Here 'bunch length' refers to the average temporal size of particle bunches within the beam. This parameter represents the temporal duration over which particles within a single bunch are distributed, hence the bunch length indicates the temporal interval during which particles are in phase with the RF oscillation. Longitudinal phase-space pertains to the distribution of positions and velocities of particles along the longitudinal axis of the accelerator, namely in the direction of beam propagation. This enables understanding of how particles propagate and interact within the accelerator.

It is of utmost importance to control the bunch length and longitudinal phase-space as they are closely tied to the efficiency of the accelerator. Indeed, a more compact group helps maximise particle interaction with the cavity electric field, thereby increasing beam stability and reducing particle dispersion. One of the fundamental parameters of the RF cavity contributing to their regulation is the RF cavity voltage, as indicated in Equation (4.1). This voltage determines the amplitude of the oscillating electric field inside the cavity, which acts on particles, either accelerating or decelerating them depending on their phase relative to the RF oscillation. By altering the cavity voltage, it is possible to influence the potential energy of particles, thereby controlling their velocity and momentum.

$$V_{RF} = \frac{\Delta E^2 \pi h |\eta|}{2q\beta^2 E_s} \tag{4.1}$$

The choice of frequency for the cavities is also important for the bunching capability. Lower frequencies will give a longer time interval over which the particles will be longitudinally 'focused'. For practicality, the frequency chosen should be one common in the industry to allow for good price and availability for the surrounding RF infrastructure. As such, the frequencies of 201 MHz and 352 MHz were chosen to investigate. If further improved bunching is required or the energy spread from the target is larger than expected then lower frequencies may need

to be investigated also.

4.2 SuperFish Simulations

For the initial design of the cavity, the simulation package SuperFish was chosen to be used. This package is a compilation of cavity design frameworks, suitable for modelling and simulating most simple cavity designs. The framework chosen to be used for designing this bunching cavity is CCLFish, a package for modelling close-coupled LINACs. Although intended for multi-cell cavities this framework was also effective for designing a single-cell bunching cavity like the ones for LhARA.

The CCLFish framework gives eight free parameters that can be adjusted to meet the design specifications. A diagram depicting these parameters on the CCLFish cavity geometry is shown in Figure 4.1. The diameter is automatically adjusted to fit the cavity to the designated frequency. Two parameters have a dominating effect on the cavity's design – length and gap length.

After running the simulation of a cavity design, a results text file is generated. Along with the data describing the field on axis for the cavity, it contains many descriptors which are key for evaluating the effectiveness of this cavity design. The most important descriptors to our design are described in Table 4.1.

Parameter	Description
Shunt impedance	The impedance equivalent for the RF. A high value indicates higher transfer of power to the beam.
Transit time factor	Indicator of the field uniformity the particle sees as it transits the cavity.
Quality factor	Describes the resonance quality of the cavity by the half-life of the oscillation amplitude.

Table 4.1: Cavity performance indicators.

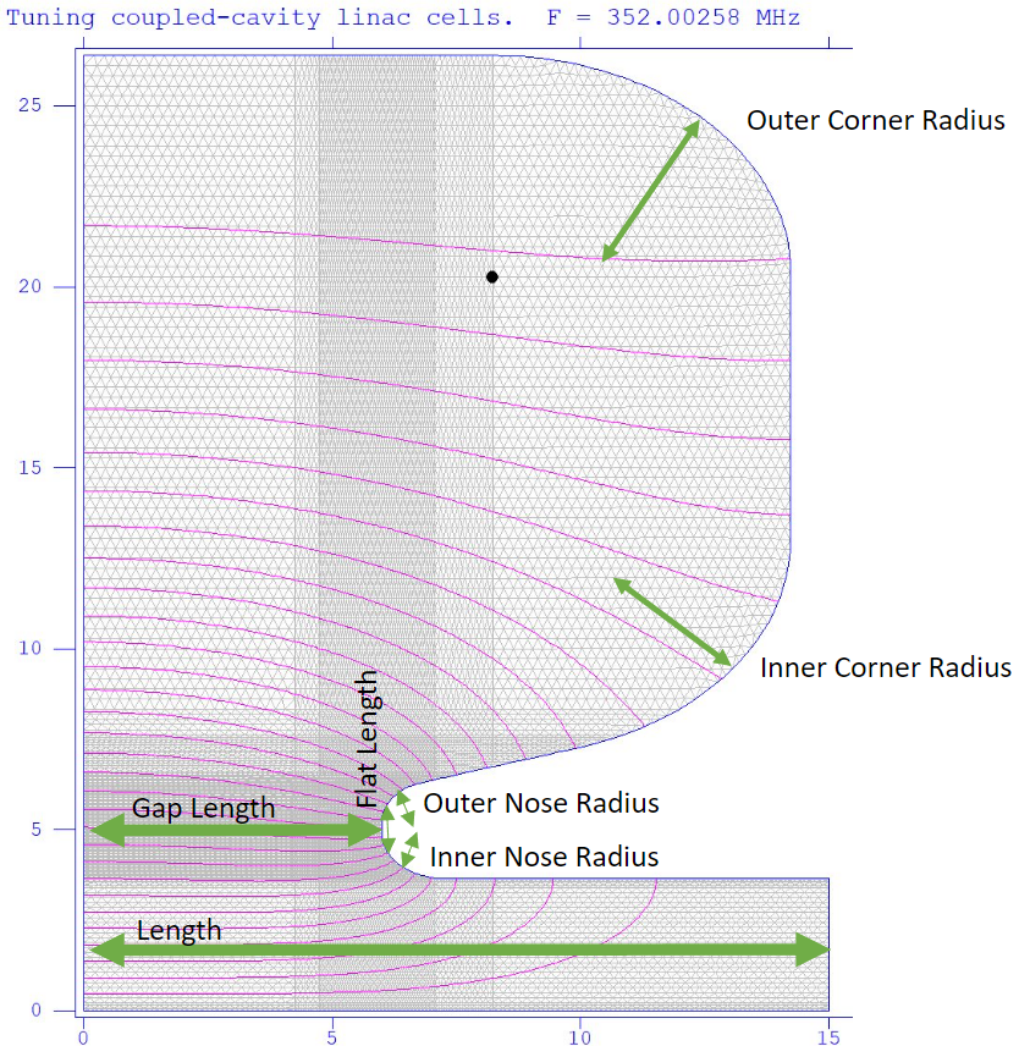

Figure 4.1: Key parameters defining the geometry of the CCLFish Framework.

4.3 Longitudinal phase-space simulation

Although the SuperFish simulation provides much useful information about the cavity's performance, fails to provide information about its overall performance in bunching the beam. To properly evaluate the design of such a cavity, we must simulate the accelerator's longitudinal phase-space. There are many simulation codes that can do this including BDSIM, the software package being used for the lattice design component of this report, however these packages are very computationally expensive and are excessively complex for this task. To iteratively test many cavity designs, a simple and fast alternative was necessary.

A custom simulation package was written in Python to measure the longitudinal phase-space and evaluate the cavity design. A set of N particles were created to represent the initial particles coming from the target, with a Gaussian distribution of energies. These were then propagated through the lattice of LhARA Stage 1, modelling all components except for the cavities as drift tubes since they are expected to have insignificant effect on the longitudinal phase-space. For a given cavity design, the field along axis was taken from the simulated CCLFish results file, and the particle was simulated to be accelerated by this field.

The time-varying nature of these fields was included in this simulation, and a cavity phase was calculated such that the on-momentum particle could pass through the cavity with zero net acceleration. All other particles received an acceleration from their delay relative to the on-momentum particle.

By simulating this set of particles as they pass through both cavities and the length of LhARA Stage 1, we can develop an understanding of how the longitudinal phase-space is transformed, and what bunch length and energy spread can be achieved at the low energy *in-vitro* end station. A visualisation of the longitudinal phase-space throughout the accelerator can be seen in Appendix I.

4.3.1 Bunch length and energy spread calculation

From the set of particles simulated, a standard approach for calculating the bunch length and energy spread would use the 1σ standard deviation. This is appropriate for Gaussian beam distributions, but as seen in the visualisation in Appendix I, the distribution can become highly non-Gaussian, with values from 'lost' particles dominating the standard deviation calculation to the point that the value is no longer representative of the distribution of the majority of particles.

An alternative to this could be using a percentile-based method. It was found that defining the bunch length and energy spread by the 10^{th} and 90^{th} percentiles gave a much better representation of the bunching, and was more successful as a metric for achieving the optimisations discussed in Section 4.4 of this report. These percentile-based spread values could be converted back into a 1σ standard deviation by assuming the distribution to once again be Gaussian, reducing the percentile value by a factor of 2.56.

For transparency and consistency in this report, the 1σ standard deviation is the method used to calculate the bunch length and energy spread throughout.

4.3.2 Simulation verification and accuracy

Since a custom simulation code was written to model the longitudinal phase-space of the particles throughout Stage 1, the results should be compared to other well-established packages to verify its accuracy. In a later stage of the simulation process an optimised cavity design was included in a simulation in BDSIM, as discussed in Section 4.7.

The longitudinal phase-space as simulated by BDSIM at the exit of the second RF cavity is shown in Figure 4.2. This was simulated using the optimised 352 MHz design. The longitudinal phase-space as simulated using the simple custom code is shown in Figure 4.3 for comparison.

The qualitative similarity between the two phase-spaces is clear between the two figures, but the quantitative measurements of bunch length and energy spread show some minor discrepancy.

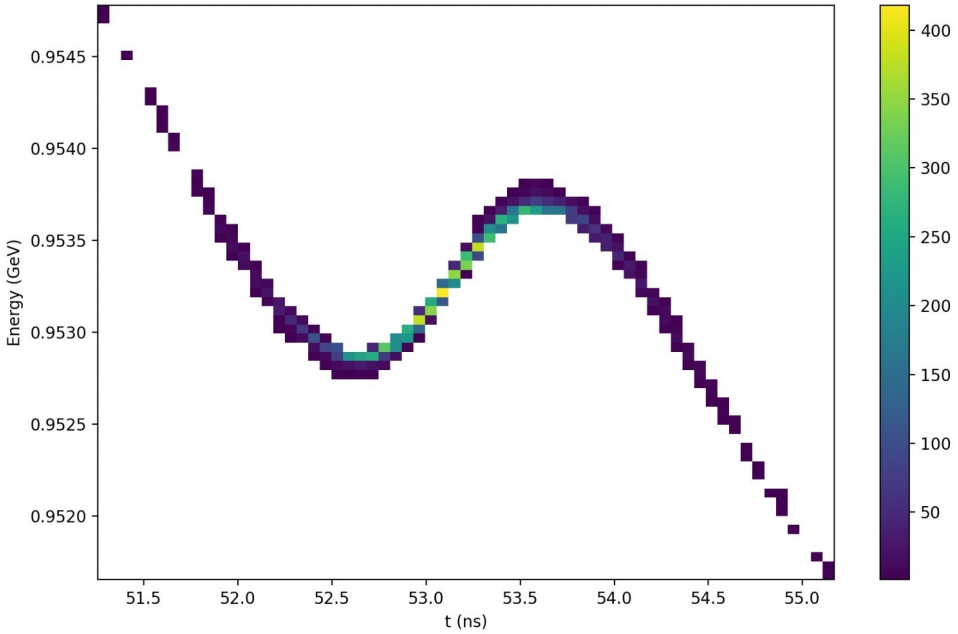

Figure 4.2: BDSIM simulation of the longitudinal phase-space after the second cavity.

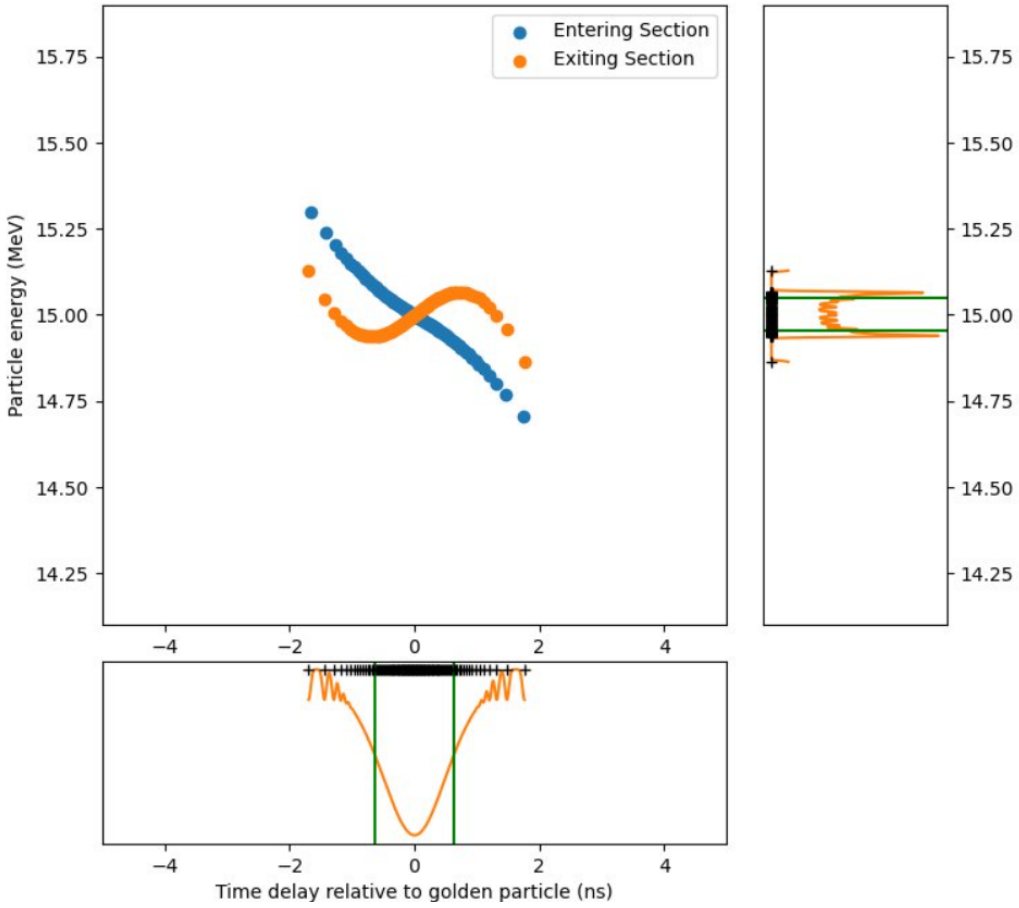

Figure 4.3: Custom simulation of the longitudinal phase-space after the second cavity (shown in orange).

The custom simulation code measured a bunch length of 1.26 ns and an energy spread of 0.05 MeV or 0.31%. The BDSIM simulation code measured a bunch length of 0.9 ns and an energy spread of 0.7%. These show qualitative agreement in that both show significant bunching compared to the particle distribution when no cavity is active (4 ns, 2%).

Key differences between the two simulation methods are the inclusion of transverse particle movement and coupling into the longitudinal phase-space from other lattice components in BDSIM, but at the cost of using a more discretised 3D model of the cavity which could affect its accuracy. In future, a 3D cavity model with a finer mesh should be used to minimise this effect.

4.4 Longitudinal phase-space optimisation

4.4.1 Optimisation heuristic

To evaluate and compare different cavity designs, it is important to define a heuristic function that clarifies the objective of the optimisation. For the purposes of LhARA stage 1, this will include the cavity's bunching capability (bunch length and energy spread), as well as the cavity's operational characteristics (shunt impedance, transit time, etc).

The nature of how this heuristic is assembled and the weightings given to each property will define what an optimal solution looks like. Throughout the optimisation process, the heuristic was iteratively tweaked and adjusted to better match the expectations of an 'optimal' cavity. The general form of heuristic function used is in Equation 4.2 below, with examples of values shown in Table 4.2. The value of a_0 is generally chosen to be the value at the optimal solution. The power is chosen such that higher values will penalise moving away from the optimal value at a faster rate. The weightings define the priorities by which the values are adjusted and minimised.

$$H = \sum_{n=1}^{N} \left(\frac{a_n}{a_{n0}} \right)^p \tag{4.2}$$

n	Name	Expected Value (a_{n0})	Power (p)
1	Bunch Length	0.4 ns	2
2	Energy Spread	0.06 MeV	2
3	Shunt Impedance	30 MΩ	−1

Table 4.2: Heuristic function example values.

4.4.2 In-simulation operational optimisation

For a given cavity geometry, there are some operational optimisations that should be performed to give a fair comparison of the designs. Firstly, in the SuperFish simulation it is necessary to specify the cavity field gradient used in the simulation. Different cavity geometries have different maximum field gradients, as indicated by the Kilpatrick factor. For the purposes of cavity design it is a fair assumption that cavities can be manufactured to sustain a maximum field giving a Kilpatrick factor of 1.5. To give a fair comparison of the cavities, the SuperFish simulation was repeated after normalising the fields to give a Kilpatrick factor of exactly 1.5.

Even after optimising the cavity simulation to represent its full field capability, it would be wrong to assume that both cavities will be operating at this maximum field gradient in order to achieve optimal bunching. As such, an initial optimisation must be performed to choose the field amplitudes of the two cavities that will give the best bunching.

The method used for this optimisation was two sets of grid searches - the first over the full combination-space of field amplitudes, and the second narrowing to give better resolution at the optimal point from the previous search. This ensures that the bunching values for the evaluation of a given cavity geometry are the best possible.

4.4.3 Cavity geometry optimisation

With a method for calculating results for a cavity and an evaluation heuristic function, the final component necessary to form an iterative optimisation loop is the optimisation method. There are many optimisation methods available, which are more suitable for different purposes depending on the nature of the parameter space being searched.

To optimise the cavity design many of these optimisation methods were tested in combination with different heuristics, and yielded mixed results. A formal study comparing the methods was out of scope for this project, and the final optimised design was reached through a mixture of the methods in Table 4.3, concluding with the Nelder-Mead algorithm.

A visualisation of the iterative optimisation process used is shown in Figure 4.4

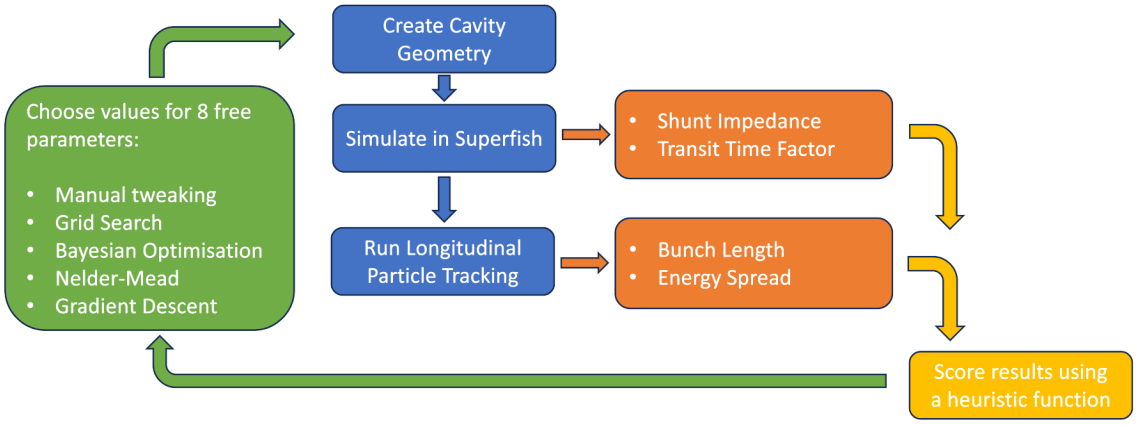

Figure 4.4: Flow chart describing the iterative workflow of the RF cavity optimisation methods used.

4.5 Final cavity geometry

It was chosen to investigate the bunching capabilities of cavities operating at both 201 MHz and 352 MHz. Both cavities had their geometry optimised based on the process stated in section 4.4.3. The results will now be compared.

Optimisation method	Description	Comparison
Manual adjustment	A human manually reviews each design and decides on the next set of adjustments to trial.	This method is very inefficient except for an experienced professional. It can lead to large sections of parameter space being unexplored.
Grid search	A set of evenly spaced points within the parameter space are decided upon. All points are evaluated then the minimum is taken.	This method is also quite inefficient since so many points are evaluated.
Gradient Descent	For a given starting point, other points are evaluated surrounding this point to calculate a gradient, then a step is taken in the direction of the minima.	This method can be quite efficient, but can get easily stuck in local minima.
Bayesian Optimisation	Each previously evaluated point is treated as a Gaussian process. The expected location of the minima is iteratively evaluated, eventually finding the global minima.	Most effective at finding the global minima, but can be slow.
Nelder-Mead Algorithm	This is a simplex method where points are iteratively evaluated to find a local minima.	This algorithm can be much faster than the other algorithms, but only converges to a local minima, and sometimes doesn't converge at all.

Table 4.3: Optimisation algorithms investigated.

4.5.1 201 MHz

Figure 4.5 shows the final 2D geometry of the 201 MHz cavity in SuperFish. The fundamental mode electric field lines are also shown. Figure 4.5 shows only the upper right corner of the cavity since, given the symmetry of the problem, only this region needs to be calculated. Cylindrical rotational symmetry is implied. In figure 4.5, the beam travels along the horizontal axis. The details of the geometry used are outlined in Table 4.4.

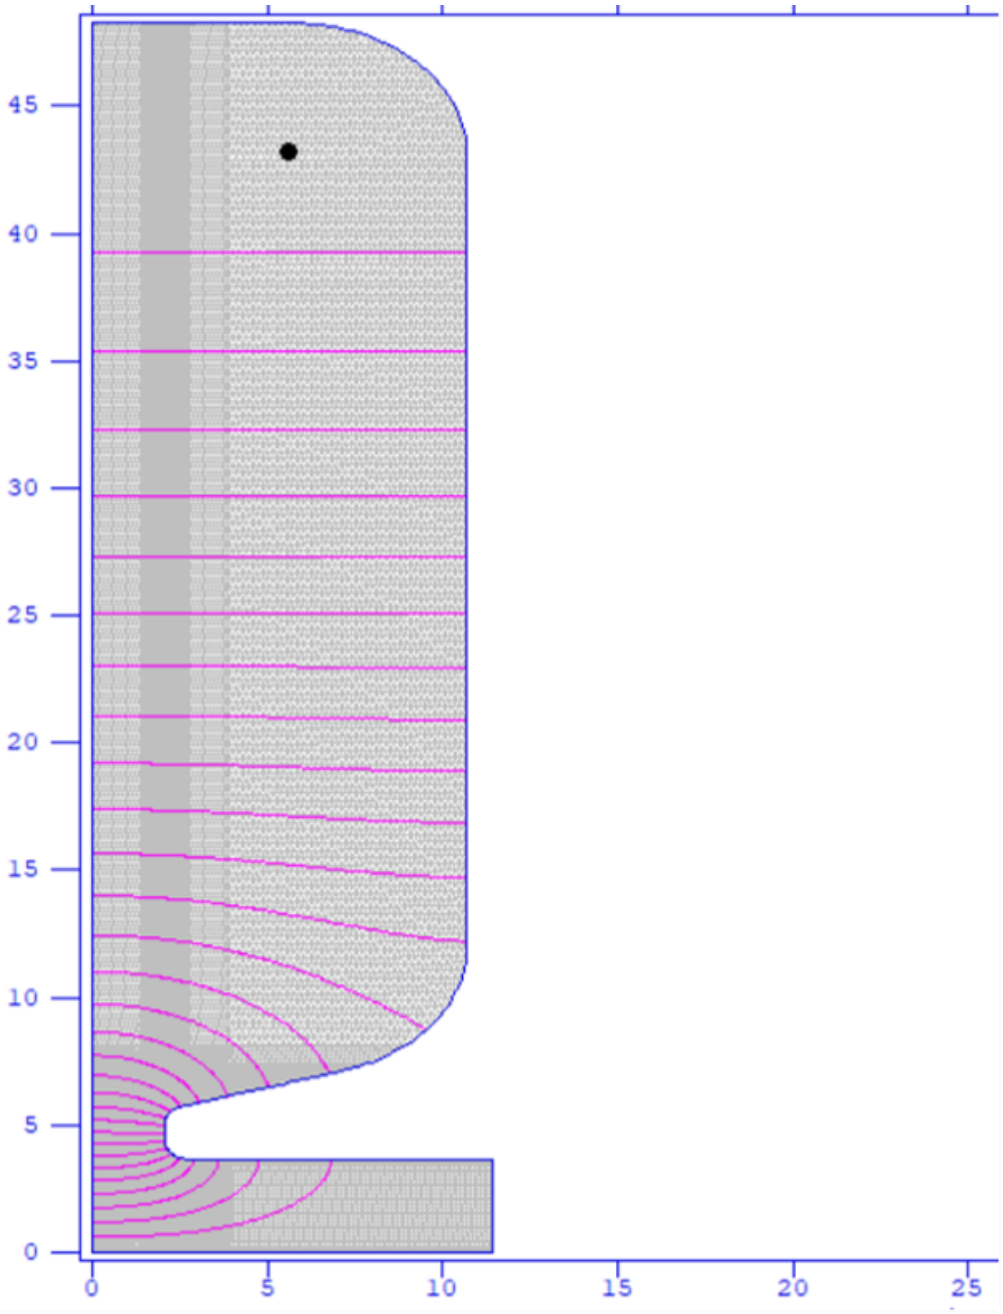

Figure 4.5: Final cavity geometry for the 201 MHz cavity in SuperFish. Axes are in cm.

Parameter	Value
Length [cm]	22.9
Gap length [cm]	4.1
Outer corner radius [cm]	5
Inner corner radius [cm]	5
Outer nose radius [cm]	0.7
Inner nose radius [cm]	0.7
Flat length [cm]	0.7
Cone angle [°]	17

Table 4.4: Geometry parameters for 2D SuperFish design of the final 201 MHz cavity.

Parameter	Value
Frequency [MHz]	201
Diameter [cm]	96.2
Shunt Impedance - Z [$\mathrm{M\Omega\,m^{-1}}$]	25.41
Transit time factor - T	0.33
ZTT [$\mathrm{M\Omega\,m^{-1}}$]	2.91
Maximum E-field [$\mathrm{MV\,m^{-1}}$] @ Kilpatrick = 1.5	22.15
Maximum E-field on axis [$\mathrm{MV\,m^{-1}}$] @ Kilpatrick = 1.5	8.08

Table 4.5: Results from 2D SuperFish simulation of the final 201 MHz cavity.

The maximum field occurs near the cutout for the nose. The Nelder-Mead optimiser generally tried to increase the curvature of this area to reduce the maximum overall field. However, the heuristic function used also encourages a higher transit time factor and a larger field on axis. These two objectives conflict so the optimiser should arrive at a balance between the two as long as the weights of the objective function are set appropriately.

From the longitudinal phase-space simulation, it was found that this cavity design was able to achieve a bunch length of 0.32 ns, with energy spread of 0.8% for protons. For carbon ions, this design was able to achieve a bunch length of 16.2 ns, with energy spread of 3.9%.

4.5.2 352 MHz

Following a similar process we arrived at a final geometry for a 352 MHz cavity. This cavity has a smaller diameter owing to the higher frequency. However, the higher frequency also means that smaller proportion of the bunch length is longitudinally captured by the cavity.

Nevertheless, it was thought that despite this, a full final design of a the higher frequency cavity was merited enable us to compare the two frequency choices.

Figure 4.6 shows the cavity geometry and field lines for the 352 MHz cavity. Similarly to the 201 MHz, the Nelder-Mead optimiser reduced the maximum field in the cavity by increasing the radius of curvature of the nose section. The details of the geometry used are outlined in Table 4.6.

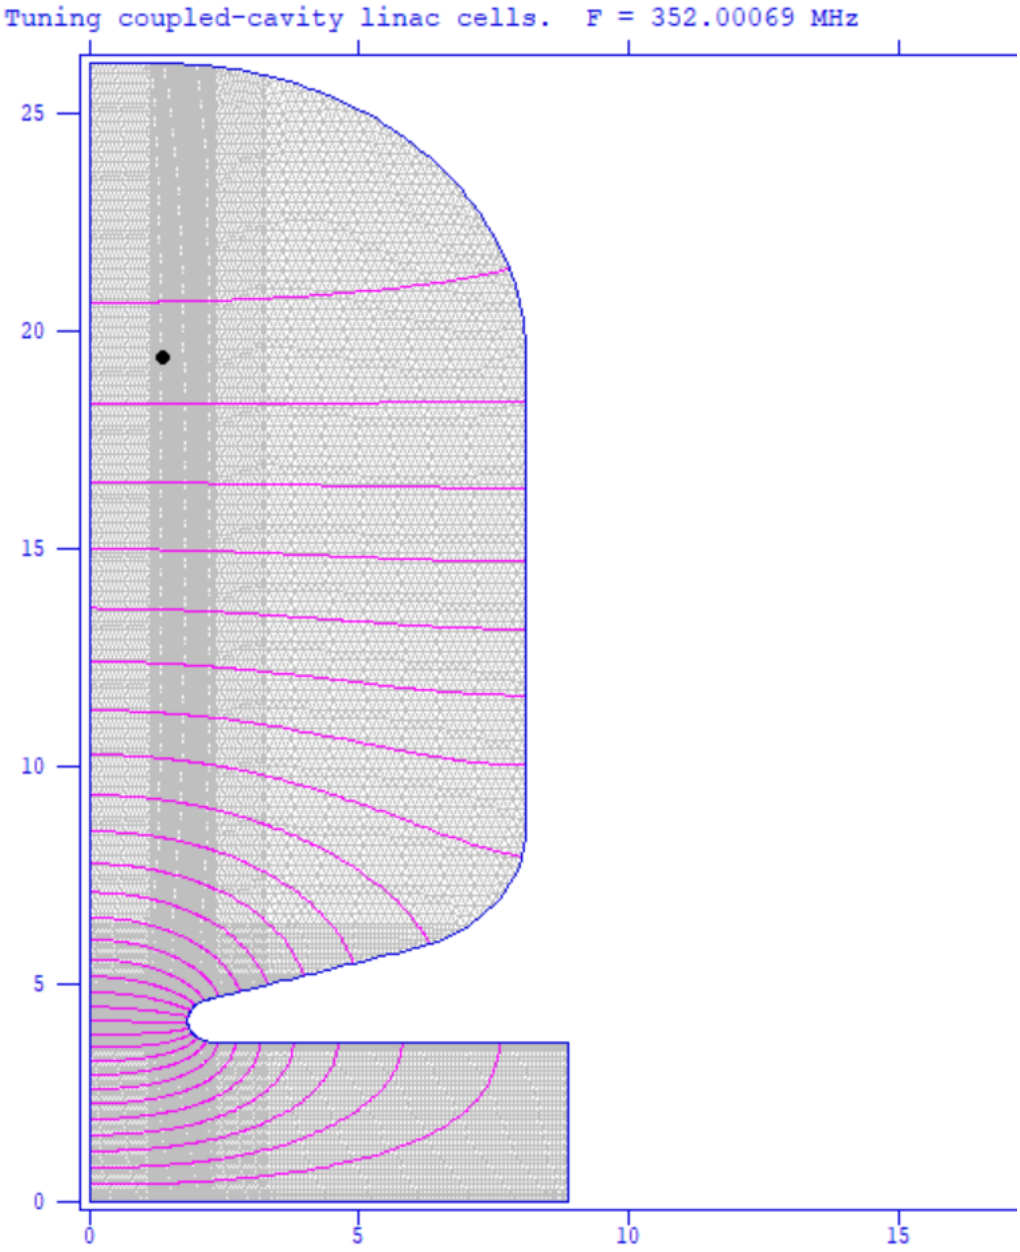

Figure 4.6: Final cavity geometry for the 352 MHz cavity in SuperFish. Axes are in cm.

From the longitudinal phase-space simulation, it was found that this cavity design was able to achieve a bunch length of 1.2 ns, with energy spread of 1.1% for protons. For carbon ions, this

Parameter	Value
Length [cm]	22.3
Gap length [cm]	6
Outer corner radius [cm]	9
Inner corner radius [cm]	9
Outer nose radius [cm]	1
Inner nose radius [cm]	1
Flat length [cm]	0.7
Cone angle [°]	17

Table 4.6: Geometry parameters for 2D SuperFish design of the final 352 MHz cavity.

design was able to achieve a bunch length of 19 ns, with energy spread of 4.7%.

Parameter	Value
Frequency [MHz]	352
Diameter [cm]	51.9
Shunt Impedance - Z [M$\Omega\,$m^{-1}]	32.6
Transit time - T	0.72
ZTT [M$\Omega\,$m^{-1}]	16.8
Maximum E-field [MV m^{-1}] @ Kilpatrick = 1.5	27.33
Maximum E-field on axis [MV m^{-1}] at Kilpatrick = 1.5	8.1

Table 4.7: Results from 2D SuperFish simulation of the final 352 MHz cavity.

4.6 3D cavity design

After settling on a design in 2D we progressed to a full 3D design using CST microwave studio. A 3D design allows us to cross check the SuperFish simulation, obtain a 3D field map of the cavity and validate our 2D optimisation. The results of the 3D CST simulation using the eigenmode solver for the 201 MHz cavity are shown in Figure 4.9, and the equivalent results for the 352 MHz cavity are shown in Figure 4.10.

In both cavity designs the maximum electric field magnitude, and the mostly likely point for a breakdown to occur, is at the curvature at the cavity nose. Although this feature increases the overall maximum field it also increases the all important field on axis. The optimisation process should have minimised this problem to a reasonable degree resulting in an effective cavity.

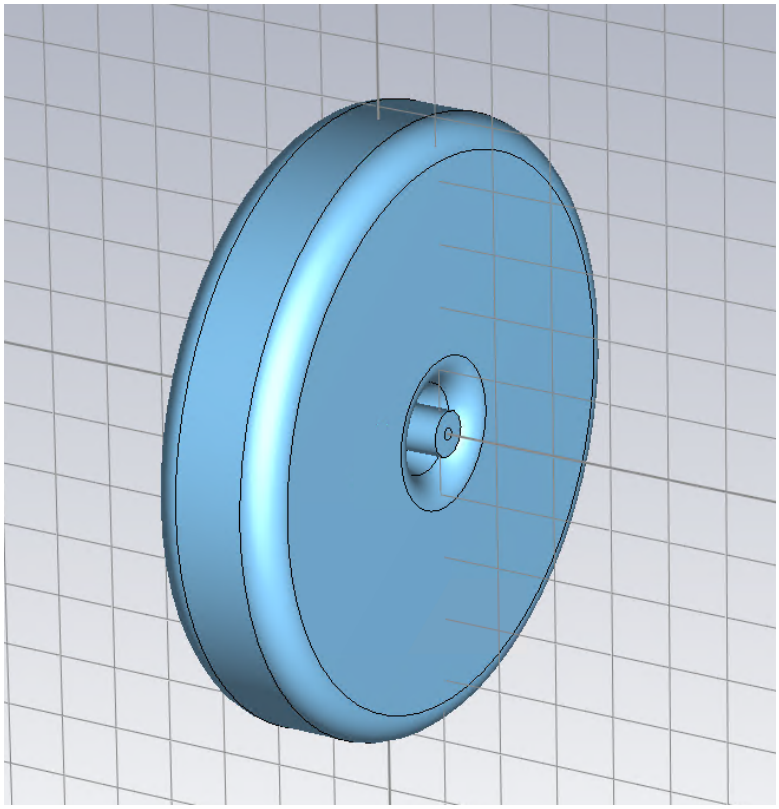

Figure 4.7: The outside of the 201 MHz cavity in CST.

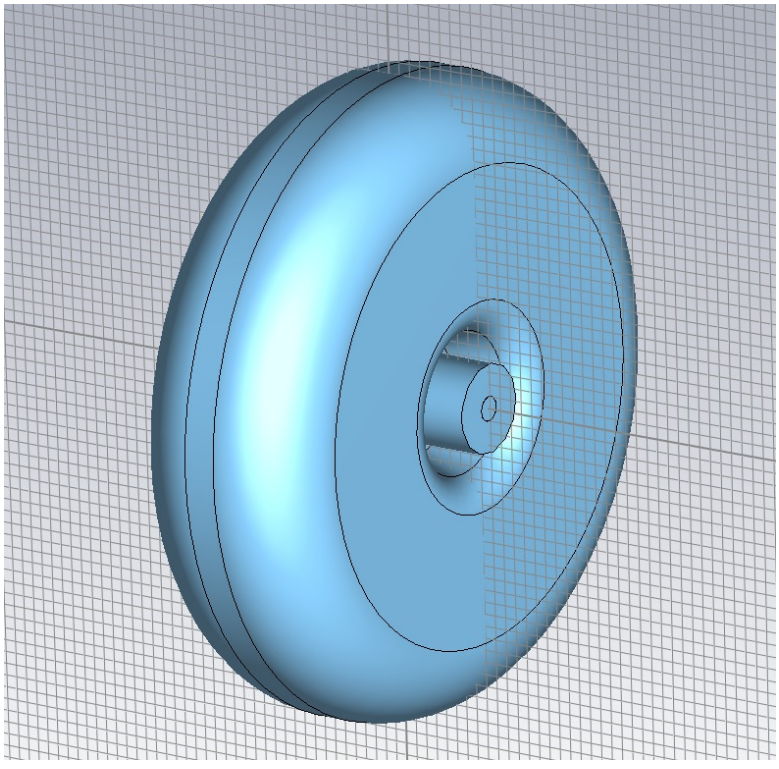

Figure 4.8: The outside of the 352 MHz cavity in CST.

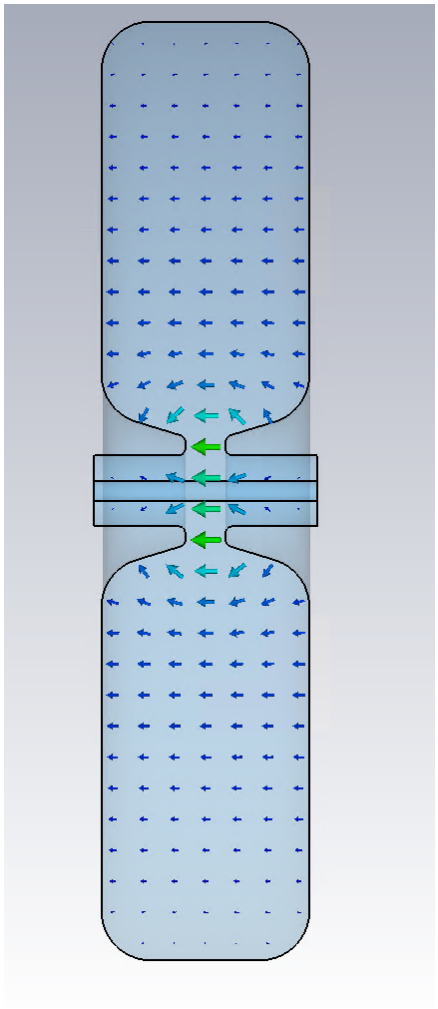

(a) Electric field line arrow plot.

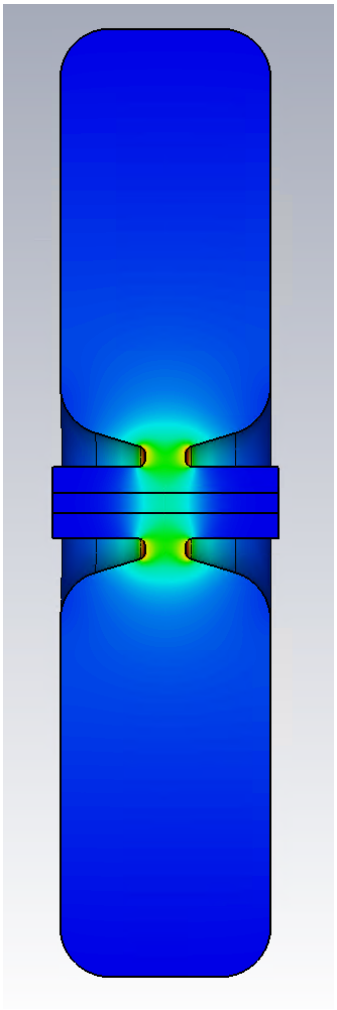

(b) Electric field magnitude contour plot.

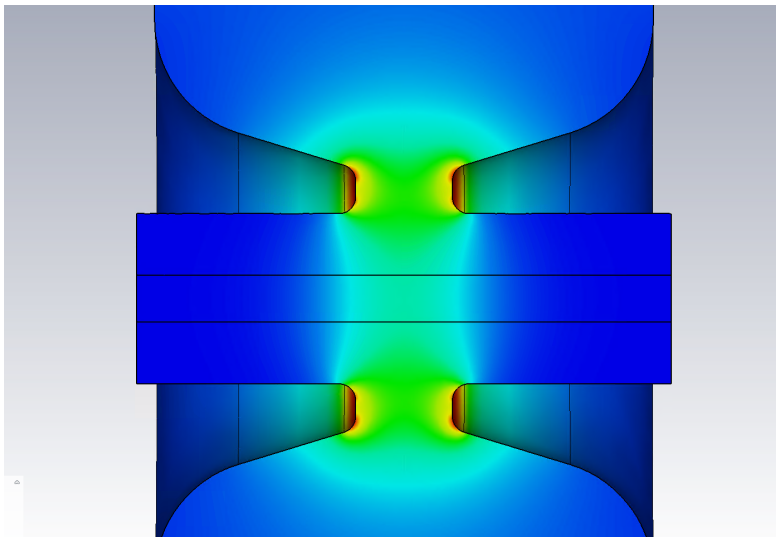

(c) Zoom into Figure 4.9b.

Figure 4.9: Cross section of the 201 MHz cavity. The absolute size of the field is somewhat arbitrary, as these plots are intended to show the relative strengths of the field within the cavity.

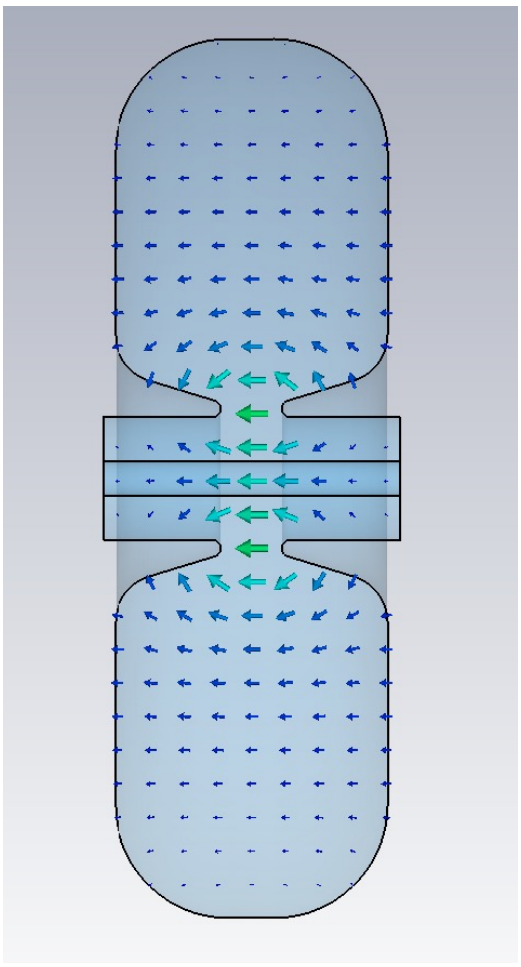

(a) Electric field line arrow plot.

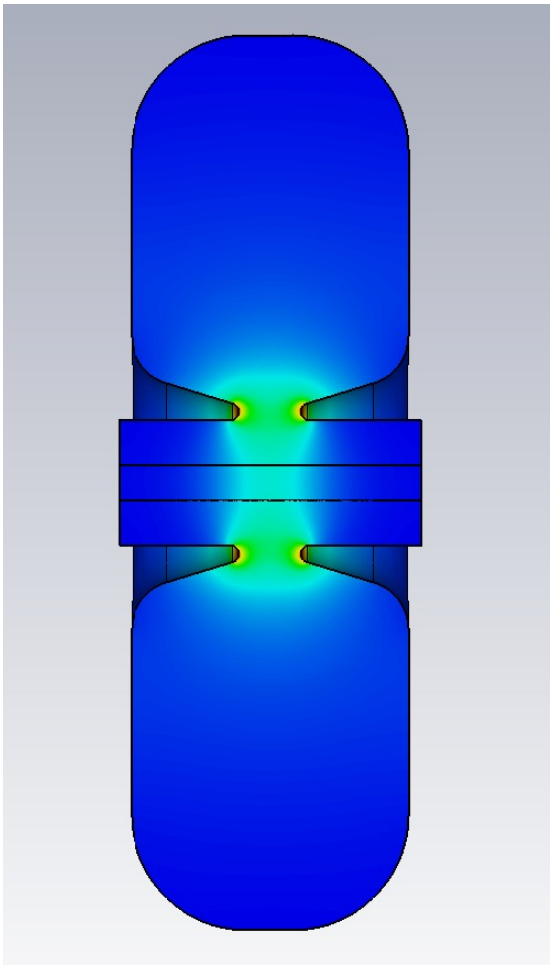

(b) Electric field magnitude contour plot.

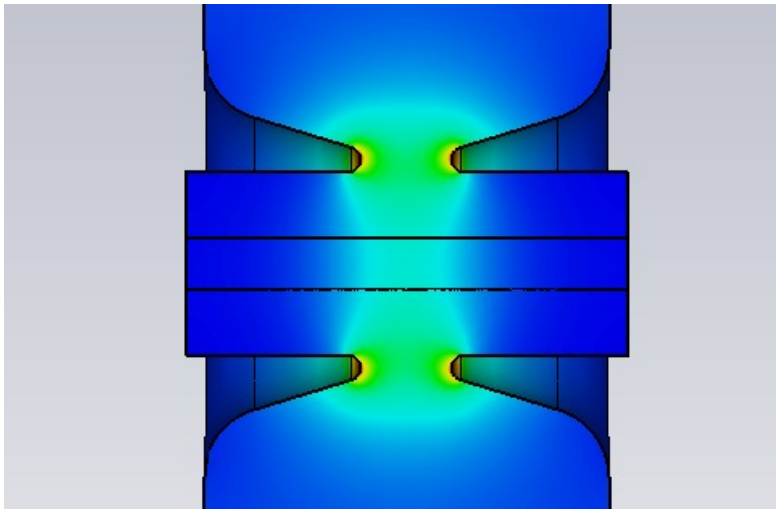

(c) Zoom into Figure 4.10b.

Figure 4.10: Cross section of the 352 MHz cavity. The absolute size of the field is somewhat arbitrary, these plots are intended to show the relative strengths of the field within the cavity.

4.6.1 Comparison with 2D

With the 2D cavity translated into a 3D CST model we can now obtain a 3D cavity field map and compare the CST result with that from SuperFish. Since the overall field normalisation used in the electric field magnitude calculation in somewhat arbitrary, the field magnitudes have been rescaled such that the stored energy in the cavity is 1 J, the CST standard.

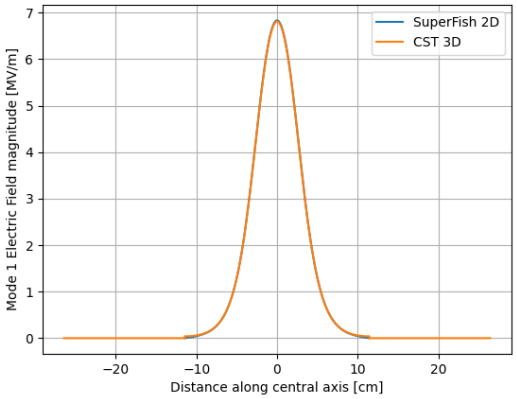

(a) Field on axis profile.

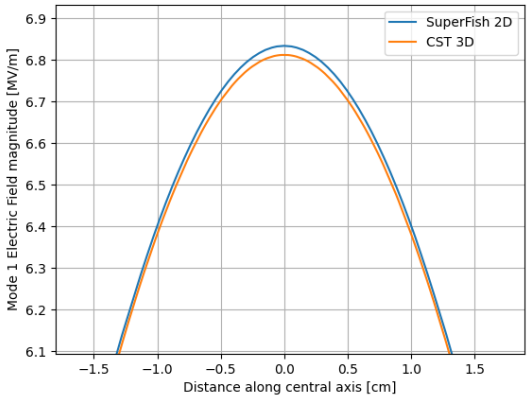

(b) Field on axis in the middle of the cavity.

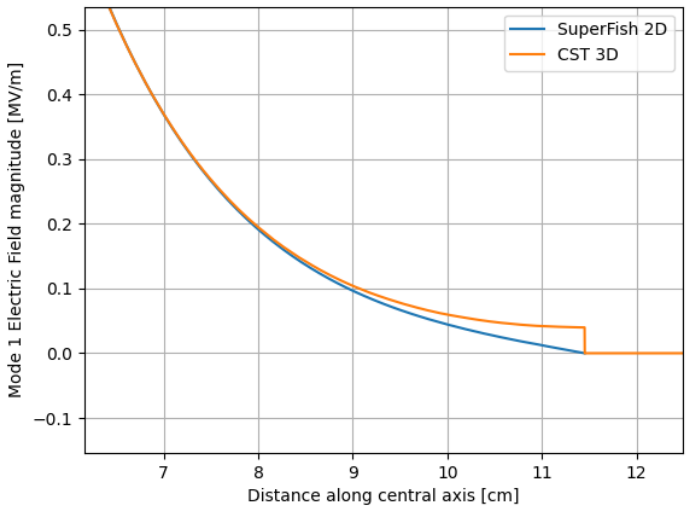

(c) Field on axis at the cavity boundary.

Figure 4.11: A comparison of the fundamental mode electric field on axis in CST and SuperFish for the 201 MHz cavity.

The field on axis for the 201 MHz and 352 MHz cavities is shown in Figures 4.11 and 4.12 respectively. The fields on axis for the two simulations agree well, particularly in the body of the cavity. As a result, the two simulations produce transit time factors within a few percent of

each other. To produce these plots, a finer mesh in CST around the beam axis had to be used. CST has no notion of the beam axis as a region of interest but this can be manually adjusted by defining a region of higher mesh density. The black lines in Figure 4.9c show the finer mesh region.

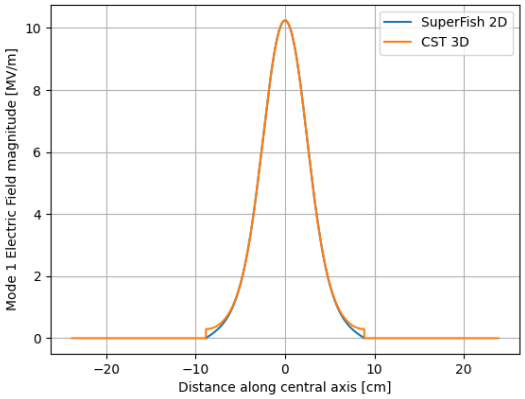

(a) Field on axis profile.

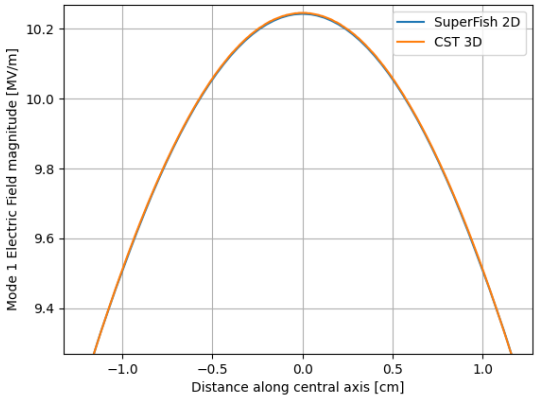

(b) Field on axis in the middle of the cavity.

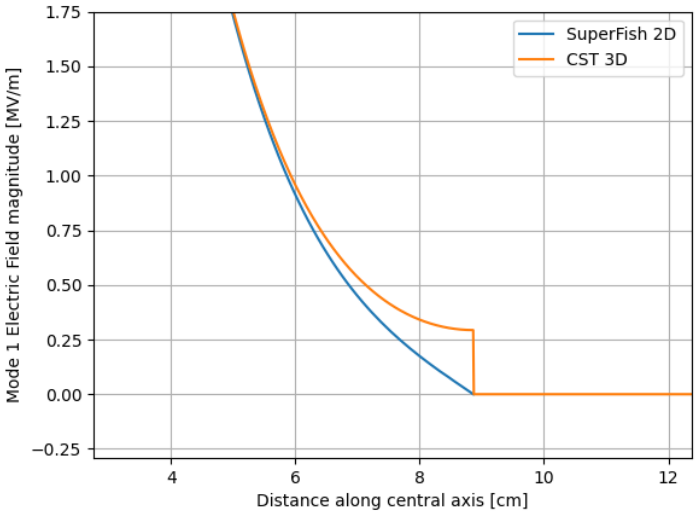

(c) Field on axis at the cavity boundary.

Figure 4.12: A comparison of the fundamental mode electric field on axis in CST and SuperFish for the 352 MHz cavity.

The main discrepancy between the field appears at the cavity boundary as shown in Figures 4.11c and 4.12c, likely caused by a difference in boundary conditions between the two simulations. In SuperFish the simulation domain is the inside of the cavity with no knowledge of anything outside of the boundary. In CST, the simulation domain is a large solid cuboid of

perfect electrical conductor (PEC). The cavity sits within the PEC cuboid. The cavity material in CST is vacuum; the cavity exists a solid piece of vacuum in the shape of the cavity within the PEC. Since the electric field lines and frequency is purely defined by the cavity geometry, no material information is needed. Specifying the material properties of the cavity is only needed for calculating power losses, shunt impedance and related quantities. In CST these quantities are calculated by inputting the material properties in post-processing. In CST the boundary conditions of the simulation are defined at the edges of the PEC. Although both simulations use the same boundary conditions, the boundary itself is in a different location, therefore in SuperFish the electric field goes to zero at the cavity boundary where as in CST it does not.

This discrepancy can be rectified in CST by increasing the length of beam pipe simulated or by moving the boundary condition to the end of the pipe. Extending the length of the beam pipe should reflect reality more closely as the real cavity will of course be installed on a beam pipe much longer than the cavity itself. At this stage this relatively minor difference was left unchanged due to time concerns, to ensure that a 3D field map could be produced for 6D tracking simulations.

4.7 BDSIM simulation

To verify our 1D model and ensure that the cavity design performs as expected, the 3D field map from CST was exported into BDSIM for a 6D tracking simulation of LhARA stage 1 with a realistic cavity. The gap voltage in the simulation was informed by the 1D simulations described earlier and the simulation particle was protons.

Figures 4.13 and 4.14 show the results of the BDSIM simulation for the 201 MHz and 201 MHz cavities respectively. The initial longitudinal distribution of the beam was a uniform beam with the nominal $\pm 2\%$ energy spread and 0 s bunch length to reflect the extremely small initial bunch length of the order $\sim$10 fs at the laser ion source.

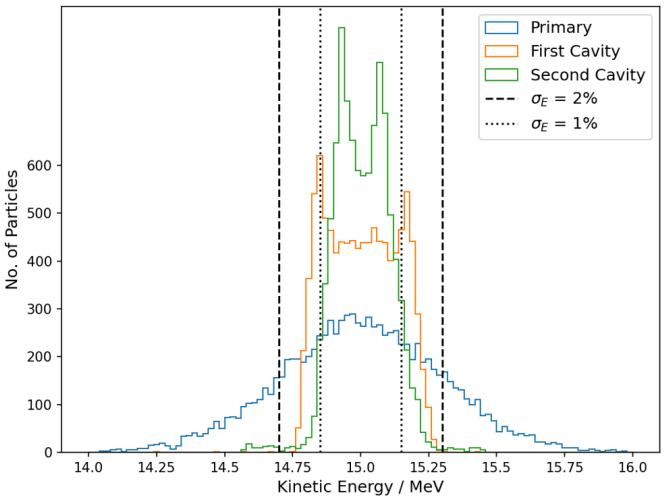

Figure 4.13: Results of the BDSIM simulation for the 201 MHz cavity. Primary refers to a beam with no RF cavities. The first and second cavity lines show the beam after the first and second cavities respectively.

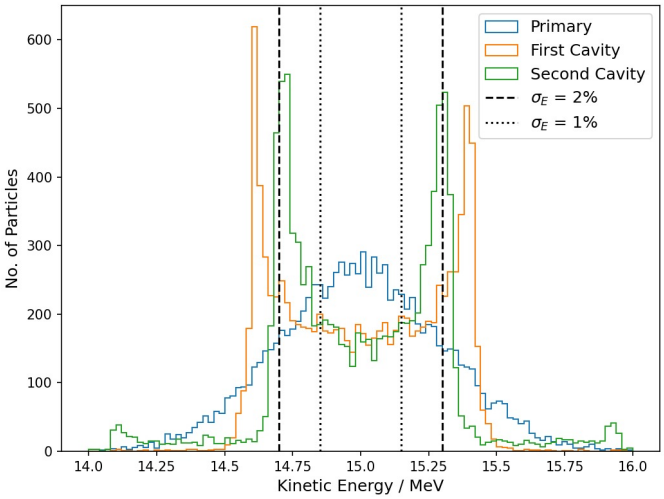

Figure 4.14: Results of the BDSIM simulation for the 352 MHz cavity. Primary refers to a beam with no RF cavities. The first and second cavity lines show the beam after the first and second cavities respectively.

4.8 Cavity frequency comparison

From the simulation results in both SuperFish and CST, it is clear that the 201 MHz cavity out-performs the 352 MHz. It is able to achieve better bunching for both particle types, with the lower frequency able to focus a larger portion of the beam. As such, it is recommended to proceed using the 201 MHz cavity design.

There was some initial investigation into the use of an even lower frequency for the cavity such as 50 MHz, but this was disregarded due to the impracticality of the larger cavity diameter required; over 2 m on the initial design. If further control over bunching is required, it may be necessary to investigate these lower frequency options further.

Chapter 5

Magnet Design

Other than the Gabor lenses, all of the LhARA stage 1 magnets are a standard design and are used in the majority of accelerators globally. This means that the general designs for such magnets are well-known, but still need optimising for the specific design parameters of LhARA. Producing accurate magnet designs is vital for determining the range of possible magnet strengths for the lattice optimisation (Chapter 3), and ensures that the magnet strengths can be met without the magnetic flux saturating in the yoke material of the magnet, with reasonable coil currents.

5.1 Magnetic Analysis Tools

The magnetic lenses for LhARA Stage 1 were designed and modelled in 2D using industry standard magnet design software. The decision was made to analyse in 2D over 3D to save time and effort when optimising the designs, since using a 3D software such as OPERA [30] is far more complicated and involved to produce even a single design, let alone to iterate and improve it.

The 2D slices shown in this chapter are assumed to be in centre of a thick magnet, where there are no fringe fields or effects that could cause distortions in the magnetic fields. Such

fringe fields are also neglected in the lattice calculations in Chapter 3, but should integrate to a negligible contribution over the full length of each magnet with precise fabrication techniques. The true fringes can be be analysed in future studies using 3D modelling software such as OPERA or BDSIM [26].

5.1.1 FEMM 4.2

For the main LhARA beamline magnets, the FEMM 4.2 [31] (Finite Element Method Magnetics) software was used for 2D modelling. Here the magnet cross-section was drawn (along with the relevant coils) in 2D, and then parameters such as the material properties, coil currents, and boundary conditions were set. A standard triangular mesh was used to discretise the problem space (with linear interpolation between mesh points) and then to solve using a preconditioned conjugate gradient (PCG) algorithm with Symmetric Successive Over-Relaxation (SSOR) preconditioner [32], [33].

FEMM 4.2 was used to model both of the vertical arc elements - the dipoles and quadrupoles - and the extraction/shaping octupoles (REF earlier diagram or put one here?) in LhARA stage 1. Symmetries in the magnets were exploited to reduce the computational power and memory required for the solver. For the quadrupoles and octupoles, a quarter-section of the magnet was drawn, and the boundary conditions were set to require $\vec{B}$ to be perpendicular to the symmetric axes. The dipole design necessitated a half-section of the magnet to be modelled, however due to the relative simplicity of the dipole design this could be compensated with a larger mesh size without losing significant precision.

5.1.2 Pandira

The Pandira solver is a sister program to the SuperFish code used for the RF cavity optimisation in Chapter 4; part of the POISSON/SUPERFISH group of codes [34]. Pandira is used here to efficiently model permanent magnet systems in 2D, where FEMM does not work. This solver was used to model the permanent magnet quadrupole nozzle in Section 5.3.6.

Pandira solves magnetostatics problems via a 'direct' method, using an efficient matrix method instead of iteration for problems that either do not converge or converge very slowly. This is possible due to the topological regularity of the mesh used, meaning that the coefficient matrix for the difference equations describing the problem has an identical structure for all problems. Although there are many thousands of equations in the system, this regularity allows the use of a highly efficient sparse matrix method to solve for the magnetic fields within the problem.

5.2 Analysis of Harmonics

While pure multipole fields are desired when optimising the lattice, all real magnets will be a superposition of a main multipole mode and other, higher-order multipole modes (HOMs). These HOMs have two main sources - systematic errors from design constraints when drawing the magnets in CAD software, and random errors from the fabrication process. In this section only the former is considered, as details of fabrication are yet to be decided and so cannot be modelled yet.

Intrinsic symmetries in the magnets and the resulting multipole fields also mean that certain harmonics may be 'allowed' or 'non-allowed' for each design; allowed harmonics are a result of finite magnet pole widths combined with the design symmetries and will be present in all magnets, while large non-allowed harmonics are indicative of fabrication and assembly errors. The allowed harmonics can be calculated in accordance with the following formula:

$$n = N(2m + 1) \tag{5.1}$$

where n is the order of the multipole component (i.e. $n = 1$ gives the dipole component, $n = 2$ the quadrupole, $n = 3$ the sextupole etc.), N is the order of the fundamental field, and m is an integer ≥ 1. Note that this does not correspond directly to typical notation for the multipole string parameter k, where k_0 is the dipole strength parameter, k_1 is the quadrupole strength parameter and so on.

There are two main ways of calculating HOMs, each with their own pros and cons. These are described below in sections 5.2.1 and 5.2.2. A more comprehensive explanation of each can be found in [35].

5.2.1 Polynomial Residual Fitting

The simplest way to calculate HOMs is by fitting a polynomial to the curve of the transverse magnetic field strength. By plotting the absolute value of B on a radial contour from $r = 0$ to $r = r_{\mathrm{GFR}}$ (the edge of the good field region in the magnetic aperture), fitting an n^{th} order polynomial will give the HOMs up to order n by looking at the coefficients of each monomial and converting to the relevant magnetic strength parameter. Equation 5.2 shows the expansion that can be used to fit such a curve, which is truncated at the desired order.

$$B(r) = k_0 + B\rho\left(k_1 r + \frac{1}{2}k_2 r^2 + \frac{1}{6}k_2 r^3 + ...\right) \tag{5.2}$$

While this method is easy to implement and visualise using e.g. a least-squares residual fit, and in principle can produce useful results, it is not very robust. Firstly, the set of monomials does not form an orthogonal basis of the parameter space. This means that it is possible to construct a set of data using a given monomial, and obtain a fit to the same data with a different combination of monomials. The coefficients obtained will depend on the order at which the fit is truncated, which leads to inaccuracies in predictions of beam behaviour.

It is natural to think this problem can be resolved by increasing the order of the fit to high n, but this results in a second problem. Due to the modelling of the magnets with finite element methods, the precision of a high-n fit is greatly affected by small inaccuracies in the mesh. The polynomial will begin to fit any noise in the mesh from interpolations or errors below the convergence threshold, resulting in over-fitting and larger-than-expected HOMs. For the octupoles and quadrupoles modelled here, these errors became very noticeable around $n = 4$, and fits for $n \geq 6$ were completely unusable.

5.2.2 Fourier Analysis

A more robust method for calculating the HOMs for a given magnet is done with a Fourier analysis. An appropriate set of basis functions for fitting the multipole harmonics are provided by the functions $r^{n-1}e^{in\theta}$. By writing the magnetic field in polar coordinates (Equation 5.3), information about the vector properties of the field can be kept, allowing the coefficients of these functions to be calculated by a Fourier transform in the following way:

$$\vec{B} = B_\theta + iB_r = \sum_{n=1}^{\infty} C_n r^{n-1} e^{in\theta} \tag{5.3}$$

Suppose that M measurements of the field are taken at equidistant intervals around a contour of constant r, at $\theta = \theta_m$, where:

$$\theta_m = 2\pi \frac{m}{M}, \qquad m = 0, 1, 2...M-1 \tag{5.4}$$

This measurement of B_m is given as a complex number, describing both the azimuthal and radial components of the field at that point. Now, summing over all the measured points for a given integer n' gives:

$$\sum_{m=0}^{M} B_m e^{-2\pi i n' m/M} = \sum_{m=0}^{M} \sum_{n=1}^{\infty} C_n r_0^{n-1} e^{2\pi i (n-n')m/M} \tag{5.5}$$

where r_0 is the radius of the contour on which the measurements of $\vec{B}$ were taken. Noticing that, unless $n = n'$, the summation over m vanishes, equation 5.5 can be simplified, and rewriting n' as n gives the coefficients C_n as:

$$C_n = \frac{1}{M r_0^{n-1}} \sum_{m=0}^{M} B_m e^{-2\pi i n m/M} \tag{5.6}$$

$$= A_n + iB_n \tag{5.7}$$

This is clearly implemented with a discrete Fourier transform over the measured points. As already mentioned, this method has the advantage of not depending on the truncation order, since all orders are automatically calculated, and also cannot over-fit the data since the random noise from the mesh will not have a periodic structure to pick up in the Fourier transform.

An additional advantage of the Fourier method over the polynomial method is that the fit actually increases in accuracy with increasing radius and n. Supposing that the field was measured with an accuracy ΔB_m, the accuracy of the calculated coefficient will then be:

$$\Delta C_n \propto \frac{\Delta B_m}{r_0^{n-1}} \tag{5.8}$$

Hence the accuracy in the coefficient calculation will be greatest when r_0 is maximised. Furthermore, the ΔB accuracy of the field predictions from the multipole expansion is increased for $r < r_0$ (although the accuracy will deteriorate rapidly for $r > r_0$, especially for large n), as this accuracy will be:

$$\Delta B \propto \Delta C_n \left(\frac{r}{r_0}\right)^{n-1} \tag{5.9}$$

A final advantage of the Fourier method is the way that it intrinsically separates the HOMs into the normal and skew components, as seen in Equation 5.7 where C_n is complex. Here the real part (A_n) refers to the normal component and the imaginary part (B_n) refers to the skew component. While practically there is no difference between the two (a magnet can simply be rotated to change it from normal to skew, and vice versa), each of the modelled magnets are designed to be either normal or skew, and so should only have one component or the other.

For the reasons discussed, the Fourier analysis method is used in Section 5.3 to calculate the harmonics of each magnet. These harmonic calculations also have a second use, allowing the field in each magnetic aperture to be stored for simulation use in a far more concise format than when saving the magnetic field at each point in the aperture (which can cause file sizes to become large very quickly if taking saving in high detail). Full tables of harmonics can be seen

in Appendix II.

5.3 Magnet Design

5.3.1 General Design Parameters

Several design parameters had to be determined before the designing of the magnets could commence, primary among which was the beam pipe diameter. The LhARA design document [24] specifies a maximum desired beam diameter of 30 mm (defined as the 2σ diameter of the beam), and so for this section of the study a beam pipe radius of 50 mm was decided, requiring a good field region (GFR) radius of 35 mm - just under the 5σ radius of the maximum beam size. This not the same as was calculated in Section 3.3.1, and differs due to requirement for a radius value early on in the design process, well before the value of 3.65 cm was reached in the lattice design studies. From the nominal energy/momentum of the beam, the beam rigidity $B\rho$ was calculated as 0.561 T m.

The magnet materials also had to be specified in order to set the B/H curve for the simulation of the magnetic response. In this analysis pure iron was used as the material for all the magnet yokes, and pure copper was used for all the coil. All gaps were simulated with air, although this has the same magnetic properties as a vacuum for the purposes of these simulations. All B/H curves were taken from the FEMM 4.2 materials library. Limits on the permitted field strengths in the magnetic yokes were heuristically imposed to be 2.0 T, meaning that any design with values greater than this was automatically rejected. This ensures that the magnets will not have saturation problems limiting the magnetic fields later on in the design process.

When considering the field purity in the magnetic apertures, a nominal purity of 99.9% was desired from each magnet. In this analysis, this means the contribution of HOMs and residuals should be less than 0.1% when compared to the measured B-field strength of the ideal component at $r = r_{\mathrm{GFR}}$. The magnets below were designed to satisfy (or at least be close to satisfying) this condition at the maximum required field strength of any particular lattice

element. The magnets can be tuned later to a lower field strength by altering the current in the coils if required.

In all simulation cases below, a mesh size of 1 mm was used in all regions except the aperture region, where the mesh size was significantly reduced to 0.03 mm. This system provided a good payoff between simulation speed and accuracy, as a very fine mesh was not necessary in areas where the field was not relevant to the aperture. The mesh shape was triangular (as specified in Section 5.1.1, and auto-generated using the usual FEMM 4.2 algorithm.

5.3.2 Magnetic Pole Tip Shapes

In this analysis, the pole tip shapes of the magnets were designed to follow the ideal equipotential lines of the desired harmonic. This ensures the magnetic field lines will come out perpendicular to the edge of the yoke, maximising the field strength in the aperture in the ideal case. As always, the pole tips are not ideal and must be truncated to some finite width which will introduce small errors, however there was not time in this analysis to optimise the pole tip shapes beyond the equipotential lines of the required multipole harmonic. Further studies would focus on optimising the width and shape of the pole tips beyond this, for example with shims on the pole tip corners. A table of the equipotential lines for each magnetic component up to octupole is given below in table 5.1:

Harmonic	Equipotential Equation
Dipole	$y = \pm r$
Quadrupole	$2xy = \pm r^2$
Sextupole	$3x^2y - y^3 = \pm r^3$
Octupole	$4\left(x^2y - y^2x\right) = \pm r^4$

Table 5.1: Table of equipotential lines for multipole harmonics up to octupole giving pole tip shapes.

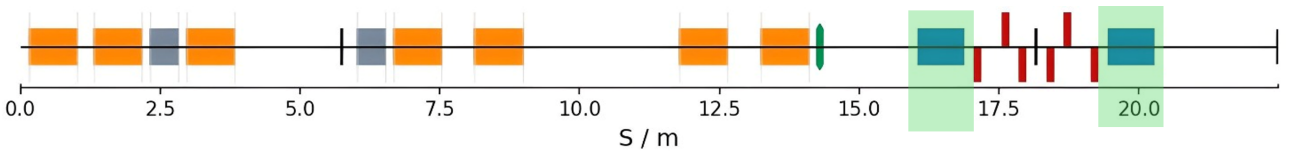

Figure 5.1: Location of the dipole magnets in the LhARA stage 1 beamline.

5.3.3 Vertical Arc Dipole

Dipole Design Parameters

The vertical arc dipoles are used in two positions within the LhARA stage 1 beamline, to switch the end station of the beam between the beam dump, the stage 1 *in-vitro* end station, and the beamline to stage 2. Due to these requirements, the dipole was chosen to be a 'C-shaped' design, allowing the beam to be switch between directions far more easily than with an 'H-shaped' design. The dipole lengths and bending radii were specified to be 0.8 m and 45 deg each, giving a bending radius of 1.019 m. Combined with the previously derived beam rigidity of 0.561 m this gives a required dipole field strength of 0.551 T.

Dipole Magnetic Simulation Output

The vertical arc dipole design is shown in figure 5.2. This shows the top half of the dipole (i.e. the full magnet would be mirrored along the bottom of the figure), and the materials can be distinguished as follows: the iron yoke is the half-C shaped area containing the vast majority of the magnetic flux, the coils are the two boxes either side of the yoke gap, and the air is everywhere else. As seen from the contour key on the right side of the figure, the maximum B-field in the yoke is $\sim$1.25 T, well below the saturation value of 2 T.

In this design, the coils have a current of 18.6 kA turn, with 50 turns over an area of 20 000 mm^2. This results in an overall current density of 0.93 $\mathrm{A\,mm^{-2}}$, which is low enough to allow the coils to be radiatively cooled instead of water cooled. This reduces magnet complexity and power consumption and is highly desirable for such magnets, but requires current densities less than 1 $\mathrm{A\,mm^{-2}}$.

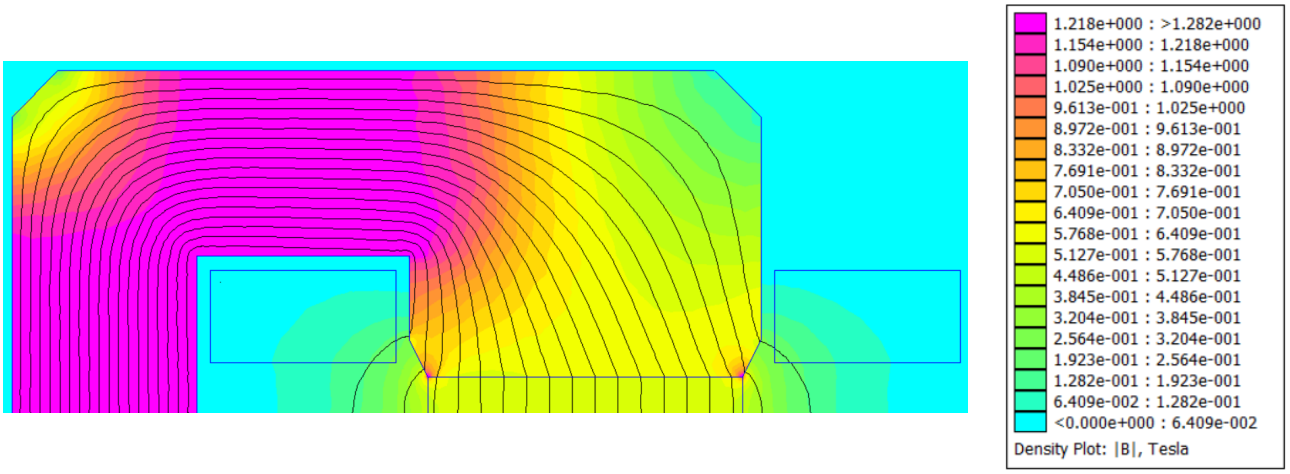

Figure 5.2: Cross-section render of the vertical arc dipole.

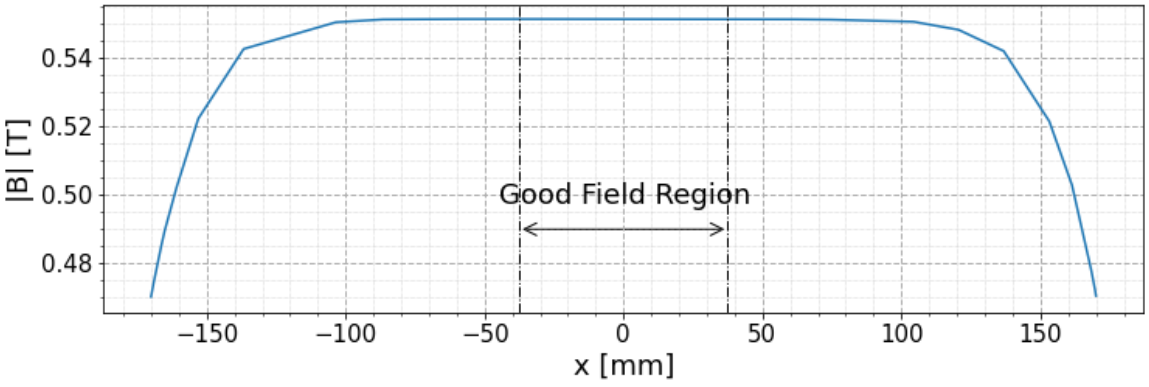

(a) B-field in the dipole gap, GFR labelled.

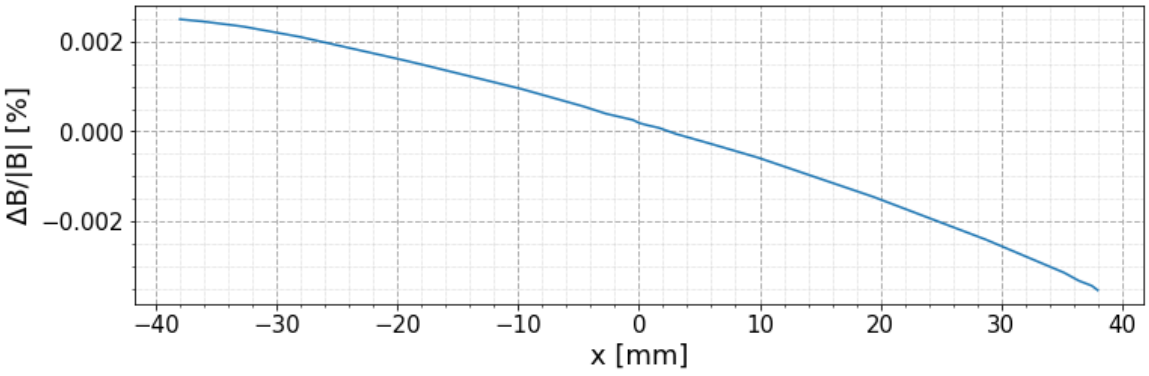

(b) B-field residuals in the GFR of the dipole gap.

Figure 5.3: B-field and residuals along the dipole gap, on the symmetric axis.

Figure 5.3 shows the magnitude of B along the symmetric axis of the dipole gap. As can be seen in Figure 5.3a, the B-field is very flat and exactly the desired strength throughout the whole of the GFR (indicated on the figure). This is confirmed with the residuals shown in Figure 5.3b, showing a field error well in excess of the required 99.9%. Given how much in excess of the required purity this is, it is possible that some purity could be sacrificed for a smaller magnet footprint in further studies, although the extent to which this is possible is yet to be determined.

Analysis of Dipole Harmonics

While every harmonic is "allowed" for this dipole design, it is clear than none have a significant effect on the overall dipole field. Figure 5.4 shows the skew and normal components of the dipole, along with a zoomed view for the HOMs between dipole and octupole for comparison of size. Values for the multipole components (k_n) of the dipole are given in Table 5.2 up to the decapole component. Higher order harmonics than this have such negligible contributions to the B-field that they are not worth listing.

It is worth noting how the value of k_n increases with n. A naive observer might see this as bad, increasing the contribution of large-n HOMs to the error in the B-field, however this does not take into account the r^{n-1} dependency of the field as seen in Equation 5.2. This dependency more than cancels the increase in k_n, as shown in Table 5.2. It should also be noted that these results don't distinguish between $\pm k$, so the contributions of the HOMs could be positive or negative relative to the main multipole component. Addionally, low precision on k may result in small inconsistencies between the computed B and the B shown in Figure 5.4.

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.551\,\mathrm{m}^{-1}$	$0.551\,\mathrm{T}$
Quadrupole (k_1)	$1.0 \times 10^{-5}\,\mathrm{m}^{-2}$	$0.0\,\mathrm{m}^{-2}$	$0.160\,\mu\mathrm{T}$
Sextupole (k_2)	$1.2 \times 10^{-4}\,\mathrm{m}^{-3}$	$1.0 \times 10^{-5}\,\mathrm{m}^{-3}$	$0.041\,\mu\mathrm{T}$
Octupole (k_3)	$4.5 \times 10^{-3}\,\mathrm{m}^{-4}$	$2.9 \times 10^{-4}\,\mathrm{m}^{-4}$	$0.018\,\mu\mathrm{T}$
Decapole (k_4)	$0.29\,\mathrm{m}^{-5}$	$0.015\,\mathrm{m}^{-5}$	$0.010\,\mu\mathrm{T}$

Table 5.2: Table of dipole higher-order harmonics up to decapole. Full table in Appendix II.

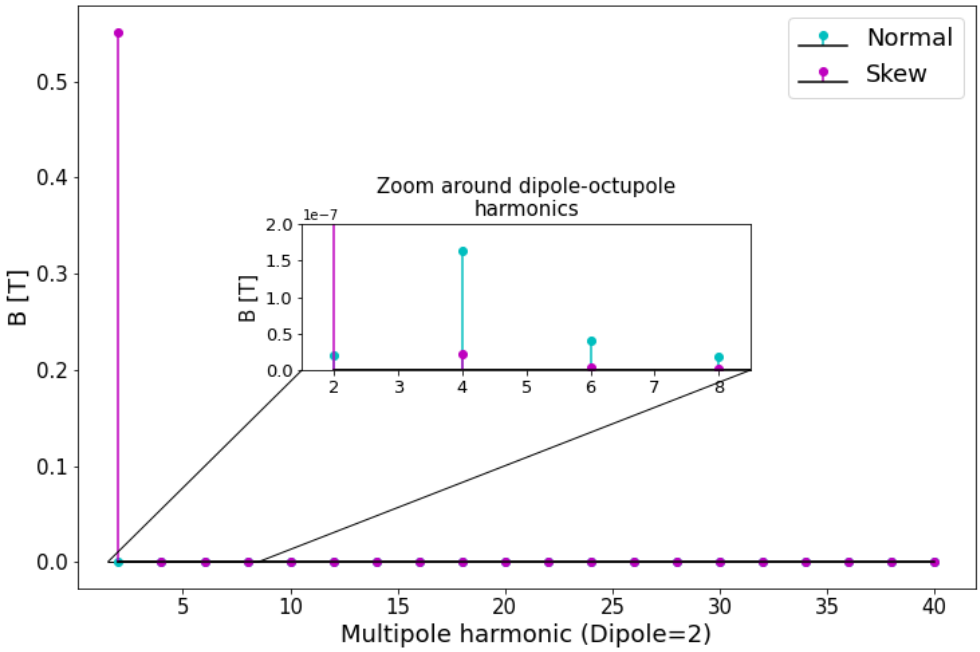

Figure 5.4: Harmonics of the vertical arc dipole, with zoom around the quadrupole, sextupole and octupole (largest HOM) components. Relative strengths given in terms of B at $r = r_{\mathrm{GFR}}$

From Table 5.2 it is clear that the dipole meets the imposed purity requirements with essentially no other components. While this is exactly what was required, it is interesting to consider whether a small quadrupole component could provide a useful focusing effect for off-energy particles through the bends (since lower energy particles will experience a smaller bending radius that higher energy particles). This would be an area for potential future study, although such a combined function magnet would have to be analysed when integrated with the full lattice to find any knock-on effects.

5.3.4 Vertical Arc Quadrupole

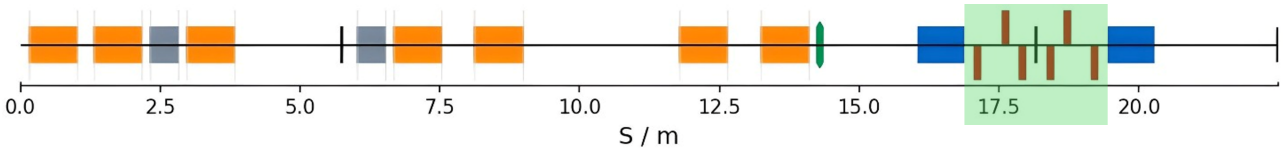

Figure 5.5: Location of the quadrupole magnets in the LhARA stage 1 beamline.

Quadrupole Design Parameters

The LhARA stage 1 beamline has 6 quadrupoles, positioned in the middle of the vertical arc. These alternate focusing/defocusing directions, and are tuned in strength to give the desired twiss parameter matching discussed earlier in Chapter 3. This allows the beam to be focused and manipulated to give the desired spot size for the *in-vitro* end station, which can be further switched simply by varying the currents in the quadrupole coils.

The lattice studies in Section 3.2.1 gave a maximum required $k_1 = 32.0\,\mathrm{m}^{-2}$, although smaller spot sizes could be reached more easily with values higher than this. Not all the quadrupoles will have to be this strength either, but in this analysis $k_1 = 32.0\,\mathrm{m}^{-2}$ was the initial target, with lower values obtained by reducing the current in the quadrupole coils. After reaching the initial k_1 target, the quadrupole strength was maximised above this subject to the requirements on magnet footprint and field purity as before.

Quadrupole Magnetic Simulation Output

Figure 5.6 shows a render of the top-right quarter of the proposed quadrupole design. Like the dipole design presented in Section 5.3.3, the iron yokes can clearly be distinguished as containing the vast majority of the magnetic flux. The coils are either side of the yokes as before. Unlike before the shape of the coil cross-section has not been optimised, and most likely does not conform to the shape of a physical coil - only the absolute cross-sectional area was considered here. Further studies would optimise the total coil shape based on available coil geometries, although this should not affect the design significantly.

In this design, it is clear that the iron yokes are approaching saturation, but there is still some leeway with the maximum field approximately 1.85 T. This is achieved with a total current of 21.6 kA turn with 18 coil turns over a coil area of $4070\,\mathrm{mm}^2$, resulting in a current density of $5.31\,\mathrm{A\,mm}^{-2}$. This is significantly higher than in the dipole design, and so will require water-cooling for the coils, but is still accessible with most commercially available water-cooled coils.

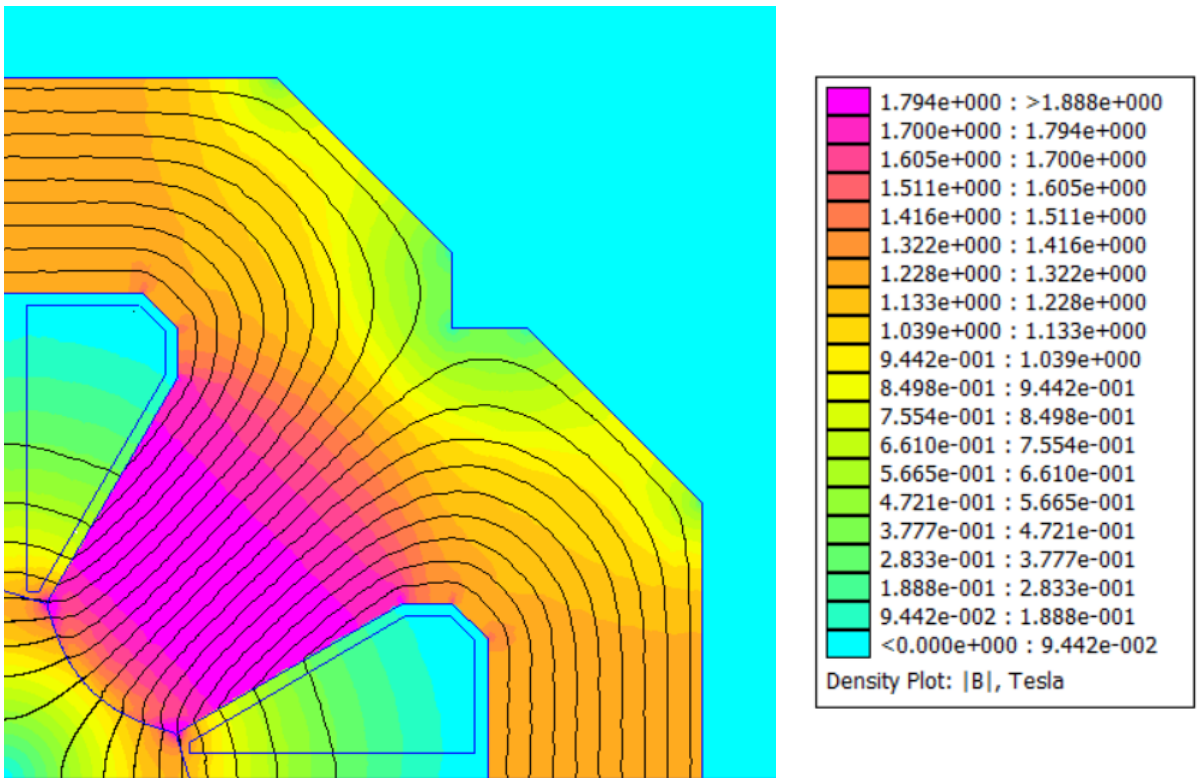

Figure 5.6: Cross-section render of the vertical arc quadrupole.

The quadrupole gradient achieved with this design is $20.7\,\mathrm{T\,m^{-1}}$, shown in Figure 5.7a, which corresponds to $k_1 = 36.96\,\mathrm{m^{-2}}$. This significantly exceeds the initial requirement of $32\,\mathrm{m^{-2}}$, and suggests that smaller spot sizes than previously calculated could be accessible with further lattice optimisation. The residuals to this quadrupole gradient can be seen in Figure 5.7b with a maximum residual of 0.5 mT at $r = 35\,\mathrm{mm}$ (upper). This corresponds to a maximum deviation of 0.1% at this point (lower), just about meeting the field purity requirements specified previously. Looking at the percentage error for $r < 5\,\mathrm{mm}$, it can be seen that the error actually exceeds 0.1% for small r, but it must be noted that this is mainly due to small mesh interpolation errors that become proportionally large compared to the true B-field at small r. This is clear when looking at the graph of the absolute residuals.

Analysis of Quadrupole Harmonics

It is clear from the shape of the residuals to the quadrupole component that there are some non-negligible HOMs present in this design. As show in Figure 5.8 and Table 5.3, these are the expected allowed harmonics from the dodecapole, 20-pole, 28-pole and 36-pole components (k_5,

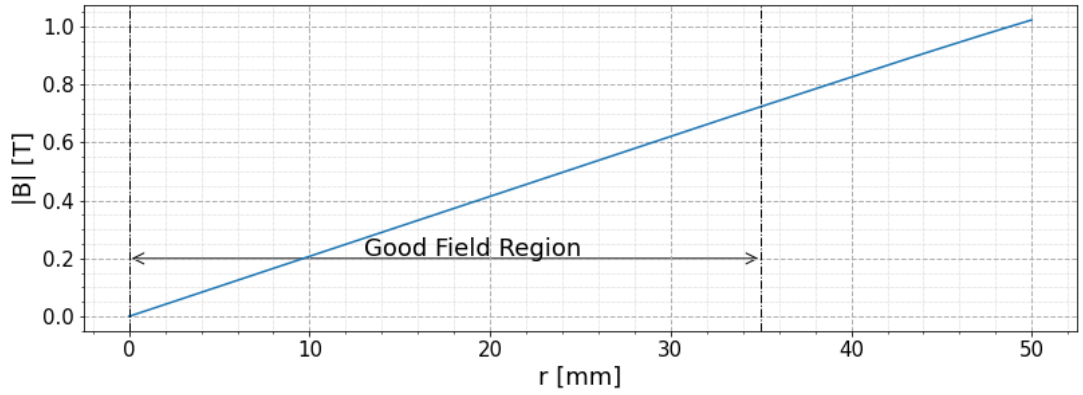

(a) B-field in the quadrupole gap, GFR labelled.

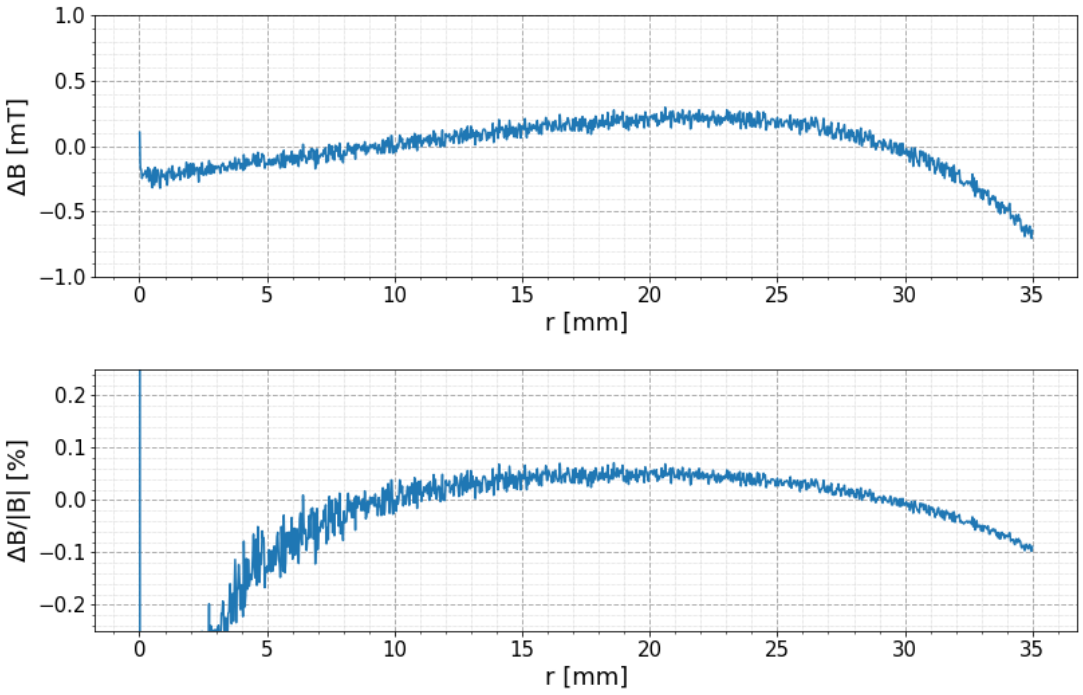

(b) Absolute B-field residuals (upper) and relative percentage residuals (lower) in the GFR of the quadrupole gap.

Figure 5.7: B-field and residuals along the quadrupole gap.

k_9, k_{13}, k_{17}). Most of the residual (∼0.13%) is from the dodecapole component; this is actually larger that the targeted 0.1%, but some of the other HOMs marginally cancel this to give this design a total residual below the field purity requirement.

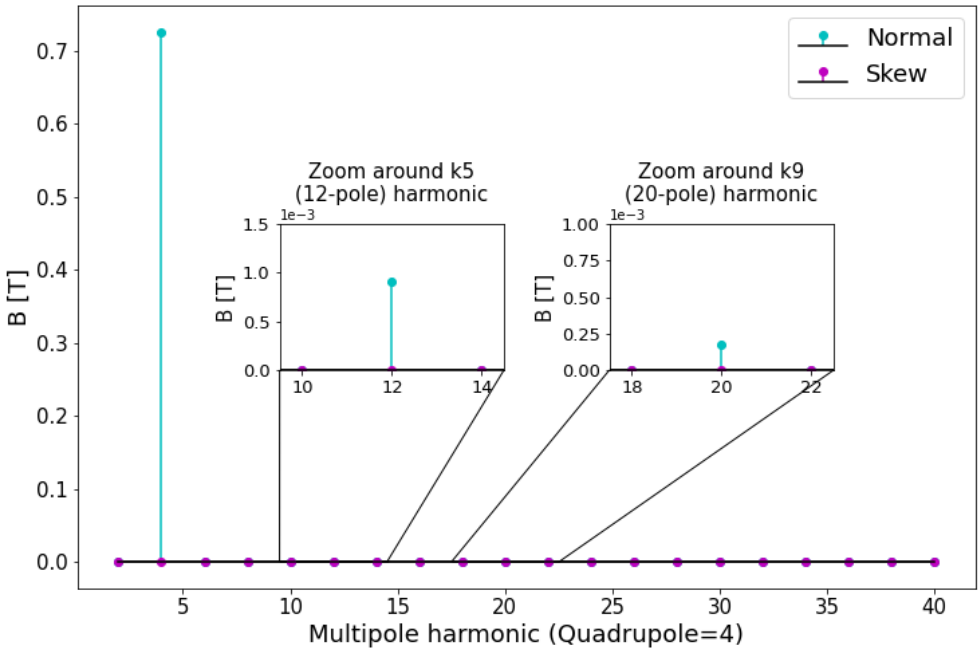

Figure 5.8: Harmonics of the vertical arc quadrupole, with zoom around the largest HOM components. Relative strengths given in terms of B at $r = r_{\mathrm{GFR}}$.

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.0\,\mathrm{m}^{-1}$	0.0 T
Quadrupole (k_1)	$36.96\,\mathrm{m}^{-2}$	$0.058\,\mathrm{m}^{-2}$	0.762 T
Sextupole (k_2)	$0.0\,\mathrm{m}^{-3}$	$0.0\,\mathrm{m}^{-3}$	0.0 T
Octupole (k_3)	$0.0\,\mathrm{m}^{-4}$	$0.0\,\mathrm{m}^{-4}$	0.0 T
Decapole (k_4)	$0.0\,\mathrm{m}^{-5}$	$0.0\,\mathrm{m}^{-5}$	0.0 T
Dodecapole (k_5)	$3.7 \times 10^{6}\,\mathrm{m}^{-6}$	$1.6 \times 10^{4}\,\mathrm{m}^{-6}$	0.908 mT
20-pole (k_9)	$1.4 \times 10^{15}\,\mathrm{m}^{-10}$	$1.7 \times 10^{13}\,\mathrm{m}^{-10}$	0.174 mT
28-pole (k_{13})	$3.7 \times 10^{24}\,\mathrm{m}^{-14}$	$5.1 \times 10^{22}\,\mathrm{m}^{-14}$	0.039 mT

Table 5.3: Table of quadrupole higher-order harmonics up to dodecapole, including allowed harmonics up to 28-pole. Full table in Appendix II.

Table 5.3 gives a clear example of how the allowed harmonics can influence the structure of the B-field harmonics. There are negilible contributions to the field strength from the low-n harmonics, however the dodecapole component contributes nearly 1 mT to the B-field

at $r = r_{\mathrm{GFR}}$, with smaller (but still noticeable) contributions from the 20-pole and 28-pole harmonics.

Areas of future study should work on reducing the error from the higer order components, potentially with a wider pole yoke, or by introducing a more optimal pole-tip shape with e.g. shims. The coil geometry must also be further explored to fit with the cross-sectional shapes of physical coils - the total current can be reduced to accommodate this, as currently the quadrupole strength exceeds the required value from the lattice optimisation.

5.3.5 Extraction Octupole

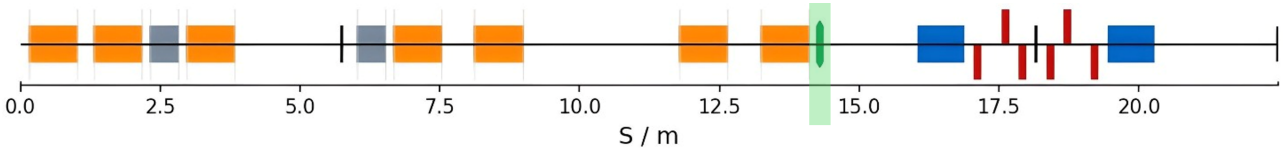

Figure 5.9: Location of the octupole magnet in the LhARA stage 1 beamline.

Octupole Design Parameters

An Octupole is used in one position in the LhARA stage 1 beamline, and is the only type of magnet other than the Gabor lenses to be included before the vertical matching arc. The role of the octupole is to manipulate the profile of the beam from an initial Gaussian profile to a flat distribution, as shown in Section 3.3, to give an appropriate dose profile for the end station. The Octupole also has the advantage of third-order focusing, which offsets the perturbations introduced by the Gabor lenses. After the octupoles, the beam profile is given a square transverse profile - this is cleaned back to the desired shape with the use of collimators, discussed in Section 3.3.1.

In the LhARA technical document [24], a nominal desired octupole strength was set at $k_3 = 60\,000\,\mathrm{m}^{-4}$. This was suggested as a reasonably achievable target and the maximum that the LhARA facility would require, but has little basis in formal studies. It is very possible that a significantly lower octupole strength would be required, as some beam simulations suggest

(REF BDSIM octupole studies) a value of $k_3 = 30\,000\,\mathrm{m}^{-4}$ may be all that is required. The ramifications of this are discussed later on in this section.

Octupole Magnetic Simulation Output

As with the quadrupole in Section 5.3.4, the octupole was simulated as a quarter-section of the full magnet, exploiting the high degree of rotational symmetry in the octupole. This quarter-symmetry was chosen over the available eight-fold symmetry as it outputs in a way that is both easier to analyse by Fourier transform and more intuitive to view in a render, despite being more intensive for computational power. The output of the simulation can be seen in Figure 5.10, and interpreted in the same way as previously.

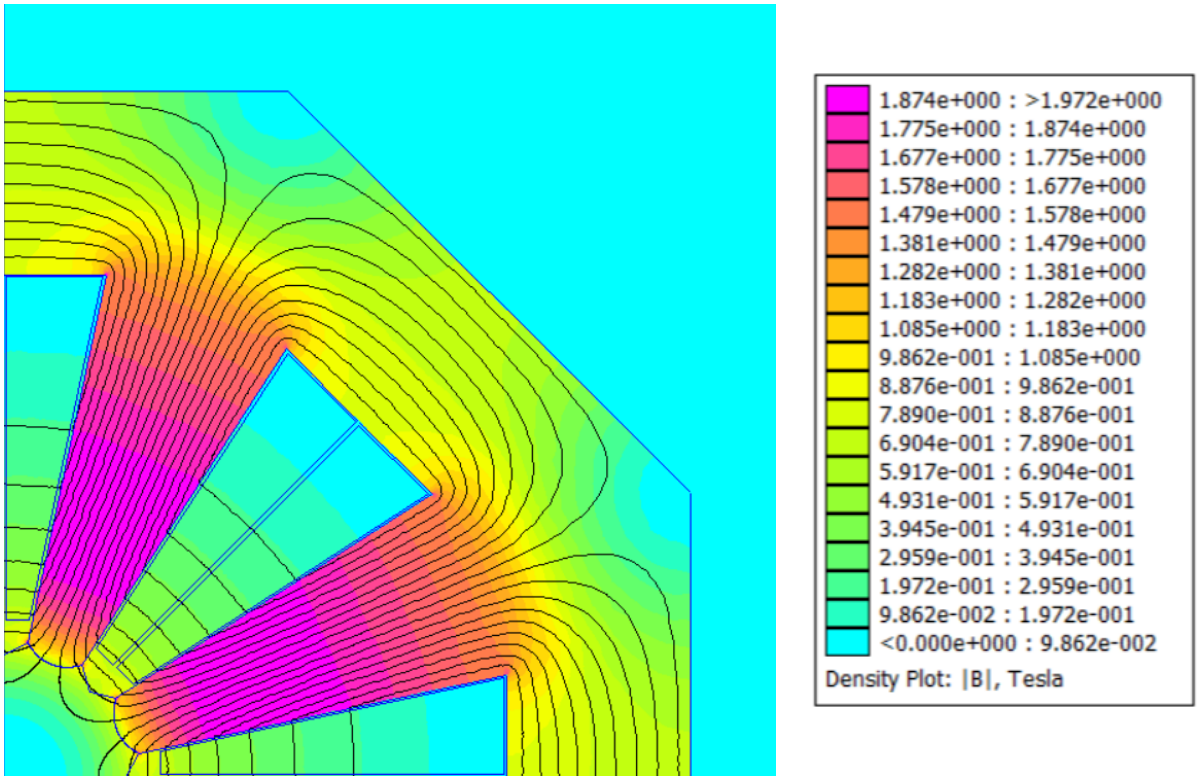

Figure 5.10: Cross-section render of the extraction octupole.

The octupole presented here gives a k_3 of $60\,273\,\mathrm{m}^{-4}$, just about meeting the previously mentioned strength requirements. This k_3 approximately the limit of what this design can handle, as the iron yokes are sustaining a maximum flux of $\sim$1.95 T, essentially reaching saturation. The limiting factor preventing further strength increases is the width of the yoke poles - to fit

in the required number of poles and coils they are significantly thinner than would be found in e.g. a quadrupole magnet. There is some scope to change this, as the field is achieved with a current density of only $2.48\,\mathrm{A\,mm^{-2}}$, meaning that a higher current density could be used to reduce the footprint of the coils and increase space for the yokes. The coils still have the same unoptimised shape as with the quadrupole coils in Section 5.3.4 however, and taking this into account may significantly reduce this possibility.

The coil current here is 8.51 kA turn, with 10 turns over a coil area of $3431\,\mathrm{mm^2}$. As previously stated, this results in a current density of $2.48\,\mathrm{A\,mm^{-2}}$, large enough to require water-cooled coils but not close to the limit of what a water-cooled coil can support. This results in an interesting conundrum: can the current density be reduced to $\sim 1\,\mathrm{A\,mm^{-2}}$ in order to allow for a simpler radiative cooling system (which would have an increased number of turns with lower current in each), or should the current density be increased to $\sim 5\,\mathrm{A\,mm^{-2}}$, a more typical value for water-cooled coils, and the coil footprint reduced to allow for wider yokes and lower magnetic fields inside the body of the magnet. Based on coil shape problem previously discussed it is tempting to say the latter, but given the advantages of radiative cooling over water cooling and the possibility of a much lower k_3 requirement the former should not be discounted as a viable option either. This needs to be studied further to provide a definitive answer.

The B-field in the octupole aperture is shown in Figure 5.11a, and the residuals in Figure 5.11b. Clearly this shows a high level of octupole harmonic purity up to $r = 25\,\mathrm{mm}$, and is within the 0.1% error target up to $r \approx 32\,\mathrm{mm}$ before decreasing to $\sim$0.2% at r_{GFR}. As before this is a result of the allowed octupole HOMs (specifically the 24-pole harmonic), which will have an r^{11} dependence unlike the r^3 dependence of the octupole harmonic, which is why there is no significant error relative to the octupole field until r approaches r_{GFR}. There are also some small contributions from the non-allowed k_7 and k_{15} harmonics which are a result of the symmetries exploited when drawing the octupole, mirroring around the 45 deg diagonal to draw the full quarter of the magnet. These would likely disappear with a more accurate render. As before there is significant noise around the small-r region of the aperture due to mesh interpolation inaccuracies where B is small - this would be significantly smoother in a physical magnet.

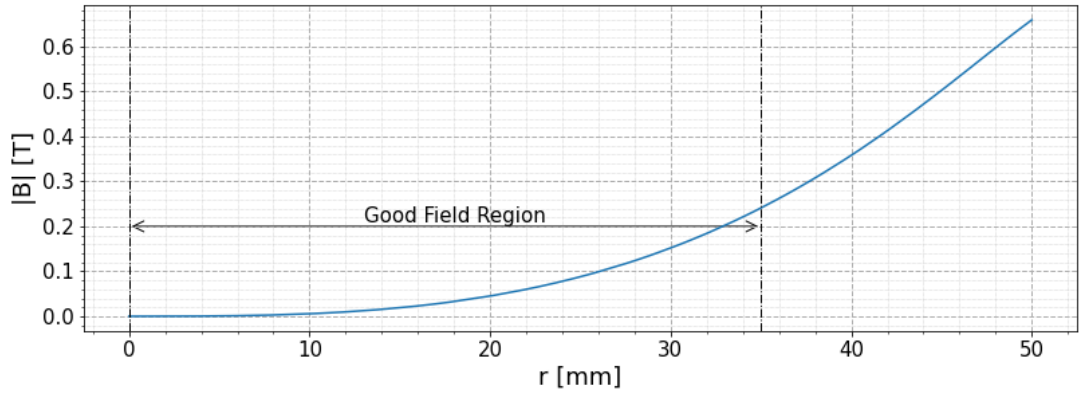

(a) B-field in the octupole gap, GFR labelled.

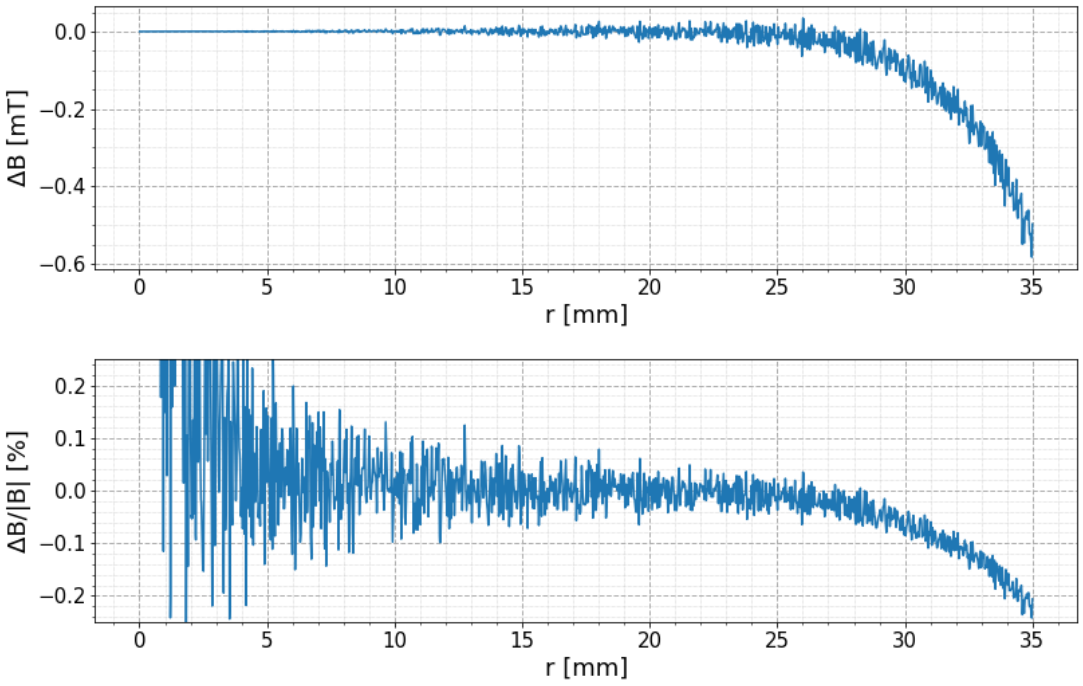

(b) Absolute B-field residuals (upper) and relative percentage residuals (lower) in the GFR of the octupole gap.

Figure 5.11: B-field and residuals along the octupole gap.

Analysis of Octupole Harmonics

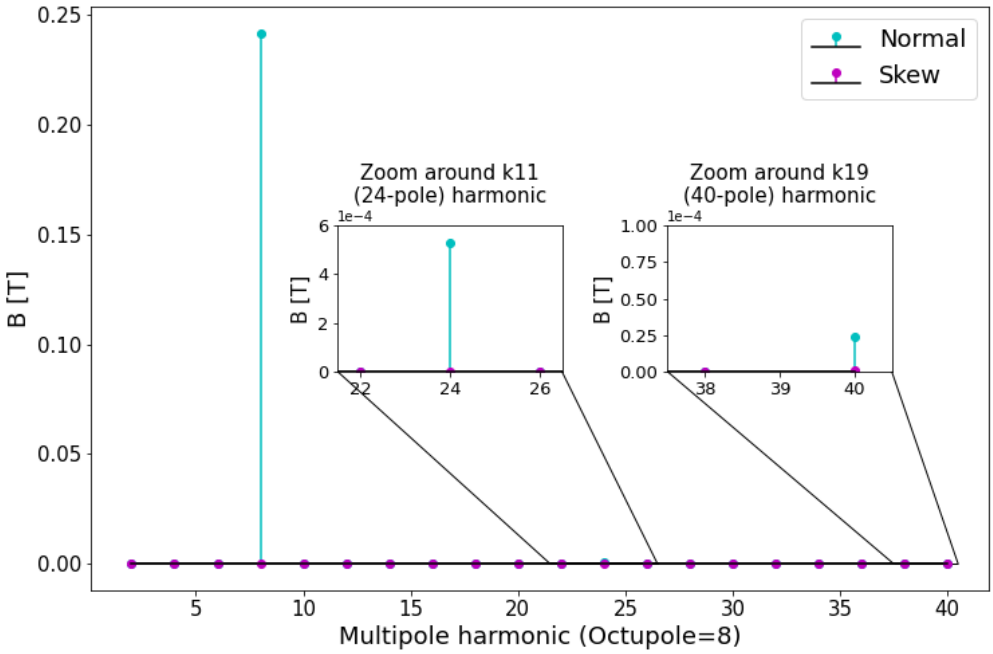

Figure 5.12: Harmonics of the extraction octupole, with zoom around the largest HOM components. Relative strengths given in terms of B at $r = r_{\mathrm{GFR}}$.

As discussed in the previous section, there are some non-negligable HOMs in the octupole field, and these can be seen clearly in Figure 5.12. Here the only non-negligible HOMs are the allowed harmonics: the main octupole mode, the 24-pole harmonic, and the 40-pole harmonic. The contributions to the B-field at $r = r_{\mathrm{GFR}}$ are listed in Table 5.4, but to summarise there is a $\sim$0.22% contribution to the field from the 24-pole harmonic, and a $\sim$0.01% contribution from the 40-pole harmonic. Due to the high-order of these harmonics, they do not become prevalent until very close to the edge of the GFR, and in this case the field purity is above 99.9% until $r = 32\,\mathrm{mm}$.

It is also possible that, due to the lack of research done into the octupole requirements so far, a purity level of 99.8% or even lower could still be acceptable in the octupole, especially if the errors come from the allowed octupole harmonics. A definitive answer to this will require more work on lattice tolerance analysis.

Further studies on the octupole should look into optimising the shape of the magnet pole yokes

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.0\,\mathrm{m}^{-1}$	0.0 T
Quadrupole (k_1)	$0.0\,\mathrm{m}^{-2}$	$0.0\,\mathrm{m}^{-2}$	0.0 T
Sextupole (k_2)	$0.0\,\mathrm{m}^{-3}$	$0.0\,\mathrm{m}^{-3}$	0.0 T
Octupole (k_3)	$60\,273\,\mathrm{m}^{-4}$	$0.167\,\mathrm{m}^{-4}$	0.242 T
Decapole (k_4)	$0.0\,\mathrm{m}^{-5}$	$0.0\,\mathrm{m}^{-5}$	0.0 T
Dodecapole (k_5)	$0.0\,\mathrm{m}^{-6}$	$0.0\,\mathrm{m}^{-6}$	0.0 T
24-pole (k_{11})	$3.9 \times 10^{20}\,\mathrm{m}^{-12}$	$2.4 \times 10^{17}\,\mathrm{m}^{-10}$	0.527 mT
40-pole (k_{19})	$2.4 \times 10^{40}\,\mathrm{m}^{-20}$	$7.9 \times 10^{38}\,\mathrm{m}^{-20}$	0.024 mT

Table 5.4: Table of octupole higher-order harmonics up to dodecapole, including allowed harmonics up to 40-pole. Full table in Appendix II.

and tips, and how the coils would fit around this. Further lattice analysis will provide a clearer picture of the requirements of the extraction octupole, and will determine the necessary coil geometry to be met by confirming the type of coil cooling method that will be necessary.

5.3.6 Permanent Magnet Quadrupole

Immediately after the laser ion source, there are very large space charge forces in the produced beam. In the first few centimetres the space charge forces are compensated, since both positive and negative charge are produced in the initial target interaction, but as the beam propagates these forces become more important. To capture more particles, the first focusing element in the lattice should be as close as possible to the source.

In the current design, the first focusing element is the Gabor lens after the vacuum nozzle. Previous studies have shown that a focusing element such as a permanent magnet quadrupole (PMQ) even earlier in the beam line could result in more particles being captured. A PMQ before the vacuum nozzle, in the vacuum environment of the source rather than the beam line could serve this purpose. Figure 5.13 shows a schematic of the laser ion source, where the PMQ could be placed in the expanding ion beam region. This PMQ would need to be very small and radiation hard to maintain its magnetisation in a region with many particle species of various energies.

The Pandira solver in SuperFish was used to investigate what focusing strengths can be achieved

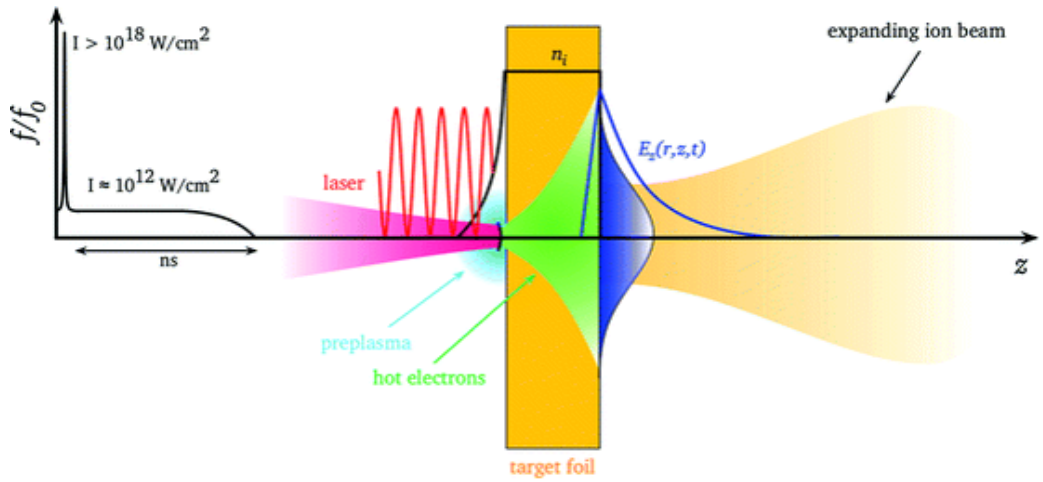

Figure 5.13: Schematic diagram of the laser ion source. A permanent magnet quadrupole could be placed in the region after the source, where the ion beam is expanding from an initially small size.

with a small samarium-cobalt PMQ with a maximum outer radius of 2 cm. To minimise the size of the PMQ, a Halbach array design seemed the most appropriate as they contain only permanent magnet material and support structure, there is no iron in the design. In a Halbach array, wedges of permanent magnet material are arranged ed in a circle to produce the desired field. The size and orientation of the wedges can be adjusted to change the relative strengths and harmonic content of the magnetic field. Since the goal here is a purely quadrupolar magnetic field, sixteen equal wedges are used to form the Halbach array.

2D design

Figure 5.14 shows a diagram of the PMQ in Pandira. The magnetic field lines inside the magnet material and in the good field region are shown.

In this particular design the outer radius of the magnet material was limited to 2 cm and the bore radius is 4.65 mm. The small aperture allows for a very high quadrupole gradient of around 340 T m^{-1}.

These simulations demonstrate that a small focusing element very close to the source with a high gradient could be realised, however, there are many caveats still to investigate. Perhaps the most pressing is how the magnetic materials will respond to radiation damage. Samarium-cobalt

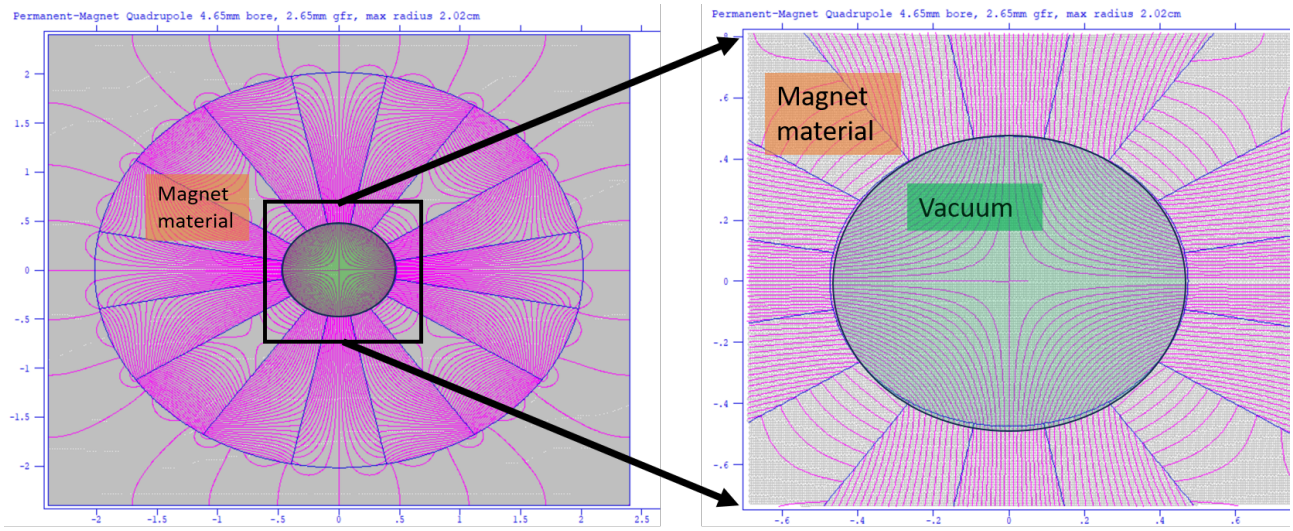

Figure 5.14: The magnetic field produced by a permanent magnet quadrupole Halbach array, with GFR highlighted in green. Axes in cm

.

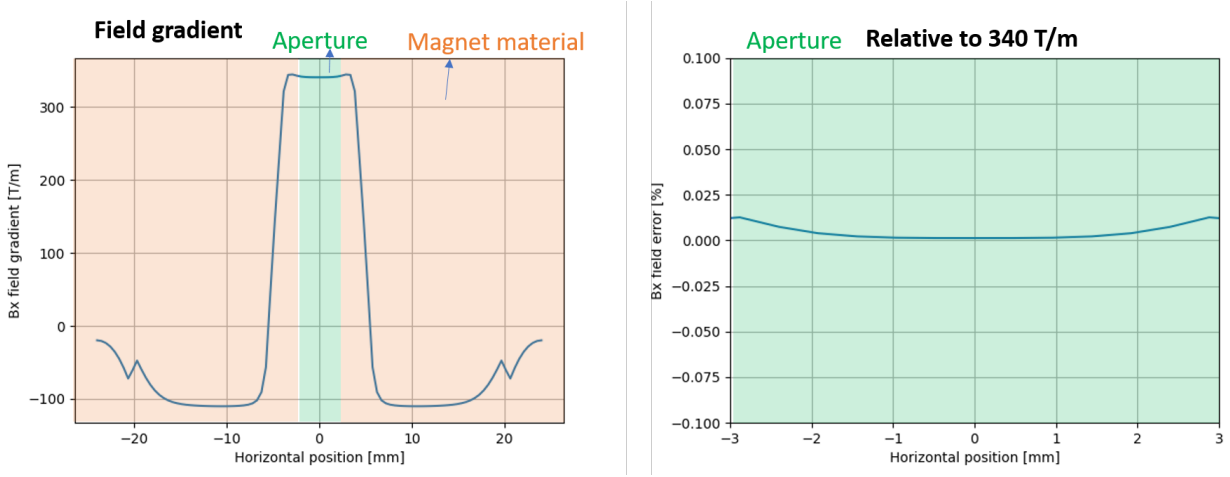

Figure 5.15: Field gradient across the magnet (left) and within the bore (right).

and other magnetic materials can demagnetise if too much energy is deposited in the material, and these magnets would likely need shielding to prevent demagnetisation. Additionally, the wedges would require some support structure; this will likely increase the outer radius and decrease the bore radius for the same magnetic field strength. With further investigation these issues could be surmountable.

Chapter 6

Conclusion

By simulating beams within the LhARA facility, this study has further developed the design of the lattice for LhARA stage 1 beamline. This has been achieved with detailed full-lattice simulations of a Gaussian beam produced by laser-driven acceleration, supported by in-depth studies of the beamline elements. The element studies focused in two areas: design of the RF bunching cavities, where the cavities were optimised in shape and frequency to achieve maximal bunching of a 15 MeV proton beam, and the design of the magnetic lens elements (dipoles, quadrupoles and octupoles) in the lattice, where the magnets are designed to meet criteria defined in the lattice simulations.

Initial lattice optimisation in MAD-X yielded sets of strengths for the Gabor lens elements (modelled as solenoids), capable of producing beam sizes with 2σ diameters of between 1 cm and 3 cm, whilst sufficiently satisfying constraints across the lattice ensuring all lattice configurations facilitate the function of elements not modelled in MAD-x (e.g. RF cavities and the end station). Beam size reduction is shown to be achievable through quadrupole strength manipulation within achievable values. Conversion of the lattice model to BDSIM facilitated several studies beginning with energy loss and deposition studies to inform and validate the geometry of elements such as the beam pipe aperture and the collimators. A final beam pipe aperture of 3.65 cm was chosen as it introduced minimal losses whilst preventing the re-entry of scattered particles with a minimal impact on the overall lattice optics. The introduction

of collimators followed the same criteria of conforming to the intended optics whilst also only removing the intended particles. All these losses were tested for significance to the beam produced for the end station by modelling of an end station in BDSIM for the purpose of dose rate calculation. The element of major significance to the calculated dose rate was the second collimator which when introduced produced a dose rate change of $-2.3\,\mathrm{Gy\,s^{-1}}$ with respect to the reference dose (no collimators, 10 cm aperture) which further motivates LhARA's aim of minimised beam sizes, which would experience even less loss across the lattice and especially in that final collimator and produce higher dose rates overall.

An octupole placed before the bending arc was successfully able to generate uniform beams over a circle with radius 30 mm out of the arc at a strength of $k_3 = 30\,000\,\mathrm{mm^{-4}}$. A methodology was developed for quantifying the bunch smoothness by using correlation coefficients on a QQ plot of the bunch as it left the arc, where the bunch quantiles were compared to a theoretical uniform distribution. A benefit of this approach is that it should be possible with real beam snapshots as well. This method also serves as a useful way to validate shot-to-shot stability from the stage one LhARA beamline in simulations of mis-steered seeds. Future work should focus on minimising the tailedness of the bunch from the high-order focusing effects, so that a greater fraction of particles occupy the central smooth region.

The bunching capabilities of the two RF cavities was optimised and evaluated for both protons and carbon ions. Frequencies of both 352 MHz and 201 MHz were investigated, with the 201 MHz design recommended as the best option. To design the cavity geometry, a variety of optimisation algorithms were employed to converge on the best configuration. A custom longitudinal phase-space simulation code was written to evaluate the bunching of each design. This code was later compared to the more sophisticated BDSIM simulation, with good consistency. The optimally bunched proton beam achieved a bunch length of 1.18 ns, compared to the initial 4 ns, and the energy spread was reduced from 2% to 0.3%. The bunching for carbon ions gave a bunch length of 13.16 ns, compared to the initial 15.5 ns, while the energy spread was reduced from 2% to 1.6%. A 3D model and simulation was created using CST, and the resulting 3D field map was used to augment the lattice simulation in BDSIM.

The study into magnet designs has proved the feasibility of the lattice design optimised in the lattice optimisation study. Designs for a dipole with strength 0.551 T met magnetic field purity requirements, and designs for a quadrupole and octupole with magnetic strength parameters $k_1 = 36.96\,\mathrm{m}^{-2}$ and $k_3 = 60.273\,\mathrm{m}^{-4}$ respectively were above 99.8% field purity - close to the required 99.9%. Studies into a potential permanent magnet quadrupole for initial beam capture and focusing resulted in a samarium-cobalt design with a field gradient of $340\,\mathrm{T\,m}^{-1}$. Further research must be done into coil design for the electromagnetic lattice elements, ensuring reasonable currents densities and coil geometries, Additional research should also be done into improvements in field purity, where the magnet pole tips could be optimised for better field purity, potentially with the use of shims or similar.

Bibliography

[1] Stephen J McMahon. "Proton RBE models: commonalities and differences". en. In: *Phys. Med. Biol.* 66.4 (Feb. 2021), 04NT02.

[2] Harald Paganetti, Andrzej Niemierko, Marek Ancukiewicz, et al. "Relative biological effectiveness (RBE) values for proton beam therapy". en. In: *Int. J. Radiat. Oncol. Biol. Phys.* 53.2 (June 2002), pp. 407–421.

[3] Bleddyn Jones. "Why RBE must be a variable and not a constant in proton therapy". In: *British Journal of Radiology* 89.1063 (May 2016), p. 20160116. ISSN: 0007-1285. DOI: 10.1259/bjr.20160116. eprint: https://academic.oup.com/bjr/article-pdf/89/1063/20160116/54459260/bjr.20160116.pdf. URL: https://doi.org/10.1259/bjr.20160116.

[4] Henning Willers, Antino Allen, David Grosshans, et al. "Toward A variable RBE for proton beam therapy". en. In: *Radiother. Oncol.* 128.1 (July 2018), pp. 68–75.

[5] Tracy S A Underwood, Aimee L McNamara, Ane Appelt, et al. "A systematic review of clinical studies on variable proton Relative Biological Effectiveness (RBE)". en. In: *Radiother. Oncol.* 175 (Oct. 2022), pp. 79–92.

[6] Osama Mohamad, Brock J Sishc, Janapriya Saha, et al. "Carbon ion radiotherapy: A review of clinical experiences and preclinical research, with an emphasis on DNA damage/repair". en. In: *Cancers (Basel)* 9.6 (June 2017).

[7] Ann A Lazar, Reinhard Schulte, Bruce Faddegon, et al. "Clinical trials involving carbon-ion radiation therapy and the path forward". en. In: *Cancer* 124.23 (Dec. 2018), pp. 4467–4476.

[8] Tadashi Kamada, Hirohiko Tsujii, Eleanor A Blakely, et al. "Carbon ion radiotherapy in Japan: an assessment of 20 years of clinical experience". en. In: *Lancet Oncol.* 16.2 (Feb. 2015), e93–e100.

[9] Roberto Pacelli and Luigi Mansi. "Eric Hall and Amato J. Giaccia: Radiobiology for the radiologist, 6th edn." In: *European Journal of Nuclear Medicine and Molecular Imaging* 34 (Mar. 2007), pp. 965–966. DOI: 10.1007/s00259-007-0383-8.

[10] Thomas Friedrich, Uwe Scholz, Thilo Elsässer, et al. "Calculation of the biological effects of ion beams based on the microscopic spatial damage distribution pattern". en. In: *Int. J. Radiat. Biol.* 88.1-2 (Jan. 2012), pp. 103–107.

[11] Pierre Montay-Gruel, Kristoffer Petersson, Maud Jaccard, et al. "Irradiation in a flash: Unique sparing of memory in mice after whole brain irradiation with dose rates above 100Gy/s." In: *Radiotherapy and oncology : journal of the European Society for Therapeutic Radiology and Oncology* 124 3 (2017), pp. 365–369. URL: https://api.semanticscholar.org/CorpusID:12880630.

[12] Marie-Catherine Vozenin, Pauline De Fornel, Kristoffer Petersson, et al. "The advantage of FLASH radiotherapy confirmed in mini-pig and cat-cancer patients". en. In: *Clin. Cancer Res.* 25.1 (Jan. 2019), pp. 35–42.

[13] Jean Bourhis, Wendy Jeanneret Sozzi, Patrik Gonçalves Jorge, et al. "Treatment of a first patient with FLASH-radiotherapy". en. In: *Radiother. Oncol.* 139 (Oct. 2019), pp. 18–22.

[14] Kenneth W. D. Ledingham, Paul Robert Bolton, Naoya Shikazono, et al. "Towards Laser Driven Hadron Cancer Radiotherapy: A Review of Progress". In: *Applied Sciences* 4 (2014), pp. 402–443. URL: https://api.semanticscholar.org/CorpusID:8977517.

[15] Hiroyuki Daido, Mamiko Nishiuchi, and Alexander S Pirozhkov. "Review of laser-driven ion sources and their applications". en. In: *Rep. Prog. Phys.* 75.5 (May 2012), p. 056401.

[16] S. M. Wiggins, M. Boyd, E. Brunetti, et al. "Application programmes at the Scottish Centre for the Application of Plasma-based Accelerators (SCAPA)". In: *Relativistic Plasma Waves and Particle Beams as Coherent and Incoherent Radiation Sources III*. Ed. by Dino A. Jaroszynski and MinSup Hur. Vol. 11036. International Society for Optics and

Photonics. SPIE, 2019, 110360T. DOI: 10.1117/12.2520717. URL: https://doi.org/10.1117/12.2520717.

[17] Giuseppe A. P. Cirrone, Daniele Margarone, Mario Maggiore, et al. "ELIMED: a new hadron therapy concept based on laser driven ion beams". In: *Laser Acceleration of Electrons, Protons, and Ions II; and Medical Applications of Laser-Generated Beams of Particles II; and Harnessing Relativistic Plasma Waves III*. Ed. by Eric Esarey, Carl B. Schroeder, Wim P. Leemans, et al. Vol. 8779. International Society for Optics and Photonics. SPIE, 2013, p. 87791I. DOI: 10.1117/12.2026530. URL: https://doi.org/10.1117/12.2026530.

[18] F. Romano, F. Schillaci, G.A.P. Cirrone, et al. "The ELIMED transport and dosimetry beamline for laser-driven ion beams". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829 (2016). 2nd European Advanced Accelerator Concepts Workshop - EAAC 2015, pp. 153–158. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2016.01.064. URL: https://www.sciencedirect.com/science/article/pii/S0168900216000929.

[19] Daniele Margarone, G. A. Pablo Cirrone, Giacomo Cuttone, et al. "ELIMAIA: A Laser-Driven Ion Accelerator for Multidisciplinary Applications". In: *Quantum Beam Science* 2.2 (2018). ISSN: 2412-382X. DOI: 10.3390/qubs2020008. URL: https://www.mdpi.com/2412-382X/2/2/8.

[20] George K Holt, Giorgio Battaglia, Enrico Brunetti, et al. "Progress Towards Laser Wakefield Acceleration and Applications at the Scottish Centre for the Application of Plasma-based Accelerators (SCAPA)". In: *Journal of Physics: Conference Series* 1596 (2020). URL: https://api.semanticscholar.org/CorpusID:224484855.

[21] Pankaj Chaudhary, Giuliana Milluzzo, Hamad Ahmed, et al. "Radiobiology Experiments With Ultra-high Dose Rate Laser-Driven Protons: Methodology and State-of-the-Art". In: *Frontiers in Physics* 9 (2021). ISSN: 2296-424X. DOI: 10.3389/fphy.2021.624963. URL: https://www.frontiersin.org/articles/10.3389/fphy.2021.624963.

[22] Dennis Gabor. "A Space-Charge Lens for the Focusing of Ion Beams". In: *Nature* 160 (1947), pp. 89–90. URL: https://api.semanticscholar.org/CorpusID:4033128.

[23] J. Pozimski and M. Aslaninejad. "Gabor lenses for capture and energy selection of laser driven ion beams in cancer treatment". In: *Laser and Particle Beams* 31.4 (2013), 723–733. DOI: 10.1017/S0263034613000761.

[24] Galen Aymar, Tobias Becker, Stewart Boogert, et al. "LhARA: The Laser-hybrid Accelerator for Radiobiological Applications". English. In: *Frontiers in Physics* 8 (Sept. 2020), pp. 1–21. ISSN: 2296-424X. DOI: 10.3389/fphy.2020.567738.

[25] H. Grote and F. Schmidt. "Mad-X - an upgrade from mad8". In: *Proceedings of the 2003 Particle Accelerator Conference*. Vol. 5. 2003, pp. 3497–3499. DOI: 10.1109/PAC.2003.1289960.

[26] L.J. Nevay, S.T. Boogert, J. Snuverink, et al. "BDSIM: An accelerator tracking code with particle–matter interactions". In: *Computer Physics Communications* 252 (2020), p. 107200. ISSN: 0010-4655. DOI: https://doi.org/10.1016/j.cpc.2020.107200. URL: https://www.sciencedirect.com/science/article/pii/S0010465520300400.

[27] L. Nevay, A. Abramov, S. Alden, et al. *pybdsim: A Python package for Beam Delivery Simulation*. https://www.pp.rhul.ac.uk/bdsim/pybdsim/. Accessed: 2024-04-05. 2024.

[28] Philip F. Meads. "A Nonlinear Lens System to Smooth the Intensity Distribution of a Gaussian Beam". In: *IEEE Transactions On Nuclear Science* 30 (1983). ISSN: 1558-1578. DOI: 10.1109/TNS.1983.4332972. URL: https://ieeexplore.ieee.org/document/4332972.

[29] M. J. D. Powell. "A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation". In: *Advances in Optimization and Numerical Analysis*. Ed. by Susana Gomez and Jean-Pierre Hennart. Dordrecht: Springer Netherlands, 1994, pp. 51–67. ISBN: 978-94-015-8330-5. DOI: 10.1007/978-94-015-8330-5_4. URL: https://doi.org/10.1007/978-94-015-8330-5_4.

[30] Dassault Systemes. *OPERA: Operating environment for Electromagnetic Research and Analysis*. 2020. URL: https://www.3ds.com/products/simulia/opera.

[31] Meeker, David. *Finite Element Method Magnetics*. Version 4.2. Feb. 28, 2018. URL: https://www.femm.info.

[32] Peter P. Silvester and Ronald L. Ferrari. *Finite Elements for Electrical Engineers*. 2nd ed. Cambridge University Press, 1990.

[33] Clive A.J. Fletcher. *Computational Techniques for Fluid Dynamics*. Springer-Verlag, 1988. DOI: https://doi.org/10.1007/978-3-642-97071-9.

[34] M T Menzel and H K Stokes. *User's guide for the POISSON/SUPERFISH group of codes*. Tech. rep. Los Alamos, NM: Los Alamos Nat. Lab., 1987. URL: https://cds.cern.ch/record/801133.

[35] Andrzej Wolski. *Beam Dynamics in High Energy Particle Accelerators*. 1st ed. IMPERIAL COLLEGE PRESS, 2014. DOI: 10.1142/p899. eprint: https://www.worldscientific.com/doi/pdf/10.1142/p899. URL: https://www.worldscientific.com/doi/abs/10.1142/p899.

Appendix I

Longitudinal Phase-Space Evolution

Figures A1.1, A1.2, A1.3, A1.4 and A1.5 display the phase-space evolution as a simulated beam passes through LhARA stage 1. This example simulation is for the optimised 352 MHz cavity, with the cavity field amplitudes optimised to give maximum bunching. The green lines in the side-plots represent 1σ in the energy spread and the bunch length respectively.

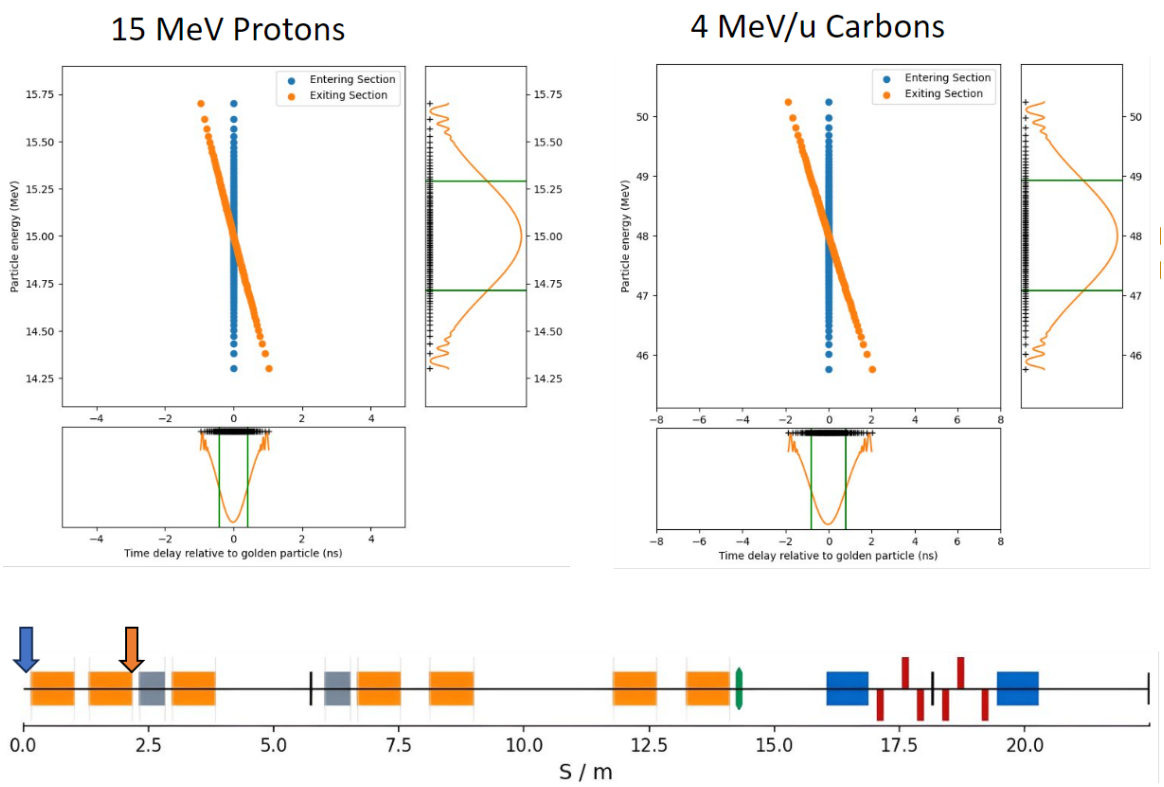

Figure A1.1: Longitudinal phase-space of particles passing through the first lattice section.

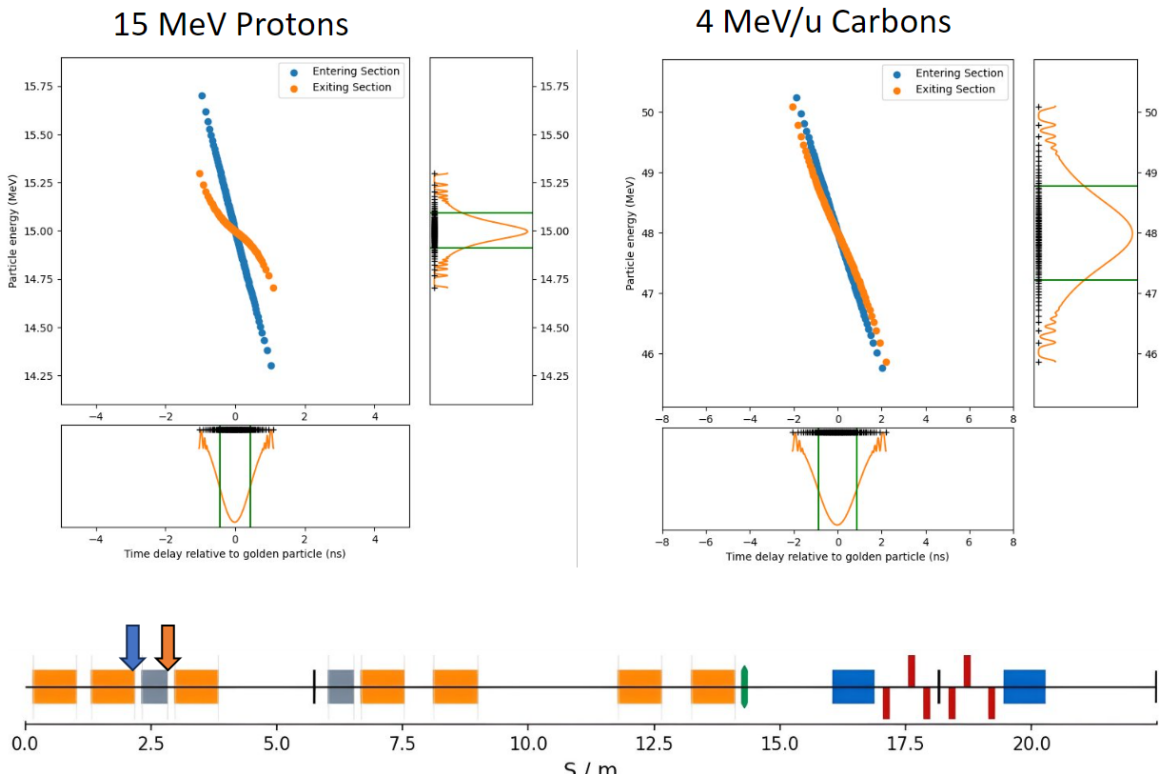

Figure A1.2: Longitudinal phase-space of particles passing through the first bunching cavity.

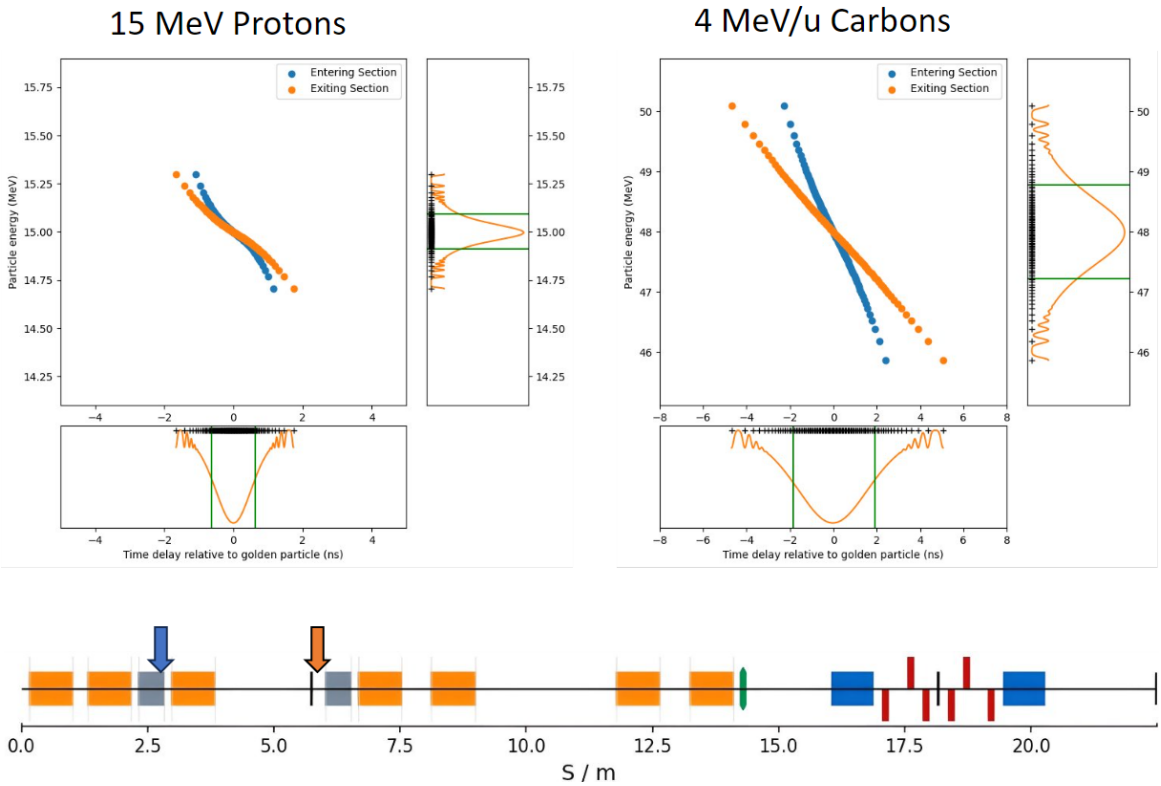

Figure A1.3: Longitudinal phase-space of particles passing through the second lattice section.

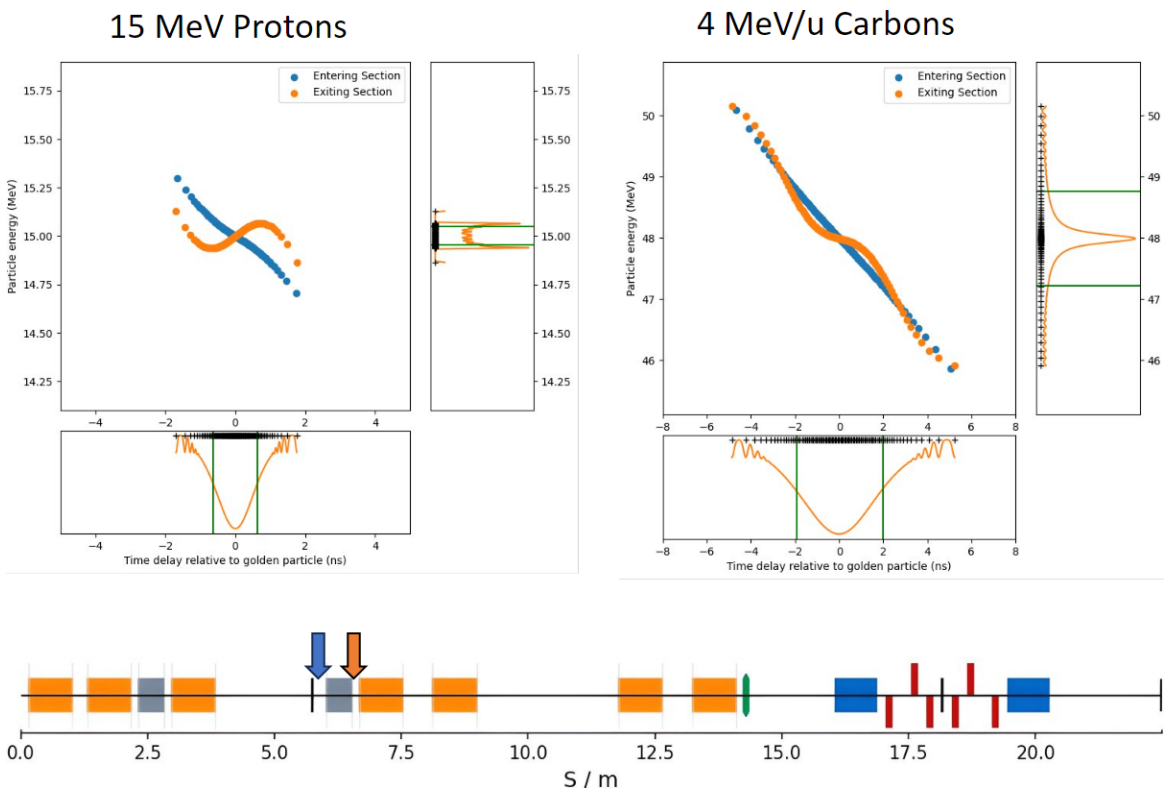

Figure A1.4: Longitudinal phase-space of particles passing through the second bunching cavity.

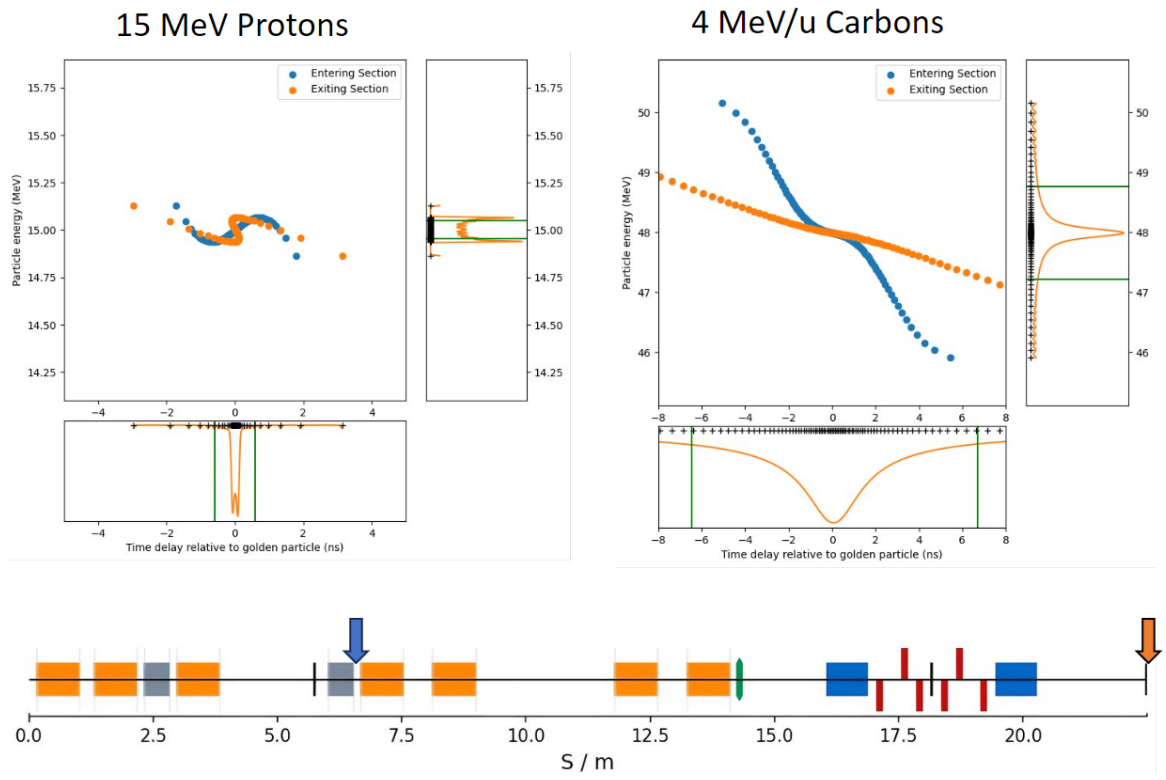

Figure A1.5: Longitudinal phase-space of particles passing through the third lattice section reaching the end station.

Appendix II

List of Magnet Harmonics

Appendix II shows the full list of magnet harmonics for the designed dipole. quadrupole and octupole, along with their contributions to the B-field at $r = r_{\mathrm{GFR}}$. All B-field contributions calculated from the larger of the two harmonics (normal or skew).

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.551\,\mathrm{m}^{-1}$	0.551 T
Quadrupole (k_1)	$1.0 \times 10^{-5}\,\mathrm{m}^{-2}$	$0.0\,\mathrm{m}^{-2}$	0.160 µT
Sextupole (k_2)	$1.2 \times 10^{-4}\,\mathrm{m}^{-3}$	$1.0 \times 10^{-5}\,\mathrm{m}^{-3}$	0.041 µT
Octupole (k_3)	$4.5 \times 10^{-3}\,\mathrm{m}^{-4}$	$2.9 \times 10^{-4}\,\mathrm{m}^{-4}$	0.018 µT
Decapole (k_4)	$0.291\,\mathrm{m}^{-5}$	$0.0146\,\mathrm{m}^{-5}$	0.010 µT
Dodecapole (k_5)	$26.58\,\mathrm{m}^{-6}$	$1.065\,\mathrm{m}^{-6}$	0.007 µT
14-pole (k_6)	$3165\,\mathrm{m}^{-7}$	$107.9\,\mathrm{m}^{-7}$	0.005 µT
16-pole (k_7)	$4.7 \times 10^{5}\,\mathrm{m}^{-8}$	$1.3 \times 10^{4}\,\mathrm{m}^{-8}$	0.003 µT
18-pole (k_8)	$8.1 \times 10^{7}\,\mathrm{m}^{-9}$	$2.0 \times 10^{6}\,\mathrm{m}^{-9}$	0.003 µT
20-pole (k_9)	$1.7 \times 10^{10}\,\mathrm{m}^{-10}$	$3.6 \times 10^{8}\,\mathrm{m}^{-10}$	0.002 µT
22-pole (k_{10})	$3.8 \times 10^{12}\,\mathrm{m}^{-11}$	$7.3 \times 10^{10}\,\mathrm{m}^{-11}$	0.002 µT
24-pole (k_{11})	$9.9 \times 10^{14}\,\mathrm{m}^{-12}$	$1.7 \times 10^{13}\,\mathrm{m}^{-12}$	0.001 µT
26-pole (k_{12})	$2.9 \times 10^{17}\,\mathrm{m}^{-13}$	$4.2 \times 10^{15}\,\mathrm{m}^{-13}$	0.001 µT
28-pole (k_{13})	$9.1 \times 10^{19}\,\mathrm{m}^{-14}$	$1.2 \times 10^{18}\,\mathrm{m}^{-14}$	0.0 T
30-pole (k_{14})	$3.1 \times 10^{22}\,\mathrm{m}^{-15}$	$3.6 \times 10^{20}\,\mathrm{m}^{-15}$	0.0 T
32-pole (k_{15})	$1.2 \times 10^{25}\,\mathrm{m}^{-16}$	$1.2 \times 10^{23}\,\mathrm{m}^{-16}$	0.0 T
34-pole (k_{16})	$4.7 \times 10^{27}\,\mathrm{m}^{-17}$	$3.6 \times 10^{25}\,\mathrm{m}^{-17}$	0.0 T
36-pole (k_{17})	$2.0 \times 10^{30}\,\mathrm{m}^{-18}$	$1.5 \times 10^{28}\,\mathrm{m}^{-18}$	0.0 T
38-pole (k_{18})	$9.2 \times 10^{32}\,\mathrm{m}^{-19}$	$5.8 \times 10^{30}\,\mathrm{m}^{-19}$	0.0 T
40-pole (k_{19})	$4.5 \times 10^{35}\,\mathrm{m}^{-20}$	$2.1 \times 10^{33}\,\mathrm{m}^{-20}$	0.0 T

Table A2.1: Full table of dipole higher-order harmonics up to 40-pole.

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.0\,\mathrm{m}^{-1}$	0.0 T
Quadrupole (k_1)	$36.96\,\mathrm{m}^{-2}$	$0.058\,\mathrm{m}^{-2}$	0.726 T
Sextupole (k_2)	$0.0\,\mathrm{m}^{-3}$	$0.0\,\mathrm{m}^{-3}$	0.0 T
Octupole (k_3)	$0.0\,\mathrm{m}^{-4}$	$0.00\,\mathrm{m}^{-4}$	0.0 T
Decapole (k_4)	$0.0\,\mathrm{m}^{-5}$	$0.0\,\mathrm{m}^{-5}$	0.0 T
Dodecapole (k_5)	$3.7 \times 10^{6}\,\mathrm{m}^{-6}$	$15\,461\,\mathrm{m}^{-6}$	0.908 mT
14-pole (k_6)	$0.0\,\mathrm{m}^{-7}$	$0.0\,\mathrm{m}^{-7}$	0.0 T
16-pole (k_7)	$0.0\,\mathrm{m}^{-8}$	$0.0\,\mathrm{m}^{-8}$	0.0 T
18-pole (k_8)	$0.017\,\mathrm{m}^{-9}$	$0.038\,\mathrm{m}^{-9}$	0.0 T
20-pole (k_9)	$1.4 \times 10^{15}\,\mathrm{m}^{-10}$	$1.7 \times 10^{13}\,\mathrm{m}^{-10}$	0.174 mT
22-pole (k_{10})	$3843\,\mathrm{m}^{-11}$	$2316\,\mathrm{m}^{-11}$	0.0 T
24-pole (k_{11})	$3.7 \times 10^{6}\,\mathrm{m}^{-12}$	$9.8 \times 10^{5}\,\mathrm{m}^{-12}$	0.0 T
26-pole (k_{12})	$5.1 \times 10^{8}\,\mathrm{m}^{-13}$	$5.2 \times 10^{8}\,\mathrm{m}^{-13}$	0.0 T
28-pole (k_{13})	$3.7 \times 10^{24}\,\mathrm{m}^{-14}$	$5.1 \times 10^{22}\,\mathrm{m}^{-14}$	0.039 mT
30-pole (k_{14})	$1.3 \times 10^{13}\,\mathrm{m}^{-15}$	$1.2 \times 10^{13}\,\mathrm{m}^{-15}$	0.0 T
32-pole (k_{15})	$3.7 \times 10^{15}\,\mathrm{m}^{-16}$	$3.8 \times 10^{16}\,\mathrm{m}^{-16}$	0.0 T
34-pole (k_{16})	$3.2 \times 10^{19}\,\mathrm{m}^{-17}$	$1.6 \times 10^{19}\,\mathrm{m}^{-17}$	0.0 T
36-pole (k_{17})	$3.3 \times 10^{35}\,\mathrm{m}^{-18}$	$1.4 \times 10^{33}\,\mathrm{m}^{-18}$	0.092 mT
38-pole (k_{18})	$3.8 \times 10^{23}\,\mathrm{m}^{-19}$	$2.6 \times 10^{24}\,\mathrm{m}^{-19}$	0.0 T
40-pole (k_{19})	$6.5 \times 10^{26}\,\mathrm{m}^{-20}$	$5.2 \times 10^{26}\,\mathrm{m}^{-20}$	0.0 T

Table A2.2: Full table of quadrupole higher-order harmonics up to 40-pole.

Harmonic	k-value (normal)	k-value (skew)	$B(r = r_{\mathrm{GFR}})$
Dipole (k_0)	$0.0\,\mathrm{m}^{-1}$	$0.0\,\mathrm{m}^{-1}$	0.0 T
Quadrupole (k_1)	$0.0\,\mathrm{m}^{-2}$	$0.0\,\mathrm{m}^{-2}$	0.0 T
Sextupole (k_2)	$0.0\,\mathrm{m}^{-3}$	$0.0\,\mathrm{m}^{-3}$	0.0 T
Octupole (k_3)	$60\,273\,\mathrm{m}^{-4}$	$0.167\,\mathrm{m}^{-4}$	0.242 T
Decapole (k_4)	$0.0\,\mathrm{m}^{-5}$	$0.0\,\mathrm{m}^{-5}$	0.0 T
Dodecapole (k_5)	$0.0\,\mathrm{m}^{-6}$	$0.0\,\mathrm{m}^{-6}$	0.0 T
14-pole (k_6)	$0.0\,\mathrm{m}^{-7}$	$0.0\,\mathrm{m}^{-7}$	0.0 T
16-pole (k_7)	$5.7 \times 10^{8}\,\mathrm{m}^{-8}$	$4.6 \times 10^{8}\,\mathrm{m}^{-8}$	4.082 µT
18-pole (k_8)	$0.0\,\mathrm{m}^{-9}$	$0.0\,\mathrm{m}^{-9}$	0.0 T
20-pole (k_9)	$0.0\,\mathrm{m}^{-10}$	$0.0\,\mathrm{m}^{-10}$	0.0 T
22-pole (k_{10})	$0.0\,\mathrm{m}^{-11}$	$0.0\,\mathrm{m}^{-11}$	0.0 T
24-pole (k_{11})	$3.9 \times 10^{20}\,\mathrm{m}^{-12}$	$2.4 \times 10^{17}\,\mathrm{m}^{-12}$	0.527 mT
26-pole (k_{12})	$0.0\,\mathrm{m}^{-13}$	$0.0\,\mathrm{m}^{-13}$	0.0 T
28-pole (k_{13})	$0.0\,\mathrm{m}^{-14}$	$0.0\,\mathrm{m}^{-14}$	0.0 T
30-pole (k_{14})	$0.0\,\mathrm{m}^{-15}$	$0.0\,\mathrm{m}^{-15}$	0.0 T
32-pole (k_{15})	$9.2 \times 10^{26}\,\mathrm{m}^{-16}$	$1.2 \times 10^{28}\,\mathrm{m}^{-16}$	746 µT
34-pole (k_{16})	$0.0\,\mathrm{m}^{-17}$	$0.0\,\mathrm{m}^{-17}$	0.0 T
36-pole (k_{17})	$0.0\,\mathrm{m}^{-18}$	$0.0\,\mathrm{m}^{-18}$	0.0 T
38-pole (k_{18})	$0.0\,\mathrm{m}^{-19}$	$0.0\,\mathrm{m}^{-19}$	0.0 T
40-pole (k_{19})	$2.4 \times 10^{40}\,\mathrm{m}^{-20}$	$7.9 \times 10^{38}\,\mathrm{m}^{-20}$	0.024 mT

Table A2.3: Full table of octupole higher-order harmonics up to 40-pole.